COMPUTATIONAL ASPECTS OF OPTIMAL INFORMATION REVELATION

by

Yu Cheng

A Dissertation Presented to the
FACULTY OF THE USC GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(COMPUTER SCIENCE)

July 2017

Acknowledgments

I first want to thank my Ph.D. advisor Shang-Hua Teng. A few months before I became Shang-Hua's student, I was at his talk on Laplacian Paradigm at Tsinghua University. The elegance of the materials, together with his research vision, motivated me to pursue graduate study in theoretical computer science. I am grateful to him for seeing the potential in me, and taking me in as his student when I was a teaching assistant at University of Southern California (USC). During my Ph.D., Shang-Hua has been a source of constant encouragement, support, and guidance, for both research and life. I thank him for his patience, and for the flexibility he offers me to work on anything that excites me. Shang-Hua is happy to let me take the wheel, but he is also always there when I need his advice. He often embeds his philosophies in his stories — stories that I understand better as time goes by.

I would like to thank Shaddin Dughmi and David Kempe for many enjoyable discussions, and for sharing their knowledge and experience with me. I am also grateful for their guidance on writing and presentations. I thank Shaddin for introducing me to mechanism design and convex optimization, which eventually lead to the works in this thesis. I thank David for teaching me randomized algorithms, especially the principle of deferred decisions; for his advice on running and hiking, and for his invaluable help in preparing this thesis.

I would also like to thank the rest of my thesis committee, Yan Liu and Ben Reichardt, for providing feedback to my thesis and giving me advice on career paths.

The bulk of this thesis occurred in 2014–2015 at USC. Shaddin Dughmi, in particular, made many invaluable contributions to the results in this thesis, including but not limited to, the mixture selection framework. I would like to thank Umang Bhaskar and Chaitanya Swamy for our long-distance collaboration, and for their contributions to the results on signaling in routing games.

I would like to thank Ilias Diakonikolas for introducing me to the world of learning and testing distributions, and Alistair Stewart for teaching me how to play with the sum of independent random vectors. I thank Ilias and Alistair for their contribution to the results on computing Nash equilibria in anonymous games.

I thank Ho Yee Cheung, Ehsan Emamjomeh-Zadeh, and Li Han for many casual discussions on theorems, proofs, and puzzles; and for their contributions to the results on mixture selection and persuading voters presented in this thesis.

I am very fortunate to have the opportunities to learn from Xi Chen. I thank him for introducing me to equilibrium computation and the complexity class PPAD. I wish I could be as calm as him, and I deeply treasure the times when we worked on “random” open questions at the Simons Institute for the Theory of Computing.

I would like to thank Richard Peng for our discussions on faster algorithms for large-scale graph and matrix problems. I thank Dehua Cheng for showing me the power of singular value decomposition. I would also like to thank the rest of my collaborators during my graduate study, from whom I learned so much: Wade Hann-Caruthers, Daniel Kane, Robert Kleinberg, Young Kun Ko, Abhishek Samanta, Aaron Sidford, Xiaorui Sun, Ravi Sundaram, Omer Tamuz, Bo Tang, Adrian Vladu, Di Wang, and Haifeng Xu.

I am grateful to Leana Golubchik for her advice and support, particularly during the first and last year of my Ph.D. I would like to thank Leonard Adleman for sharing his stories and being a constant source of inspiration; I admire Len's passion for mathematics and boxing.

I thank my fellow students in the theory group at USC: Brendan Avent, Joseph Bebel, Hsing-Hau Chen, Xinran He, Ho Yee Cheung, Ehsan Emamjomeh-Zadeh, Lian Liu, Anand Kumar Narayanan, Ruixin Qiang, Alana Shine, and Haifeng Xu. All of them together made Theoroom a wonderful place to be in during my time.

I am grateful to Vincent Conitzer, Rong Ge, Kamesh Munagala, and Debmalya Panigrahi for offering me a postdoctoral researcher position at Duke University. I thank the Department of Computer Science at USC for partially supporting my graduate studies through teaching assistantships; and I thank the wonderful staff in the department, especially Lizsl De Leon, Lifeng (Mai) Lee and Kusum Shori, for always being so friendly and helpful to me and many other students.

I would like to thank Chin-Yew Lin, Vahab Mirrokni, Konstantin Voevodski, and Wen Xu, for offering me the opportunity to experience industry and work with real data, and for their help and guidance during my internships. I thank Yong Yu for building up the ACM Class at Shanghai Jiao Tong University, which has enabled a generation of young people, including me, to pursue their dreams.

I am indebted to my parents, Ningqiu Cheng and Xiaoli Lyu, for encouraging me to learn more about maths and programming at an early age. I thank them for allowing me to play computer games as much as I want as a kid, for respecting my life choices, and for their support and unconditional love.

Abstract

Strategic interactions often take place in environments rife with uncertainty and information asymmetry. Understanding the role of information in strategic interactions is becoming more and more important in the age of information we live in today. This dissertation is motivated by the following question: What is the optimal way to reveal information, and how hard is it computationally to find an optimum? We study the optimization problem faced by an informed principal, who must choose how to reveal information in order to induce a desirable equilibrium, a task often referred to as *information structure design*, *signaling* or *persuasion*.

Our exploration of optimal signaling begins with *Bayesian network routing games*. This widely studied class of games arises in several real-world settings. For example, millions of people use navigation services like Google Maps every day. Is it possible for Google Maps (the principal) to partly reveal the traffic conditions to reduce the latency experienced by selfish drivers? We show that the answer to this question is two-fold: (1) There are scenarios where the principal can improve selfish routing, and sometimes through the careful provision of information, the principal can achieve the best-coordinated outcome; (2) Optimal signaling is computationally hard in routing games. Assuming P $\neq$ NP, there is no polynomial-time algorithm that does better than full revelation in the worst case.

We next study the optimal signaling problem in one of the most fundamental classes of games: *Bayesian normal form games.* We settle the complexity of (approximately) optimal signaling in normal form games: We give the first quasipolynomial time approximation scheme for signaling in normal form games; and complementing this, we show that a fully polynomial time approximation scheme for optimal signaling is NP-hard, and rule out any polynomial time approximation scheme assuming the planted clique conjecture. It is worth noting that our algorithm works for games with a constant number of players, and for a large and natural class of objective functions including social welfare, while our hardness results hold even in the simplest Bayesian two-player zero-sum games.

Complementing our results for signaling in normal form games, we continue to investigate the optimal signaling problem in two special cases of *succinct games*: (1) Second-price auctions in which the auctioneer wants to maximize revenue by revealing partial information about the item for sale to the bidders before running the auction; and (2) Majority voting when the voters have uncertainty regarding their utilities for the two possible outcomes, and the principal seeks to influence the outcome of the election by signaling. We give efficient approximation schemes for all these problems under one unified algorithmic framework, by identifying and solving a common optimization problem that lies at the core of all these applications. Finally, we present the currently best algorithm (asymptotically) for computing Nash equilibria in complete-information anonymous games. Compared to all other games we study in this thesis, anonymous games are the only class of games whose complexity of equilibrium computation is still open. We present the currently best algorithm for computing Nash equilibria in anonymous games, and we also provide some evidence suggesting our algorithm is essentially tight.

Contents

Chapter 1

Introduction

1.1 Information Structure Design

What is the best way to reveal information to other strategic players? This is a question we want to solve during a poker game at home, as well as in billion-dollar industries like online ad auctions. The strategic decisions of the players depend crucially on the information available to them, and the act of exploiting an informational advantage to influence the behavior of others is indeed universal.

In Bayesian games, players' payoffs often depend on the state of nature, which may be hidden from the players. Instead, players receive a *signal* regarding the state of nature which they use to form beliefs about their payoffs, and choose their strategies. Thus the strategic decisions and payoffs of the players depend crucially on the information available from the signal they receive.

In this thesis, we study the optimization problem faced by an informed principal, who must choose how to reveal information in order to induce a desirable equilibrium, a task often referred to as *information structure design*, *signaling* or *persuasion*. Similar to classic mechanism design[1], we have a principal who is interested in the outcome of the game, but the difference is that the principal influences

[1] Mechanism design is a field in economics and game theory that studies the design of mechanisms or incentives in strategic settings: A "principal" may choose the rules/structures of the game to induce a desirable outcome, given that other players act rationally.

the players by designing the information structure, rather than through designing the game structure. We focus on the design of *public* information structures that reveal the same information to all players. The study of *private* signaling schemes is interesting in its own right, but falls beyond the scope of this thesis.

Like mechanism design, the information structure design question is inherently algorithmic: How hard is it computationally to find the optimal information structure? In this thesis, we settle the computational complexity of optimal signaling in several fundamental game-theoretic settings: Bayesian normal form games, and Bayesian succinct games including network routing games, second-price auctions and voting with threshold rule.

1.2 Motivating Examples

To motivate the questions we investigate in this thesis, we first give three examples that are incomplete-information variants of classic examples studied in game theory. The first example, presented in [47], is a Bayesian network routing game adapted from Braess' paradox. The second example, presented in [46], is a Bayesian normal form game adapted from the prisoner's dilemma. The third example is a variation of a probabilistic single-item second-price auction in [52].

These three examples raise several interesting observations. First, designing the optimal signaling scheme is an important task, because revealing the right information can lead to much better results compared to trivial schemes like no revelation and full revelation. Second, as opposed to the "Market for Lemons" [2] example[2],

[2] Akerlof [2] uses the market for used cars as an example to illustrate that information deficiency can lead to worse outcomes. In his example, the market degrades in the presence of information

sometimes more information can also degrade the payoffs of all players and/or the principal, and the optimal information structure may reveal some but not all the information available. Third, they put forward the modeling question of how we should formulate the signaling problem, as well as the algorithmic question of how hard it is computationally to find the optimal signaling scheme.

1.2.1 An Informational Braess' Paradox

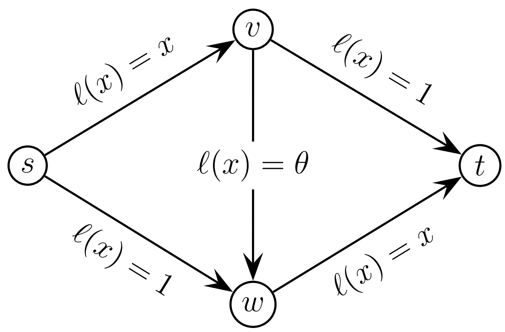

Figure 1.1: The informational Braess' paradox.

Consider the informational Braess' paradox given in Figure 1.1. It is a non-atomic Bayesian network routing game of incomplete information, a variant of the classic Braess' paradox (See, e.g. [19, 82, 85]). We call it informational Braess' paradox because in this example, somewhat counter-intuitively, all the drivers can do worse if they have more information about the uncertainty in the traffic conditions. In Figure 1.1, one unit of flow wants to travel from the source node s to the sink node t in the network, and the latency function $\ell(x)$ on each edge describes the delay experienced by drivers on that edge as a function of the fraction of overall

asymmetry between buyers and sellers, if the owner of the good used cars ("peaches") cannot distinguish himself from the owner of defective used cars ("lemons").

traffic using that road. The latency of the vertical edge is parameterized by a state of nature θ, which is drawn uniformly from $\{0, 1\}$.

When $\theta = 0$, this reverts back to the traditional Braess' paradox. Selfish drivers would all take the zig-zag path $s \to v \to w \to t$, where all drivers experience a latency of 2. If the drivers can cooperate rather than being selfish, the socially optimal routing is to send half of the drivers on the top path $s \to v \to t$, and the other half on the bottom path $s \to w \to t$. This way, all drivers experience a latency of 1.5, but then each individual would benefit from switching to the zig-zag path, and hence Braess' paradox occurs.

Now suppose that every morning nature flips a coin and decides the latency of the edge $v \to w$ to be either $\theta = 0$ or $\theta = 1$. The principal (e.g., navigation services like Google Maps) knows the exact value of θ, but the drivers only know the prior distribution of θ. Suppose that the principal wants to minimize the expected latency experienced by the selfish drivers at equilibrium, assuming that the agents are risk neutral[3].

Observe that all players would drive along the zig-zag path when $\mathbb{E}[\theta] \leq 0.5$. When $\mathbb{E}[\theta] \geq 0.5$, the delay on the edge $v \to w$ is enough to deter the drivers away from using the edge. Consider the following signaling schemes:

- Full information: When $\theta = 0$ the drivers experience latency 2, and when $\theta = 1$ the drivers experience latency 1.5. The expected latency is 1.75.

[3] The decision of a risk neutral player is not affected by the degree of uncertainty in a set of outcomes. A rational risk-neutral player always chooses the strategy with the highest expected payoff.

- No information (Optimal signaling scheme): Risk neutral players treat θ as its expectation $\mathbb{E}[\theta] = 0.5$, which is enough to incentivize the drivers to split and play the socially optimal solution, so the latency is always 1.5.

1.2.2 Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	$-1+\theta$ (upper right); $-1+\theta$ (lower left)	0 (upper right); $-5+\theta$ (lower left)
Defect	$-5+\theta$ (upper right); 0 (lower left)	-4 (upper right); -4 (lower left)

Figure 1.2: Prisoner's dilemma of Incomplete Information.

Consider a two-player game with incomplete information given by its normal-form representation in Figure 1.2, a variant of the prisoner's dilemma. Two members of a criminal gang are arrested and imprisoned. Without communicating with each other, they must choose to either cooperate with the other by remaining silent, or to defect by betraying the other and testifying. The two players (row player and column player) move simultaneously, and receive payoffs given in the cell specified by the combination of their actions. In each cell, the lower-left number represents the payoff of the row player, and the upper-right number represents the payoff of the column player. The payoffs are also parameterized by a state of nature θ, which affects players' payoffs together with their actions.

When $\theta = 0$, the game reverts back to the classic prisoner's dilemma, in which the socially optimal outcome is for both players to cooperate. However, the only

Nash equilibrium of the game is when both players defect, because it is strictly better for each player to defect, no matter what the other player does.

Suppose nature flips a coin and decides to give the prisoners θ extra reward for cooperating, where θ can be negative and is drawn uniformly from $\{0, 2\}$. The principal knows the exact value of θ, but both players only know the prior distribution. Suppose the principal cares about maximizing social welfare, which in this game is equivalent to maximizing the probability of both players cooperating, again assuming the players are risk neutral.

Observe that risk neutral players would play cooperate as an equilibrium if the expected value of θ was at least 1. Consider the following signaling schemes:

- Full information: Players would defect when $\theta = 0$. When $\theta = 2$, players would prefer to cooperate and therefore they cooperate 1/2 of the time.
- No information (Optimal signaling scheme): Both players treat θ as its expectation $\mathbb{E}[\theta] = 1$, which is just enough to incentivize them to cooperate. So by revealing no information, the principal gets the players to always cooperate and obtain the highest expected social welfare.

1.2.3 A Probabilistic Second-Price Auction

Consider a single-item second-price auction of a probabilistic good whose actual instantiation is known to the auctioneer but not to the bidders. This problem was considered in [52] and [21], and Example 1.3 is adapted from [21]. We use θ to denote the possible types of the probabilistic good. As shown in Figure 1.3, the item has four possible types and we have four bidders participating in the auction.

	Item type 1	Item type 2	Item type 3	Item type 4
Bidder 1	1	0	0	0
Bidder 2	0	1	0	0
Bidder 3	0	0	1	0
Bidder 4	0	0	0	1

Figure 1.3: An example of a probabilistic single-item auction.

Each bidder is only interested in one of the four possible item types with a valuation of 1, and not interested in any other types.

Assume that the auction format is a second-price auction, that is, the person who bids highest wins the item, but he needs to pay only the second highest bid. Since we are running a second-price auction, risk neutral players play the dominant-strategy truth-telling equilibrium, by bidding the expected value for the item. The principal (auctioneer) can choose to reveal some information about the actual realization θ, with the goal of maximizing her revenue.

Consider the following signaling schemes:

- Full information: If the principal reveals full information about θ, then only one bidder bids 1 and everyone else bids 0. Since the principal is running a second-price auction, the revenue of the auction is 0.
- No information: Without further information, the expected value of the item is 1/4 to all bidders, so everyone bids 1/4, and the revenue of the auction is 1/4.
- Optimal signaling scheme: One of the possible signaling schemes that maximizes revenue is to reveal whether $\theta \in \{1, 2\}$ or $\theta \in \{3, 4\}$. Upon learning

this information, two out of four bidders remain interested in the item, and their expected value for the item increases from 1/4 to 1/2, because the item now has only two possible types. The revenue of the auction is 1/2 since the principal receives two bids of 1/2 and two bids of 0.

1.3 A Frontier of Computational Game Theory

Information revelation has been widely studied by game theorists and economists, exploring how to reveal information strategically to other selfish agents (e.g., see [2, 5, 13, 14, 16, 64, 73, 77, 83]). However, most of the studies on information disclosure have been non-algorithmic, even though the demand for efficient algorithms has never been higher in the age of information we live in today.

Whereas understanding the role of information in influencing strategies is a classical problem in game theory, the computational problem of *designing* optimal information structures for Bayesian games, commonly called the *signaling problem*, has received mostly recent attention [7, 21, 46, 48, 49, 50, 52, 62]. These exciting developments during the past decade have brought new insights towards a much better understanding of the role of information through a computational lens.

Complexity theory, through concepts like NP-Completeness, aims to distinguish between the problems that admit efficient algorithms and those that are intractable. In this thesis, we focus on the computational aspects of information structure design. Our goals are to mathematically formulate the signaling problem, to develop algorithms and prove matching lower bounds, and to characterize the exact time complexity of optimal information revelation in different game-theoretic scenarios.

For designing efficient algorithms, we analyze the structural properties of the optimal signaling scheme, identify an optimization problem that arises naturally when one seeks to craft posterior beliefs, and develop a powerful algorithmic framework that solves the (approximately) optimal signaling problem in a variety of games. For deriving hardness results, our goal is to design the information asymmetry of the game to *encode computation* in the optimal information structure. For example, if we have an oracle for optimal signaling, can we use it to find maximum independent set of a graph, or to recover a hidden clique in a random graph? Examining the computational complexity of optimal signaling in different games helps us recognize the essence of these problems.

Besides improving our understanding about the role of information in game theory, the investigation of optimal information revelation has also led to powerful algorithmic frameworks and basic open questions. For example, in Chapter 5 we extract a common optimization problem that lies at the core of several signaling problems, identify two "smoothness" properties which seem to govern the complexity of near-optimal signaling schemes, and resolve a number of open problems under one algorithmic framework. Another example is in Section 4.4, where we utilize the equivalence between separation and optimization to show hardness of signaling in normal form games. However, the proof would be much more elegant if we had a better understanding of this equivalence in the approximate sense, in particular, if we could resolve Open Problem 7.2 in the positive.

1.4 Our Contributions and Thesis Organization

We study the optimal signaling problem in several fundamental game-theoretic settings: Bayesian network routing games, Bayesian normal form games, probabilistic second-price auctions, and majority voting with incomplete information. We focus on public signaling schemes where the principal reveals the same information to all players. Our main contribution is to derive efficient approximation algorithms, as well as hardness results for these classes of games that close the gap between what is achievable in polynomial time (or quasipolynomial time) and what is intractable.

We start from an example of using information to battle selfishness in routing games. We then continue to systematically study the signaling problem in both normal form games with a bounded number of players, as well as succinct games with many players.

Signaling in network routing games: In Chapter 3, we consider the signaling problem in (nonatomic, selfish) *Bayesian network routing games*, wherein the principal seeks to reveal partial information to *minimize* the average latency of the equilibrium flow. We show that it is NP-hard to obtain any multiplicative approximation better than $\frac{4}{3}$, even with linear latency functions (Theorem 3.2). This yields an *optimal* inapproximability result for linear latencies, since we show that full revelation obtains the *price of anarchy* of the routing game as its approximation ratio (Theorem 3.1), which is $\frac{4}{3}$ for linear latency functions [85]. These are the *first* results for the complexity of signaling in Bayesian network routing games.

Signaling in normal form games: In Chapter 4, we consider signaling in *Bayesian normal form games*, where the payoff entries are parametrized by a state of nature. Dughmi [46] initiated the computational study of this problem and obtained various hardness results. On the algorithms side, we give two different approaches (Theorem 4.1 and Theorem 5.12) for obtaining a bi-criteria QPTAS for normal-form games with a constant number of players, and for a large and natural class of objective functions like social welfare and weighted combination of players' utilities [26, 27]. In other words, we can in quasipolynomial time approximate the optimal reward from signaling while losing an additive ϵ in the objective as well as in the incentive constraints.

For hardness results, [46] considered the special case of signaling in *Bayesian (two-player) zero-sum games*, in which the principal seeks to maximize the equilibrium payoff of one of the players, and ruled out a fully polynomial time approximation scheme (FPTAS) for this problem assuming planted clique hardness. We show that it is NP-hard to obtain an additive FPTAS (Theorem 4.7), settling the complexity of the problem with respect to NP-hardness. Moreover, we show that assuming the planted clique conjecture (Conjecture 2.6), there does not exist a polynomial time approximation scheme (PTAS) for the signaling problem (Theorem 4.10).

Mixture selection framework In Chapter 5, we pose and study an algorithmic problem which we term *mixture selection*, a problem that arises naturally in the design of optimal information structures. The mixture selection problem is closely related to the optimal signaling problem. We identify two "smoothness" property of Bayesian games that seem to dictate the complexity of mixture selection and optimal signaling: Lipschitz continuity and a *noise stability* notion that we define. We present an algorithmic framework that (approximately) solves mixture selection (Theorem 5.6) and optimal signaling (Theorem 5.10) in a number of different

Bayesian games. The approximation guarantee of our algorithm degrades gracefully as a function of the two smoothness parameters, in particular, when the game is $O(1)$-Lipschitz continuous and $O(1)$-stable, we obtain an additive PTAS optimal signaling. We also show that neither assumption suffices by itself for a PTAS (Theorems 5.18 and 5.19). We give a new QPTAS for signaling in normal form game using our algorithmic framework (Theorem 5.12). Moreover, our algorithms for signaling in multi-player games also follow from the powerful mixture selection framework.

Signaling in anonymous games In Chapter 6, we consider signaling in *anonymous games*. In contrast to the normal form games we study in Chapter 4, anonymous games form an important class of *succinct games*, capturing a wide range of game-theoretic scenarios, including auctions and voting. We start with two special cases of anonymous games, both admitting a PTAS.

In Section 6.1, we consider signaling in the context of a probabilistic second-price auction. In this setting, the item being auctioned is probabilistic, and the instantiation of the item is known to the auctioneer but not to the bidders. The auctioneer must decide what information to reveal in order to maximize her revenue in this auction. Emek et al. [52] and Miltersen and Sheffet [21] considered several special cases of this problem and presented polynomial-time algorithms when bidder types are fixed. [52] showed that in the general setting, where the auctioneer holds probabilistic knowledge on the bidders' valuations, an FPTAS for signaling becomes NP-hard. We resolve the approximation complexity of optimal signaling in the Bayesian setting by giving an additive PTAS (Theorem 6.2).

In Section 6.2, we study the persuasion in voting problem proposed by Alonso and Câmara [5]. Consider a binary outcome election — say whether a ballot measure is passed — when voters are not fully informed of the consequences of the measure, and hence of their utilities. Each voter casts a Yes/No vote, and the measure passes if the fraction of Yes votes exceeds a certain pre-specified threshold. We consider a principal who has control over which information regarding the measure is gathered and shared with voters, and looks to maximize the probability of the measure passing. We present a multi-criteria PTAS for this problem (Theorem 6.5).

Section 6.3 takes a detour and studies anonymous games with complete information. We give the first polynomial time algorithm for computing Nash equilibria of inverse polynomial precision in anonymous games with more than two strategies (Theorem 6.6), and present evidence suggesting that our algorithm is essentially tight (Theorem 6.7). This gets us closer to pinning down the computational complexity of Nash equilibria in anonymous games, and the only question left is whether there is a FPTAS for computing equilibria or not.

Chapter 2

Background and Notation

We use $\mathbb{R}_+$ for the set of nonnegative reals. For an integer n, let $[n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$. If $n \geq 1$, we use Δ_n to denote the $(n-1)$-dimensional simplex $\{x \in \mathbb{R}_+^n : \sum_i x_i = 1\}$. We refer to a distribution $y \in \Delta_n$ as *s-uniform* if and only if it is the average of a multiset of s standard basis vectors in n-dimensional space. Let $\mathbb{1}_n \in \mathbb{R}^n$ be the vector with 1 in all its entries, $I_{n\times n}$ be the $n \times n$ identity matrix, and e_i be the i-th standard basis vector containing 1 as its i-th entry and 0 elsewhere.

We use "iff" to abbreviate "if and only if". For two functions $f(n)$ and $g(n)$, we write $f(n) = O(g(n))$ iff there exists constants C and n_0 such that $|f(x)| \leq C|g(x)|$ for all $n \geq n_0$; we write $g(n) = \Omega(f(n))$ iff $f(n) = O(g(n))$, and $g(n) = \Theta(f(n))$ iff $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$. We say $f(n) = o(g(n))$ iff $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0$, and $g(n) = \omega(f(n))$ iff $f(n) = o(g(n))$. We use $\text{poly}(n)$ to denote a polynomial function of n. When we say with high probability, we mean with probability at least $1 - \frac{1}{n^\alpha}$ for some constant $\alpha > 0$; and the parameter n will be clear from the context.

Let $|\mathcal{I}|$ denote the description size of the instance $\mathcal{I}$. An (additive) *polynomial time approximation scheme* (PTAS) is an algorithm that runs in time $\text{poly}(|\mathcal{I}|)$, and returns a solution of value at least $\text{OPT}(\mathcal{I}) - \epsilon$ for every instance $\mathcal{I}$ and constant $\epsilon > 0$. An *fully polynomial time approximation scheme* (FPTAS) is a PTAS whose running time for an instance $\mathcal{I}$ and parameter ϵ is $\text{poly}\left(|\mathcal{I}|, \frac{1}{\epsilon}\right)$. An *quasipolynomial time approximation scheme* (QPTAS) is an algorithm that runs in time $|\mathcal{I}|^{O(\log|\mathcal{I}|)}$ and returns an ϵ-optimal solution for every instance $\mathcal{I}$ and constant $\epsilon > 0$.

2.1 Bayesian Games and Signaling

A Bayesian game is a game in which the players have incomplete information on the payoffs of the game. In this thesis, we consider a number of Bayesian games in which payoffs are parametrized by θ, the *state of nature*. We use Θ to denote the set of all states of nature, and assume $\theta \in \Theta$ is drawn from a common-knowledge prior distribution which we denote by λ. We consider Bayesian games given in the explicit representation:

- An integer M denoting the number of states of nature. We index states of nature by the set $\Theta = [M] = \{1, \ldots, M\}$.
- A common-knowledge prior distribution $\lambda \in \Delta_M$ on the states of nature.
- A set of M games of complete information, one for each state of nature θ, describing the payoff structure of the game.

Note that a game of complete information is the special case with $M = 1$, i.e., the state of nature is fixed and known to all.

In all our applications, we assume players a priori know nothing about θ other than its prior distribution λ, and examine policies whereby a principal with access to the realized value of θ may commit to a policy revealing information to the players regarding θ. The goal of the principal is then to commit to revealing certain information about θ — i.e., a signaling scheme — to induce a favorable equilibrium over the resulting Bayesian subgames. This is often referred to as *signaling*, *persuasion*, or *information structure design*. The recent survey by Dughmi [47] contains a nice summary of the work on information structure design in the algorithmic game theory community.

2.1.1 Signaling Schemes

A *signaling scheme* is a policy by which a principal reveals (partial) information about the state of nature. We call a signaling scheme *public* if it reveals the same information to all the players, and *private* when different signals are sent to the players through private channels. In this thesis, we focus on the design of public signaling schemes.

Let $M = |\Theta|$. A signaling scheme specifies a set of signals Σ and a (possibly randomized) map $\varphi : \Theta \to \Delta_{|\Sigma|}$ from the states of nature Θ to distributions over the signals in Σ. Abusing notation, we use $\varphi(\theta, \sigma)$ to denote the probability of announcing signal $\sigma \in \Sigma$ conditioned on the state of nature being $\theta \in \Theta$.

Each signal σ yields a posterior distribution $\mu_\sigma \in \Delta_M$. It was observed by Kamenica and Gentzkow [69] that signaling schemes are in one-to-one correspondence with *convex decompositions* of the prior distribution $\lambda \in \Delta_M$: Formally, a signaling scheme $\varphi : \Theta \to \Sigma$ corresponds to the convex decomposition $\lambda = \sum_{\sigma \in \Sigma} p_\sigma \cdot \mu_\sigma$, where (1) $p_\sigma = \Pr_{\theta \sim \Theta}[\varphi(\theta) = \sigma] = \sum_{\theta \in \Theta} \lambda(\theta)\varphi(\theta, \sigma)$ is the probability of announcing signal σ, and (2) $\mu_\sigma(\theta) = \Pr_{\theta \sim \Theta}[\theta | \varphi(\theta) = \sigma] = \frac{\lambda(\theta)\varphi(\theta,\sigma)}{p_\sigma}$ is the *posterior belief distribution* of θ conditioned on signal σ. The converse is also true: every convex decomposition of $\lambda \in \Delta_M$ corresponds to a signaling scheme. Alternatively, the reader can view a signaling scheme φ as the $M \times |\Sigma|$ matrix of pairwise probabilities $\varphi(\theta, \sigma)$ satisfying conditions (1) and (2) with respect to $\lambda \in \Delta_M$. Sometimes we also describe a signaling scheme as $\{p_\mu\}_{\mu \in \Delta_M}$, where $\sum_{\mu \in \Delta_M} p_\mu \mu = \lambda$[1]. The signals Σ in such a signaling scheme are described implicitly, and correspond to the posteriors μ for which $p_\mu > 0$.

[1] We only deal with signaling schemes with finitely many signals in all of our algorithms and analyses, so we use $\sum$ even though we are summing over the uncountable simplex.

Note that each posterior distribution $\mu \in \Delta_M$ defines a complete-information subgame: for every outcome s of the game (i.e., every pure strategy profile), risk neutral players take the expected payoff over $\theta \sim \mu$ as their expected payoff under s. The principal's utility depends on the outcome of the subgames. Given a suitable equilibrium concept and selection rule, we let $f : \Delta_M \to \mathbb{R}$ denote the principal's utility as a function of the posterior distribution μ. For example, in an auction game $f(\mu)$ may be the social welfare or principal's revenue at the induced equilibrium, or any weighted combination of players' utilities, or something else entirely. The principal's objective as a function of the signaling scheme φ can be mathematically expressed by $F(\varphi, \lambda) = \sum_\sigma p_\sigma \cdot f(\mu_\sigma)$.

Fix a Bayesian game, and let $f^+(\lambda)$ denote the value of the optimal signaling scheme when the prior is λ. We note that $f^+(\lambda)$ is a concave function of the prior λ, since if λ^1 and λ^2 form a convex decomposition of λ, so do the optimal posteriors for λ^1 and λ^2. Therefore, the optimal choice of a signaling scheme is related to the *concave envelope* f^+ of the function f ([46, 69]).

Definition 2.1. *The concave envelope f^+ of a function f is the point-wise lowest concave function h for which $h(x) \geq f(x)$ for all x in the domain. Equivalently, the hypograph of f^+ is the convex hull of the hypograph of f.*

Specifically, such a signaling scheme achieves $\sum_\sigma p_\sigma \cdot f(\mu_\sigma) = f^+(\lambda)$. Thus, there exists a signaling scheme with $(M+1)$ signals that maximizes the principal's objective, by applying Carathéodory's theorem[2] to the hypograph of f.

[2] In convex geometry, Carathéodory's theorem [22] states that if a point $x \in \mathbb{R}^d$ lies in the convex hull of a set P, then there exists a subset $P' \subseteq P$ with $|P'| \leq d+1$ such that x is in the convex hull of P'.

2.1.2 Normal Form Games

A normal form game is defined by the following parameters:

- An integer k denoting the number of players, indexed by the set $[k] = \{1, \dots, k\}$.
- An integer n bounding the number of pure strategies of each player. Without loss of generality, we assume each player has exactly n pure strategies, and index them by the set $[n] = \{1, \dots, n\}$.
- A family of payoff tensors $\mathcal{A} = \{A_1, \dots, A_k\}$ with $A_i : [n]^k \to [-1, 1]$, where $A_i(s_1, \dots, s_k)$ is the payoff to player i when each player j plays strategy s_j.

A *Bayesian normal form game* is described by payoff tensors $A_i^\theta : [n]^k \to [-1, 1]$, one per player i and state of nature θ, where $A_i^\theta(s_1, \dots, s_k)$ is the payoff to player i when the state of nature is θ and each player j plays strategy s_j. For a mixed strategy profile $x_1, \dots, x_k \in \Delta_n$, we use $A_i(x_1, \dots, x_k) = \sum_{s_1,\dots,s_k \in [n]} \left(T(s_1, \dots, s_k) \cdot \prod_{i=1}^k x_i(s_i) \right)$ to denote player i's expected payoff over the pure strategy profiles drawn from $(x_1, \dots, x_k)$.

In a general Bayesian normal form game, absent any information about the state of nature beyond the prior λ, risk neutral players will behave as in the complete information game $\mathbb{E}_{\theta \sim \lambda}\left[\mathcal{A}^\theta\right]$. We consider signaling schemes which partially and symmetrically inform players by publicly announcing a signal σ, correlated with θ; this induces a common posterior belief on the state of nature for each value of σ. When players' posterior beliefs over θ are given by $\mu \in \Delta_M$, we use $\mathcal{A}^\mu$ to denote the equivalent complete information game $\mathbb{E}_{\theta \sim \mu}\left[\mathcal{A}^\theta\right]$. As shorthand, we use $A_i^\mu(x_1, \dots, x_k)$ to denote $\mathbb{E}\left[A_i^\theta(s_1, \dots, s_k)\right]$ when $\theta \sim \mu \in \Delta_M$ and $s_i \sim x_i \in \Delta_n$.

The principal's objective is described by a family of tensors $A_0^\theta : [n]^k \to [-1, 1]$, one for each state of nature $\theta \in \Theta$. Equivalently, we may think of the objective as describing the payoffs of an additional player in the game. For a distribution μ over states of nature, we use $A_0^\mu = \mathbb{E}_{\theta\sim\mu}\left[A_0^\theta\right]$ to denote the principal's expected utility in a subgame with posterior beliefs μ, as a function of players' strategies.

Extended Security Games

An *extended security game* is a family of *Bayesian zero-sum games*. A *Bayesian zero-sum game* is specified by a tuple $\left(\Theta, \{A^\theta\}_{\theta\in\Theta}, \lambda\right)$. For each state of nature $\theta \in \Theta$, $A^\theta \in [-1, 1]^{n\times n}$ specifies the payoffs of the row player in a zero-sum game.

Let **Row** and **Col** denote the row player and the column player respectively. An extended security game can be viewed as a polymatrix game between three players: Nature, **Row**, and **Col**. Formally, the payoff matrix for state θ is given by

$$A^\theta \stackrel{\text{def}}{=} \overline{A} + b^\theta \mathbb{1}_n^T + \mathbb{1}_n (d^\theta)^T, \qquad \text{where} \qquad \overline{A} \in \mathbb{R}^{n\times n}, b^\theta \in \mathbb{R}^n, d^\theta \in \mathbb{R}^n. \quad (2.1)$$

Let B and D be matrices having columns $\{b^1, \ldots, b^M\}$, and $\{d^1, \ldots, d^M\}$ respectively. The payoff of the row player is the sum of her payoffs in three separate games: a game $\overline{A}$ between **Row** and **Col**, a game B between **Row** and Nature, and a game D between Nature and **Col**. We obtain the following expressions for A^μ and $f(\mu)$ for $\mu \in \Delta_M$.

$$A^\mu = \overline{A} + (B\mu)\mathbb{1}_n^T + \mathbb{1}_n(\mu^T D^T), \quad f(\mu) = \max_{x\in\Delta_n}\left\{x^T B\mu + \min_{j\in[n]}\left(x^T\overline{A} + \mu^T D^T\right)_j\right\}. \quad (2.2)$$

A special case of an extended security game (and the reason for this terminology) is the *network security game* defined by [46]. Given an undirected graph $G = (V, E)$ with $n = |V|$ and a parameter $\rho \geq 0$, the states of nature correspond to the vertices of the graph. The row and column players are called attacker and defender respectively. The attacker and defender's pure strategies correspond to nodes of G. Let B be the adjacency matrix of G, and set $\overline{A} = D^T = -\rho I_{n\times n}$. Then, for a given state of nature $\theta \in V$, and pure strategies $a, d \in V$ of the attacker and defender, the payoff of the attacker is given by $e_a^T B e_\theta - \rho(e_a^T + e_\theta^T)e_d$. The interpretation is that the attacker gets a payoff of 1 if he selects a vertex a that is adjacent to θ. This payoff is reduced by ρ if the defender's vertex d lies in $\{\theta, a\}$, and by 2ρ if $d = \theta = a$.

2.1.3 Nash Equilibria

The celebrated theorem of Nash [80] states that every finite game has an equilibrium point. The solution concept of Nash equilibrium (NE) has been tremendously influential in economics and social sciences ever since (e.g., see [65]).

For a game with k players and n strategies per player, a *mixed strategy* is an element of $\Delta_{[n]}$, and a *mixed strategy profile* $x = (x_1, \ldots, x_k)$ maps every player i to her mixed strategy $x_i \in \Delta_{[n]}$. Throughout this thesis, we adopt the (approximate) Nash equilibrium as our equilibrium concept. There are two variants. We define them below in normal form games and note that the concept of Nash equilibria is universal in games — it simply states that each player plays a best response to other players' strategies and has no incentive to deviate. We use x_{-i} to denote the strategies of players other than i in x. The support of a vector, *supp*(x), is the set of indices i such that $x_i > 0$.

Definition 2.2. *Let $\epsilon \geq 0$. In a k-player n-action normal form game with expected payoffs in $[-1, 1]$ given by tensors $A_1, \ldots, A_k$, a mixed strategy profile $x_1, \ldots, x_k \in \Delta_n$ is an ϵ-Nash Equilibrium (ϵ-NE) if*

$$A_i(x_1, \ldots, x_k) \geq A_i(t_i, x_{-i}) - \epsilon$$

for every player i and alternative pure strategy $t_i \in [n]$.

Definition 2.3. *Let $\epsilon \geq 0$. In a k-player n-action normal form game with expected payoffs in $[-1, 1]$ given by tensors $A_1, \ldots, A_k$, a mixed strategy profile $x_1, \ldots, x_k \in \Delta_n$ is an ϵ-well-supported Nash equilibrium (ϵ-WSNE) if*

$$A_i(s_i, x_{-i}) \geq A_i(t_i, x_{-i}) - \epsilon$$

for every player i, strategy s_i in the support of x_i, and alternative pure strategy $t_i \in [n]$.

Clearly, every ϵ-WSNE is also an ϵ-NE. When $\epsilon = 0$, both correspond to the exact Nash Equilibrium. Note that we omitted reference to the state of nature in the above definitions — in a subgame corresponding to posterior beliefs $\mu \in \Delta_M$, we naturally use tensors $A_1^\mu, \ldots, A_k^\mu$ instead.

Fixing an equilibrium concept (NE, ϵ-NE, or ϵ-WSNE), a Bayesian game $(\mathcal{A}, \lambda)$, and a signaling scheme $\varphi : \Theta \to \Sigma$, an *equilibrium selection rule* distinguishes an equilibrium strategy profile $(x_1^\sigma, \ldots, x_k^\sigma)$ to be played in each subgame σ. Together with the prior λ, the Bayesian equilibria $\mathcal{X} = \{x_i^\sigma : \sigma \in \Sigma, i \in [k]\}$ induce a distribution $\Gamma \in \Delta_{\Theta \times [n]^k}$ over states of play — we refer to Γ as a *distribution of play*. We say Γ is *implemented* by signaling scheme φ in the chosen equilibrium concept (be it NE, ϵ-NE, or ϵ-WSNE). This is analogous to implementation of allocation rules in traditional mechanism design.

For a signaling scheme φ and associated (approximate) equilibria $\mathcal{X} = \{x_i^\sigma : \sigma \in \Sigma, i \in [k]\}$, our objective function can be written as

$$F(\varphi, \mathcal{X}) = \mathop{\mathbb{E}}_{\theta\sim\lambda}\left[\mathop{\mathbb{E}}_{\sigma\sim\varphi(\theta)}\left[\mathop{\mathbb{E}}_{s\sim x^\sigma}[A_0(\theta, s)]\right]\right].$$

When φ corresponds to a convex decomposition $\{(\mu_\sigma, p_\sigma)\}_{\sigma\in\Sigma}$ of the prior distribution, this can be equivalently written as $F(\varphi, \mathcal{X}) = \sum_{\sigma\in\Sigma} p_\sigma A_0^{\mu_\sigma}(x^\sigma)$. Let $\text{OPT} = \text{OPT}(\mathcal{A}, \lambda, A_0)$ denote the maximum value of $F(\varphi^*, \mathcal{X}^*)$ over signaling schemes φ^* and (exact) Nash equilibria $\mathcal{X}^*$. We seek a signaling scheme $\varphi : \Theta \to \Sigma$, as well as corresponding Bayesian ϵ-equilibria $\mathcal{X}$ such that $F(\varphi, \mathcal{X}) \geq OPT - \epsilon$.

2.1.4 Network Routing Games

A *network routing game* is a tuple $\Big(G = (V, E), \{\ell_e\}_{e\in E}, \{(s_i, t_i, d_i)\}_{i\in[k]}\Big)$, where G is a directed graph with latency function $\ell_e : \mathbb{R}_+ \to \mathbb{R}_+$ on each edge e. Each (s_i, t_i, d_i) denotes a *commodity*; d_i specifies the volume of flow routed from s_i to t_i by self-interested agents, each of whom controls an infinitesimal amount of flow and selects an s_i-t_i path as her strategy. A strategy profile thus corresponds to a multicommodity flow composed of s_i-t_i flows of volume d_i for all i; we call any such flow a *feasible flow*.

The latency on edge e due to a flow x is given by $\ell_e(x_e)$, where x_e is the total flow on e. The latency of a path P is $\ell_P(x) \stackrel{\text{def}}{=} \sum_{e\in P} \ell_e(x_e)$. The total latency of a flow x is $L(\ell, x) \stackrel{\text{def}}{=} \sum_{e\in E} x_e \ell_e(x_e)$. An optimal flow x_{OPT} is a feasible flow with minimum latency. A *Nash flow* (also called a *Wardrop flow*) x_{NE} is a feasible flow where every player chooses a minimum latency path; that is, for all i, all s_i-t_i paths P, Q with $x_{\mathrm{NE}}(e) > 0$ for all $e \in P$, we have $\ell_P(x_{\mathrm{NE}}) \leq \ell_Q(x_{\mathrm{NE}})$. All Nash flows have the same total latency (see, e.g., [85]).

The *price of anarchy (PoA)* measures how the efficiency of a system degrades due to selfish behavior of its agents. In network routing games, the price of anarchy is defined as the ratio between latencies of the Nash flow and the optimal flow: $\mathrm{PoA} = L(\ell, x_{\mathrm{NE}})/L(\ell, x_{\mathrm{OPT}})$. The price of anarchy for a class of latency functions is the maximum ratio over all instances involving these latency functions.

In a *Bayesian network routing game*, the edge latency functions $\{\ell_e^\theta\}_{e\in E}$ may depend on the state of nature $\theta \in \Theta$ (and, as before, we have a prior $\lambda \in \Delta_\Theta$). The principal seeks to *minimize* the latency of the Nash flow. Given $\mu \in \Delta_\Theta$, the expected latency function on each edge e is $\ell_e^\mu(x_e) \stackrel{\mathrm{def}}{=} \sum_{\theta\in\Theta} \mu(\theta)\ell_e^\theta(x_e)$. Define $f(\mu) \stackrel{\mathrm{def}}{=} L(\ell^\mu, x_{\mathrm{NE}}^\mu)$, where x_{NE}^μ is a Nash flow for latency functions $\{\ell_e^\mu\}$. The *signaling problem in a Bayesian routing game* is to determine $(p_\mu)_{\mu\in\Delta_M} \geq 0$ of finite support specifying a convex decomposition of λ (i.e., $\sum_{\mu\in\Delta_M} p_\mu \mu = \lambda$) that minimizes the expected latency of the Nash flow, $\sum_{\mu\in\Delta_M} p_\mu f(\mu)$.

In Bayesian network routing games, Vasserman et al. [90] study the problem of signaling to reduce the average latency. They define the *mediation ratio* as the average latency at equilibrium for the best private signaling scheme, to the average latency for the social optimum, and give tight bounds on the mediation ratio for a special family of Bayesian routing games. In contrast, we study the *computational complexity* of obtaining the best public signaling scheme, and conclude that finding an $(\frac{4}{3} - \epsilon)$-approximation is NP-hard.

2.1.5 Second-Price Auctions

The following parameters describe a variant of an *single-item auction*[3]:

- An integer n denoting the number of bidders, indexed by the set $[n]$.
- A common-knowledge prior distribution $\mathcal{D}$ on *bidders' valuations* $v \in [0,1]^n$, where v_i denotes the value of player i for the item.

A *probabilistic auction* has M possible states of nature $\theta \in \Theta$, and each θ represents a possible instantiation of the item being sold. Instead of having an n-dimensional vector as bidders' valuations, we now have valuation matrices $V \in [0,1]^{n\times M}$, where $V_{i,j}$ denotes the value of player i for the item corresponding to the state of nature j. Again, we have a common-knowledge prior distribution $\mathcal{D}$ on V, given either explicitly or as a "black-box" sampling oracle.

We examine signaling in *probabilistic second-price auctions*, as considered by Emek et al. [52] and Miltersen and Sheffet [21]. In this setting, the item being auctioned is probabilistic, and the instantiation of the item is known to the auctioneer but not to the bidders. The auctioneer commits to a signaling scheme for (partially) revealing information about the item for sale before subsequently running a second-price auction.

As an example, consider an auction for an umbrella: the state of nature θ can be the weather tomorrow, which determines the utility $V_{i,\theta}$ of an umbrella to player i. We assume that λ and $\mathcal{D}$ are independent. We also emphasize that a bidder

[3] In this thesis, for normal form games, we use n to denote the number of strategies and k for the number of players; for multi-player succinct games, we use n to denote the number of players and k for the number of strategies. For both families of games, we are interested in games with $k = O(1)$ and we want to bound the running time as a function of n.

knows nothing about θ other than its distribution λ and the public signal σ, and the auctioneer knows nothing about V besides its distribution $\mathcal{D}$ prior to running the auction.

The game being played is the following:

1. The auctioneer first commits to a signaling scheme $\varphi : \Theta \to \Sigma$;
2. A state of nature $\theta \in \Theta$ is drawn according to λ and revealed to the auctioneer but not the bidders;
3. The auctioneer reveals a public signal $\sigma \sim \varphi(\theta)$ to all the bidders;
4. A valuation matrix $V \in [0,1]^{n \times M}$ is drawn according to $\mathcal{D}$, and each player i learns his value $V_{i,j}$ for each potential item j;
5. Finally, a second-price auction for the item is run.

Note that step (4) is independent of steps (1-3), so they can happen in no particular order.

We adopt the (unique) dominant-strategy truth-telling equilibrium as our solution concept. Specifically, given a signaling scheme $\varphi : \Theta \to \Sigma$ and a signal $\sigma \in \Sigma$, in the subgame corresponding to σ it is a dominant strategy for player i to bid $\mathbb{E}_{\theta \sim \lambda}[V_{i\theta} | \varphi(\theta) = \sigma]$ — his posterior expected value for the item conditioned on the received signal σ. Therefore the item goes to the player with maximum posterior expected value, at a price equal to the second-highest posterior expected value. The algorithmic problem we consider is the one faced by the auctioneer in step (a) — namely computing an optimal signaling scheme — assuming the auctioneer looks to maximize expected revenue.

2.1.6 Majority Voting

Consider an election with n voters and two possible outcomes, 'Yes' and 'No'. For example, voters may need to choose whether to adopt a new law or social policy; board members of a company may need to decide whether to invest in a new project; and members of a jury must decide whether a defendant is declared guilty or not guilty. The new policy passes if the fraction of 'Yes' votes exceeds a certain pre-specified threshold. We index the voters by the set $[n] = \{1, \dots, n\}$. Without loss of generality, we assume utilities are normalized: voter i has a utility of $u_i \in [-1, 1]$ for the 'Yes' outcome, and 0 for the 'No' outcome. In most voting systems with a binary outcome, including threshold voting rules, it is a dominant strategy for voter i to vote 'Yes' if the utility u_i is at least 0.

We study the signaling problems encountered in the context of voting introduced by Alonso and Câmara [5]. In this setting, voters have uncertainty regarding their utilities for the two possible outcomes (e.g., the risks and rewards of the new project). Specifically, voters' utilities are parameterized by an a priori unknown state of nature θ drawn from a common-knowledge prior distribution. We assume voters' preferences are given by a matrix $U \in [-1, 1]^{n \times M}$, where $U_{i,j}$ denotes voter i's utility in the event of a 'Yes' outcome in state of nature j. A voter i who believes that the state of nature follows a distribution $\mu \in \Delta_M$ has expected utility $u(i, \mu) = \sum_{j \in \Theta} U_{i,j}\mu_j$ for a 'Yes' outcome.

We adopt the perspective of a principal — say a moderator of a political debate — with access to the realization of θ, who can determine the signaling scheme through which information regarding the measure is gathered and shared with voters. Alonso and Câmara [5] consider a principal interested in maximizing the probability that at least 50% (or some given threshold) of the voters vote 'Yes', in expectation over states of nature. They characterize optimal signaling schemes analytically, though stop short of prescribing an algorithm for signaling.

2.1.7 Anonymous Games

We study anonymous games $(n, k, \{u_a^i\}_{i\in[n],a\in[k]})$ with n players labeled by $[n] = \{1, \ldots, n\}$, and k common strategies labeled by $[k]$ for each player. The payoff of a player depends on her own strategy, and how many of her peers choose which strategy, but not on their identities. When player $i \in [n]$ plays strategy $a \in [k]$, her payoffs are given by a function u_a^i that maps the possible outcomes (partitions of all other players) Π_{n-1}^k to the interval $[0, 1]$, where $\Pi_{n-1}^k = \{(x_1, \ldots, x_k) \mid x_j \in \mathbb{Z}_+ \wedge \sum_{j=1}^k x_j = n - 1\}$.

Approximate Equilibria in Anonymous Games. We define *ϵ-approximate Nash equilibrium* for anonymous games. The definition is essentially equivalent to ϵ-equilibrium in normal form games (Definition 2.2), except now the game has a succinct representation. A mixed strategy profile s is an ϵ-approximate Nash equilibrium in an anonymous game if and only if

$$\forall i \in [n], \forall a' \in [k], \quad \mathop{\mathbb{E}}_{x\sim s_{-i}}\left[u_{a'}^i(x)\right] \le \mathop{\mathbb{E}}_{x\sim s_{-i}, a\sim s_i}\left[u_a^i(x)\right] + \epsilon,$$

where $x \in \Pi_{n-1}^k$ is the partition formed by $n - 1$ random samples (independently) drawn from $[k]$ according to the distributions s_{-i}. Note that given a mixed strategy profile s, we can compute a player's expected payoff by straightforward dynamic programming (see, e.g., [84]).

Poisson Multinomial Distributions. A *k-Categorical Random Variable* (k-CRV) is a vector random variable supported on the set of k-dimensional basis vectors $\{e_1, \dots, e_k\}$. A k-CRV is i-maximal if e_i is its most likely outcome (break ties by taking the smallest index i). A k-Poisson multinomial distribution of order n, or an (n, k)-PMD, is a vector random variable of the form $X = \sum_{i=1}^{n} X_i$ where the X_i's are independent k-CRVs. The case of $k = 2$ is usually referred to as Poisson Binomial Distribution (PBD).

Note that a mixed strategy profile $s = (s_1, \dots, s_n)$ of an n-player k-strategy anonymous game corresponds to the k-CRVs $\{X_1, \dots, X_n\}$ where $\Pr[X_i = e_a] = s_i(a)$. The expected payoff of player $i \in [n]$ for playing pure strategy $a \in [k]$ can also be written as $\mathbb{E}[u_a^i(X_{-i})] = \mathbb{E}\Big[u_a^i\Big(\sum_{j \neq i, j \in [n]} X_j\Big)\Big]$.

Let $X = \sum_{i=1}^{n} X_i$ be an (n, k)-PMD such that for $i \in [n]$ and $j \in [k]$ we denote $p_{i,j} = \Pr[X_i = e_j]$, where $\sum_{j=1}^{k} p_{i,j} = 1$. For $m = (m_1, \dots, m_k) \in \mathbb{Z}_+^k$, we define the m^{th}-parameter moments of X to be $M_m(X) \stackrel{\text{def}}{=} \sum_{i=1}^{n} \prod_{j=1}^{k} p_{i,j}^{m_j}$. We refer to $\|m\|_1 = \sum_{j=1}^{k} m_j$ as the *degree* of the parameter moment $M_m(X)$.

Total Variation Distance and Covers. The total variation distance between two distributions P and Q supported on a finite domain A is

$$\mathrm{d_{TV}}(P, Q) \stackrel{\text{def}}{=} \max_{S \subseteq A} |P(S) - Q(S)| = (1/2) \cdot \|P - Q\|_1.$$

If X and Y are two random variables ranging over a finite set, their total variation distance $\mathrm{d_{TV}}(X, Y)$ is defined as the total variation distance between their distributions. For convenience, we will often blur the distinction between a random variable and its distribution.

Let $(\mathcal{X}, d)$ be a metric space. Given $\epsilon > 0$, a subset $\mathcal{Y} \subseteq \mathcal{X}$ is said to be a proper ϵ-cover of $\mathcal{X}$ with respect to the metric $d : \mathcal{X}^2 \to \mathbb{R}_+$, if for every $X \in \mathcal{X}$ there exists some $Y \in \mathcal{Y}$ such that $d(X, Y) \leq \epsilon$. We will be interested in constructing ϵ-covers for high-variance PMDs under the total variation distance metric.

Multidimensional Fourier Transform. For $x \in \mathbb{R}$, we will denote $e(x) \stackrel{\text{def}}{=} \exp(-2\pi i x)$. The (continuous) Fourier Transform of a function $F : \mathbb{Z}^k \to \mathbb{C}$ is the function $\widehat{F} : [0,1]^k \to \mathbb{C}$ defined as $\widehat{F}(\xi) = \sum_{x \in \mathbb{Z}^k} e(\xi \cdot x) F(x)$. For the case that F is a probability mass function, we can equivalently write $\widehat{F}(\xi) = \mathbb{E}_{x \sim F}[e(\xi \cdot x)]$.

Let $X = \sum_{i=1}^n X_i$ be an (n, k)-PMD with $p_{i,j} \stackrel{\text{def}}{=} \Pr[X_i = e_j]$. To avoid clutter in the notation, we will sometimes use the symbol X to denote the corresponding probability mass function. With this convention, we can write that $\widehat{X}(\xi) = \prod_{i=1}^n \widehat{X_i}(\xi) = \prod_{i=1}^n \sum_{j=1}^k e(\xi_j) p_{i,j}$.

2.2 The Posterior Selection Problem

The signaling problem can be formulated as the mathematical program (P),

$$\max \sum_{\mu \in \Delta_M} \alpha_\mu f(\mu) \quad \text{s.t.} \quad \sum_{\mu \in \Delta_M} \alpha_\mu \mu(\theta) = \lambda_\theta \quad \text{for all } \theta \in \Theta, \quad \alpha \geq 0. \qquad \text{(P)}$$

Notice that any feasible α must also satisfy $\sum_{\mu \in \Delta_M} \alpha_\mu = 1$; hence, α is indeed a distribution over Δ_M, and a feasible solution to (P) yields a signaling scheme. Let $\text{opt}(\lambda)$ denote the optimal value of (P), and note that this is a *concave* function of λ. Although (P) has a linear objective and linear constraints, it is not quite a linear program (LP) since there are an infinite number of variables. Ignoring this issue for now, we consider the following dual of (P).

$$\min \quad w^T\lambda \qquad \text{s.t.} \qquad w^T\mu \geq f(\mu) \quad \text{for all } \mu \in \Delta_M, \quad w \in \mathbb{R}^M. \tag{D}$$

The separation problem for (D) motivates the following *dual signaling problem.*

Definition 2.4 (Dual signaling with precision parameter ϵ)**.** *Given an objective function $f : \Delta_M \to [-1, 1]$, $w \in \mathbb{R}^M$, and $\epsilon > 0$, distinguish between:*

(i) $f(\mu) \geq w^T\mu + \epsilon$ for some $\mu \in \Delta_M$; if so return $\mu \in \Delta_M$ s.t. $f(\mu) \geq w^T\mu - \epsilon$;

(ii) $f(\mu) < w^T\mu - \epsilon$ for all $\mu \in \Delta_M$.

The posterior selection problem is the special case of dual signaling where $w = \eta\mathbb{1}_M$ for some $\eta \in \mathbb{R}$.

Definition 2.5 (Posterior selection with precision parameter ϵ)**.** *Given an objective function $f : \Delta_M \to [-1, 1]$ and $\epsilon > 0$, find $\mu^* \in \Delta_M$ such that $f(\mu^*) \geq \max_{\mu\in\Delta_M} f(\mu) - \epsilon$.*

2.3 Planted Clique Conjecture

Some of our hardness results are based on the hardness of the *planted-clique* and *planted clique cover* problems.

In the *planted clique problem* $\mathbf{PClique}(n, p, k)$, one must distinguish the n-node Erdős-Rényi random graph $\mathcal{G}(n, \frac{1}{2})$ in which each edge is included independently with probability $\frac{1}{2}$, from the graph $\mathcal{G}(n, \frac{1}{2}, k)$ formed by "planting" a clique in $\mathcal{G}(n, \frac{1}{2})$ at a randomly (or, equivalently, adversarially) chosen set of k nodes. This problem was first considered by Jerrum [66] and Kučera [72], and has been the subject of a large body of work since. A quasipolynomial time algorithm exists when $k \geq 2\log n$, and the best polynomial-time algorithms only succeed for $k = \Omega(\sqrt{n})$

(see, e.g., [4, 30, 40, 56]). Several papers suggest that the problem is hard for $k = o(\sqrt{n})$ by ruling out natural classes of algorithmic approaches (e.g., [55, 57, 66]). The planted clique problem has therefore found use as a hardness assumption in a variety of applications (e.g., [3, 46, 63, 67, 78]).

Conjecture 2.6 (Planted-clique conjecture)**.** *For some $k = k(n)$ satisfying $k = \omega(\log n)$ and $k = o(\sqrt{n})$, there is no probabilistic polynomial time algorithm that solves* $\textbf{PClique}\left(n, \frac{1}{2}, k\right)$ *with constant success probability.*

The planted clique cover problem was introduced in [46]. Multiple cliques are now planted and one seeks to recover a constant fraction of them.

Definition 2.7 (Planted clique cover problem $\textbf{PCover}(n, p, k, r)$ [46])**.** *Let $G \sim \mathcal{G}(n, p, k, r)$ be a random graph generated by:*

(1) including every edge independently with probability p; and

(2) for $i = 1, \ldots, r$, picking a set S_i of k vertices uniformly at random, adding all edges having both endpoints in S_i.

We call the S_i's the planted cliques and p the background density. We seek to recover a constant fraction of the planted cliques $S_1, \ldots, S_r$, given $G \sim \mathcal{G}(n, p, k, r)$.

Dughmi [46] showed that the planted clique cover problem is at least as hard as the planted clique problem. Given an instance G of $\textbf{PClique}(n, p, k)$, we can generate an instance G' of $\textbf{PCover}(n, p, k, r)$ by planting $r - 1$ additional random k-cliques into G (as in step (2) of Definition 2.7). Because the cliques $S_1, \ldots, S_r$ are indistinguishable, recovering a constant fraction of the planted cliques from G' would recover each of $S_1, \ldots, S_r$ with constant probability. In particular, doing so would recover the original planted clique with constant probability.

2.4 The Ellipsoid Method

The *ellipsoid method* is an iterative method for minimizing convex functions, and it does so by generating a sequence of ellipsoids whose volume decreases in each step. The *shallow-cut ellipsoid method* is a variant of the ellipsoid method, which is useful when we only have access to an approximate separation oracle. The classic ellipsoid method, in each iteration, cuts the current ellipsoid using a hyperplane that goes through the center of the ellipsoid. At a high level, the shallow-cut method takes a "shallower" cut that is *close* to the center of the ellipsoid, removing slightly less than half of the current ellipsoid. It turns out that this is still sufficient for finding an (approximately) optimal solution.

We utilize the shallow-cut ellipsoid method to translate hardness results for the posterior selection problem to hardness results for optimal signaling. Formally, we use the following lemma.

Lemma 2.8 (Chapters 4, 6 in [61]; Section 9.2 in [81])**.** *Let $X \subseteq \mathbb{R}^n$ be a polytope described by constraints having encoding length at most L. Suppose that for each $y \in \mathbb{R}^n$, we can determine in time* poly(*size of* y, L) *if $y \notin X$ and if so, return a hyperplane of encoding length at most L separating y from X.*

(i) The ellipsoid method can find a point $x \in X$ or determine that $X = \emptyset$ in time poly(n, L)*.*

(ii) Let $h : \mathbb{R}^n \to \mathbb{R}$ be a concave function and $K = \sup_{x\in X} h(x) - \inf_{x\in X} h(x)$. Suppose we have a value oracle *for h that for every $x \in X$ returns $\psi(x)$ satisfying $|\psi(x) - h(x)| \leq \delta$. There exists a polynomial $p(n)$ such that for any $\epsilon \geq p(n)\delta$, we can use the shallow-cut ellipsoid method to find $x^* \in X$*

such that $h(x^*) \geq \max_{x \in X} h(x) - 2\epsilon$ *(or determine* $X = \emptyset$*) in time* $T = \mathrm{poly}\Big(n, L, \log(\frac{K}{\epsilon})\Big)$ *and using at most* T *queries to the value oracle for* h*.*

Chapter 3

Signaling in Network Routing Games

Navigation services (e.g. Google Maps) on smartphones have changed people's lives during the past few years. While using Google Maps is as easy as typing in your destination and following its navigation, there are three key aspects of this real-life scenario that are relevant to the design of information structures: (1) Driving to work is a game with incomplete information. Traffic has uncertainty, and the congestion of every road is different each day. (2) The principal (Google Maps) knows more about the real-time traffic than the drivers. One of the main reasons why drivers use Google Maps is to learn more about real-time traffic conditions so they can choose a better route. (3) Drivers are selfish and prefer to take the shortest paths to their destination. It is well known that efficiency of network routing degrades due to the selfish behavior of the drivers (see, e.g., [85]). These three aspects motivate us to study signaling in routing games. At a high level, we view the traffic conditions learned by Google Maps as an informational advantage, and we ask if (and how) Google Maps can utilize this advantage to help selfish drivers.

In this chapter, we consider information revelation in (non-atomic, selfish) Bayesian network routing games. We are interested in the most natural setting in which the principal seeks to minimize the average latency experienced by a driver in the system, knowing that the players would act selfishly after learning more about the traffic. Is it possible for the principal to carefully reveal information to reduce the latency of equilibrium flow? And if so, how much can information revelation help selfish routing?

We show that the answer to this question is two-fold.

(i) There are scenarios where the principal can improve selfish routing. Sometimes through the careful provision of information, the principal can achieve the best-coordinated outcome.

(ii) Optimal information revelation is hard in routing games in the worst case: Assuming P $\neq$ NP, there is no polynomial-time algorithm that does better than full revelation.

More specifically, we show that full revelation obtains the *price of anarchy* (defined in Section 2.1.4) of the routing game as its approximation ratio (Theorem 3.1), which is $\frac{4}{3}$ for linear latency functions [85]. We then settle the approximability of the problem by showing that it is NP-hard to obtain a multiplicative approximation better than $\frac{4}{3}$, even with linear latency functions (Theorem 3.2).

The results in this chapter appeared in [15], and my coauthors Umang Bhaskar and Chaitanya Swamy did most of the work.

3.1 Informational Braess' Paradox Revisited

We use a slight variation of the informational Braess' paradox (Example 1.1) to illustrate the problem we are trying to solve, and show that sometimes the principal must reveal some but not all information to minimize the latency of selfish routing. In Example 3.1, the states of nature θ_1, θ_2 are independent random variables, and

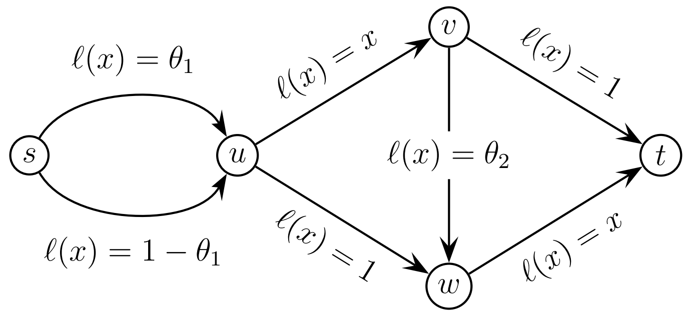

Figure 3.1: A Bayesian network routing game adapted from Example 1.1.

both are drawn uniformly from the set $\{0, 1\}$. There are two edges from s to u, and exactly one of these two edges is going to have latency 0, while the other edge has latency 1. The edge $v \to w$ has latency 0 half of the time, and latency 1 otherwise. In the optimal signaling scheme, the principal reveals the realization of θ_1, but hides the value of θ_2. This is because the drivers have no externality from s to u, so they can all take the shorter edge. But for driving from u to t, the drivers are better off not learning the value of θ_2, otherwise when $\theta_2 = 0$ all of them deviate to the zig-zag path and experience a longer travel time.

In the example above, by revealing no information about the edge $v \to w$, it is as if the principal can remove the edge from the graph. It turns out that this intuition is the key to showing hardness for this problem. We show that the optimal signaling problem in routing games is as hard as the *network design problem* studied by Cole et al. [31], where the principal puts taxes on the edges to minimize the total drivers' disutility (latency + tax). It was shown in [31] that the problem is equivalent to deciding which edges to remove to minimize the latency of the Nash flow, so we call it the network design problem. Our reduction constructs a Bayesian routing game from an instance of the network design problem, in which we can translate principal's signaling scheme back to a set of taxes on the edges — the principal puts more taxes on an edge if he reveals less information about it (and vice versa).

3.2 Full Revelation Is a (4/3)-Approximation

In this section, we prove that full revelation is a $\frac{4}{3}$-approximation for signaling in Bayesian network routing games with linear latency functions. Recall that the price of anarchy (PoA) for a class of latency functions is the maximum ratio, over all instances involving these latency functions, of the latencies of the Nash flow and optimal flow. For linear latency functions, the PoA is $\frac{4}{3}$ [85].

Intuitively, the result follows because full revelation is the best signaling scheme if one seeks to minimize the expected latency of the *optimal* flow, and the multiplicative error that results from this change in objective (from the latency of the Nash flow to that of the optimal flow) cannot exceed the price of anarchy. Our result directly generalizes to arbitrary latency functions and multi-commodities, and the approximation ratio of full revelation is bounded by the PoA of the set of allowable latency functions.

Theorem 3.1. *In Bayesian routing games, the full revelation signaling scheme has the price of anarchy for the underlying latency functions as its approximation ratio. In particular, for linear latencies, full revelation achieves a $\frac{4}{3}$-approximation.*

Proof. Given a state of nature $\theta \in \Theta$, we use x^θ_{NE} and x^θ_{OPT} to denote the Nash flow and the optimal flow with respect to the latency functions $\{\ell^\theta_e\}$ respectively. Let ρ be the price of anarchy for the collection $\{\ell^\theta_e\}_{e\in E,\theta\in\Theta}$ of latency functions, so we have $L(\ell^\theta, x^\theta_{\mathrm{OPT}}) \geq L(\ell^\theta, x^\theta_{\mathrm{NE}})/\rho)$ for all $\theta \in \Theta$.

The full revelation signaling scheme has latency $L \stackrel{\text{def}}{=} \sum_{\theta\in\Theta} \lambda_\theta L(\ell^\theta, x^\theta_{\mathrm{NE}})$. We show that the average latency of any signaling scheme $\{p_\mu\}$ is at least (L/ρ).

$$\begin{aligned}
\sum_{\mu \in \Delta_M} p_\mu f(\mu) &= \sum_{\mu \in \Delta_M} p_\mu L(\ell^\mu, x^\mu_{\text{NE}}) \\
&= \sum_{\mu \in \Delta_M} p_\mu \sum_{\theta \in \Theta} \mu(\theta) L(\ell^\theta, x^\mu_{\text{NE}}) \\
&\geq \sum_{\mu \in \Delta_M} p_\mu \sum_{\theta \in \Theta} \mu(\theta) L(\ell^\theta, x^\theta_{\text{OPT}}) \\
&= \sum_{\theta \in \Theta} \lambda_\theta L(\ell^\theta, x^\theta_{\text{OPT}}) \\
&\geq \sum_{\theta} \lambda_\theta L(\ell^\theta, x^\theta_{\text{NE}})/\rho. \qquad \square
\end{aligned}$$

3.3 NP-hard to Approximate Better Than 4/3

We now prove matching hardness result to show that Theorem 3.1 is tight. The proof of Theorem 3.2 is a direct reduction from the problem of computing edge tolls that minimize the total (latency + toll)-cost of the resulting equilibrium flow.

Theorem 3.2. *For any $\epsilon > 0$, obtaining a $(\frac{4}{3} - \epsilon)$-approximation for the signaling problem in Bayesian routing games is NP-hard, even in single-commodity games with linear latency functions.*

Let $\big(G, s, t, d\big)$ be a single-commodity routing game. By scaling latency functions suitably, we may assume that $d = 1$ and omit it from now on. We reduce from the problem of determining edge tolls $\tau \in \mathbb{R}^E_+$ that minimize $L(\ell + \tau, x^{\ell+\tau}_{\text{NE}})$, where $\ell + \tau$ denotes the collection of latency functions $\{\ell_e(x) + \tau_e\}_e$, and $x^{\ell+\tau}_{\text{NE}}$ is the Nash flow for $\ell + \tau$. Note that $L(\ell + \tau, x) = \sum_e x_e(\ell_e(x_e) + \tau_e)$ takes into account the contribution from tolls; we refer to this as the total *cost* of x. The problem of computing optimal tolls that minimizes (latency + toll) is inapproximable within a factor of $\frac{4}{3}$.

Theorem 3.3 ([31]). *There are optimal tolls where the toll on every edge is* 0 *or* ∞. *For every* $\epsilon > 0$, *there is no* $\left(\frac{4}{3} - \epsilon\right)$*-approximation algorithm for the problem of computing optimal tolls in networks with linear latency functions, unless* $P = NP$.

Let $\left(G = (V, E), \ell, s, t\right)$ be an instance of a network design problem with linear latencies. Let $m = |E| \geq 5$. Let $L_0 = L(\ell, x_{\text{NE}}^{\ell})$ be the latency of the Nash flow for the original graph ℓ. Let τ^* be optimal $\{0, \infty\}$-tolls, $L^* = L(\ell + \tau^*, x_{\text{NE}}^{\ell+\tau^*})$ be the optimal cost, and $B^* \stackrel{\text{def}}{=} \{e \in E : \tau_e^* = \infty\}$. We can view τ^* as removing the edges in B^*.

We create a Bayesian routing game as follows. Let $\left(G_1 = (V_1, E_1), s_1, t_1\right)$ and $\left(G_2 = (V_2, E_2), s_2, t_2\right)$ be two copies of (G, ℓ, s, t). Add vertices s, t, and edges (s, s_1), (s, s_2) and (t_1, t), (t_2, t). Call the graph thus created H. For $e \in E_1 \cup E_2$ with corresponding edge $e' \in E$, we set the latency function in the new graph to be $h_e(x) = \ell_{e'}(x)$, and we set $h_e(x) = 0$ for $e \in \{(s, s_1), (s, s_2), (t_1, t), (t_2, t)\}$.

Next, we introduce uncertainty to this game by randomly removing one edge in H. Each state of nature θ corresponds to an edge e in H (i.e., $\Theta = E_H$), which is going to be effectively removed from the graph. Formally, we set:

$$\lambda_\theta = \begin{cases} \frac{1}{m^2} & \text{if } \theta \in E_1 \cup E_2, \\ \frac{1}{2} - \frac{1}{m} & \text{if } e = (s, s_1) \text{ or } (s, s_2), \\ 0 & \text{if } e = (t_1, t) \text{ or } (t_2, t). \end{cases} \qquad h_e^\theta(x) = \begin{cases} h_e(x) + m^2 L_0 & \text{if } \theta = e, \\ h_e(x) & \text{otherwise.} \end{cases}$$

Our Bayesian routing game is $\left((H, \{h_e^\theta\}_{\theta,e}, s, t), \lambda\right)$.

The idea is that state θ encodes the removal of edge θ: specifically, if $\mu(\theta) \geq \frac{1}{m^2}$ for a posterior μ, then h^μ simulates the edge θ breaking down due to the large constant term $m^2 L_0$. Let B_1, B_2 be the edge-sets corresponding to B^* in G_1 and G_2

respectively. The prior λ is set up to satisfy two important properties: (1) it admits a convex decomposition into posteriors μ_1 and μ_2, where μ_1 simulates that $G_1 \setminus B_1$ is connected to s and G_2 is disconnected from s (similarly for μ_2); and (2) any reasonable signaling scheme must put most of the probability mass into posteriors μ, where h^μ connects only one of G_i to s, so that $\{\mu_e m^2 L_0\}_{e\in E_i}$ yields tolls τ for edges in the network design problem with $L(\ell+\tau, x_{\mathrm{NE}}^{\ell+\tau}) \leq f(\mu)$. Lemma 3.4 and 3.5 make the statements in (1) and (2) precise, and Theorem 3.2 follows immediately from Lemmas 3.4, 3.5 and Theorem 3.3.

Lemma 3.4. *Let L^* be the cost of optimal tolls for a network design instance (G, ℓ, s, t). There is a signaling scheme for the above Bayesian routing game $\big((H, \{h_e^\theta\}_{\theta,e}, s, t), \lambda\big)$ with expected latency L^*.*

Proof. We partition the edges of H into two sets: $B_1 \cup (E_2 \setminus B_2) \cup \{(s, s_2)\}$ as one set, and the remaining edges as the other. We claim the signaling scheme that reveals which set contains the broken edge θ has expected latency L^*.

Formally, define posterior $\mu_1 \in \Delta_{E_H}$ as:

$$\mu_1(\theta) = \begin{cases} \frac{2}{m^2} & \text{if } \theta \in B_1 \text{ or } \theta \in (E_2 \setminus B_2), \\ 1 - \frac{2}{m} & \text{if } \theta = (s, s_2), \\ 0 & \text{otherwise.} \end{cases}$$

We can define μ_2 symmetrically and check that $\lambda = (\mu_1 + \mu_2)/2$. We will show that $f(\mu_1) = f(\mu_2) = L^*$, proving the lemma.

Consider distribution μ_1. The argument for μ_2 is symmetrical. The idea is that an edge e with $\mu_1(e) > 0$ has $h_e^{\mu_1}(x) \geq 2L_0$, which effectively deletes e from H; other edges have $h_e^{\mu_1}(x) = h_e(x)$. So μ_1 retains edges in $G_1 \setminus B_1$, and disconnects

G_2 from s. The remaining graph corresponds to the optimal solution of the network design problem on G_1, removing the bad edges B_1 and adding two extra edges (s, s_1) and (t_1, t) both with latency 0. Therefore, the latency of the Nash flow under μ_1 is exactly L^*. ☐

Lemma 3.5. *Given a signaling scheme* $\{p_\mu\}_{\mu \in \Delta_M}$ *for the Bayesian routing game* $\left((H, \{h_e^\theta\}_{\theta,e}, s, t), \lambda\right)$ *with expected latency* L*, one can obtain tolls* τ *such that the routing game* $(G, \ell + \tau, s, t)$ *has Nash latency at most* $\frac{L}{1-3/m}$.

Proof. Assume $L < L_0$, otherwise $\tau = 0$ suffices. By Markov's inequality, at least $\left(1 - \frac{2}{m-1}\right)$ of the probability mass of p must be on posteriors μ with $\mu_{(s,s_1)} + \mu_{(s,s_2)} \geq 1/m$. There must exist such a posterior μ with $f(\mu) \leq \frac{L}{1-2/(m-1)} \leq \frac{L}{1-3/m}$.

Fix such a posterior μ. Without loss of generality, we assume $\mu_{(s,s_1)} \geq \frac{1}{2m}$. Let $x = x^\mu_{\mathrm{NE}}$ be the Nash flow for latency functions h^μ. Since $m \geq 5$ and $L < L_0$, we have $h^\mu_{(s,s_1)} \geq mL_0/2 > \frac{L}{1-3/m} \geq f(\mu)$, so we must have $x_{(s,s_1)} = 0$, i.e., x is supported on G_2.

For $e \in E_2$, we also use e to denote the corresponding edge in E. For every $e \in E_2$, we have $h^\mu_e(x) = \ell_e(x) + \mu_e m^2 L_0$. Thus, defining $\tau_e = \mu_e m^2 L_0$ for all $e \in E$, we obtain that x restricted to E_2 (with s_2 corresponding to s) is a Nash flow for $(G, \ell + \tau, s, t)$. The latency of the restricted flow is equal to $f(\mu)$, because under the posterior μ, every s-t path in G corresponds to an s_2-t_2 path in H with the same latency. ☐

Chapter 4

Signaling in Normal Form Games

In the previous chapter, we studied how information revelation can help selfish routing, and we showed that the principal must solve NP-hard problems to do even slightly better than full revelation. In this chapter, we examine the complexity of optimal signaling in one of the most fundamental classes of games: normal form games. As we will see, the problem of (approximately) optimal signaling in normal form games is computationally easier than optimal signaling in routing games — the principal can obtain an ϵ-additive optimal signaling scheme in *quasipolynomial time* for any constant $\epsilon > 0$. and this cannot be improved to polynomial time assuming the planted clique conjecture.

Recall that in Bayesian normal form games, we have a principal and a game whose payoff entries depend on the state of nature θ. Players only know the common prior distribution of θ, while the principal knows the realization of θ and seeks to reveal partial information about θ to induce a desirable equilibrium. For the formal definition of signaling in Bayesian normal form games, see Section 2.1.2.

4.1 Two Examples

4.1.1 Bayesian Prisoner's Dilemma Revisited

We start with a variation of the Prisoner's dilemma given in Example 1.2.

	Cooperate	Defect
Cooperate	$-2, -2$	$-5, 0$
Defect	$0, -5$	$-4, -4$

(a) Payoff when $\theta = 1$.

	Cooperate	Defect
Cooperate	$1, 1$	$-5, 0$
Defect	$0, -5$	$-4, -4$

(b) Payoff when $\theta = 2$.

Figure 4.1: A Bayesian normal form game adapted from the Prisoner's dilemma.

In Figure 4.1, the payoff of the game depends on the state of nature θ, which is drawn from $\{1, 2\}$ uniformly by nature. We are given two normal form games as input, one for each possible state of nature. In each cell, the first number represents the payoff of the row player, and the second number represents the payoff of the column player. The principal is interested in maximizing the (expected) social welfare. If the principal reveals full information, the players defect when $\theta = 1$ and cooperate when $\theta = 2$. The expected social welfare is $(-8 + 2)/2 = -3$.

Consider a signal σ and the corresponding posterior belief μ over θ. Let $\mu(1) = \Pr[\theta = 1]$ and $\mu(2) = \Pr[\theta = 2]$. The expected payoff for both players to cooperate is $\Pr[\theta = 1] \cdot (-2) + \Pr[\theta = 2] \cdot 1 = \mu(2) - 2\mu(1) = 1 - 3\mu(1)$. When $\mu(1) \leq 1/3$, the payoff is enough to incentivize risk neutral players to cooperate.

The optimal signaling scheme uses two signals, σ_1 (the "defect" signal) and σ_2 (the "cooperate" signal). The principal announces σ_2 whenever $\theta = 2$. When $\theta = 1$, the principal announces σ_2 with probability $1/2$ and announces σ_1 otherwise. Conditioned on the signal being σ_2, we have $\Pr[\theta = 1] = 1/3$. Based on the analysis above, players will cooperate under σ_2 and the expected social welfare of the optimal signaling scheme is

$$\begin{aligned} & 2\Big(\Pr[\theta = 1, \sigma_1] \cdot (-4) + \Pr[\theta = 1, \sigma_2] \cdot (-2) + \Pr[\theta = 2] \cdot 1\Big) \\ = {} & 2\left(\frac{1}{4} \cdot (-4) + \frac{1}{4} \cdot (-2) + \frac{1}{2} \cdot 1\right) = -2. \end{aligned}$$

Example 3.1 is similar to the informational Braess' paradox (Example 1.1), in which the principal tries to help the players fight their selfishness through careful provision of information. This is merely one of the many facets of optimal information revelation, and we will see a different perspective in the next example.

4.1.2 Helping a Friend in a Poker Game

	Fold	Call
Check	1	1
Bet	1	2

(a) **Row**'s payoff when $\theta = 1$.

	Fold	Call
Check	-1	-1
Bet	1	-2

(b) **Row**'s payoff when $\theta = 2$.

Figure 4.2: A Bayesian poker game.

Consider a Bayesian zero-sum game given in Figure 4.2 [1]. We use **Row** and **Col** to denote the row and column players respectively. The game proceeds as follows:

1. Each player puts \$1 on the table, then gets a card. Players do not get to see the cards (including their own).
2. **Row** goes first, and she can choose to bet another \$1 (Bet) or not (Check).
3. If **Row** bets, **Col** can choose to put in \$1 as well (Call) or to give up (Fold).

If **Row** bets and **Col** folds, **Row** wins automatically. Otherwise both cards are revealed, and the player with the higher card wins and takes all the money. Note that in this game, **Col** does not have the option to bet if **Row** checks.

[1] This example is inspired by Will Ma's talk on Poker at Google NYC in summer 2013.

To simplify the problem, we assume nature flips a coin to decide who has the higher card. If $\theta = 1$ then **Row** has a higher card, and if $\theta = 2$ then **Col** has a higher card. This game can be represented in normal form as in Figure 4.2. When **Row** bets and **Col** folds, the value of the cards are irrelevant; **Row** takes the money, winning \$1 from **Col**. In all other cases, the cards are revealed and the player with higher card wins and gets either \$1 or \$2 from her opponent.

If **Row** and **Col** play this game without extra information about θ, the value of the game is 0, because **Col** will never fold and neither player has an advantage. Now suppose there is a principal who knows the exact value of θ and wants to help **Row** in this game by sending a signal to both players. If the principal reveals full information, then **Row** wins \$1 when $\theta = 1$ and loses \$1 when $\theta = 2$, so the expected payoff of **Row** is again 0.

Consider a signal σ and the corresponding posterior belief μ. Let $\mu(1)$ denote the probability that $\theta = 1$.

- If $\mu(1) < 1/2$, then **Col** has a higher chance of winning if the cards are revealed, so **Col** always calls and **Row** has no incentive to bet. The equilibrium of the game is (Check, Call), and **Row**'s expected payoff is

$$\Pr[\theta = 1] - \Pr[\theta = 2] = \mu(1) - (1 - \mu(1)) = 2\mu(1) - 1.$$

- For $1/2 \leq \mu(1) < 3/4$, **Row** has a higher chance of winning if the cards are revealed, so she prefers to bet. If **Col** folds, she loses \$1 for sure. If **Col** calls, her expected payoff is

$$2\,(\Pr[\theta = 2] - \Pr[\theta = 1]) = 2 - 4\mu(1) > -1,$$

so the equilibrium of the game is (Bet, Call), and **Row**'s expected payoff is $4\mu(1) - 2$.

- When $\mu(1) \geq 3/4$, the analysis in similar to that in the previous case, except that the probability of **Row** winning the showdown ($\theta = 1$) is large enough that it is better for **Col** to fold. The equilibrium of the game is (Bet, Fold), and **Row**'s payoff is 1.

Recall that $f(\mu)$ is the principal's objective (in this case, **Row**'s expected payoff) under the posterior μ. We can plot out $f(\mu)$ as a function of $\mu(1)$ (Figure 4.3). The optimal signaling scheme (shown in decomposition form in Figure 4.3) uses two signals, one revealing **Row** has the lower card ($\mu(1) = 0$, **Row** loses \$1), and one just enough to force **Col** to fold when **Row** bets ($\mu(1) = 3/4$, **Row** wins \$1). The first signal appears with probability 1/3, and the second signal appears with probability 2/3, so **Row** and the principal's expected utility is 1/3.

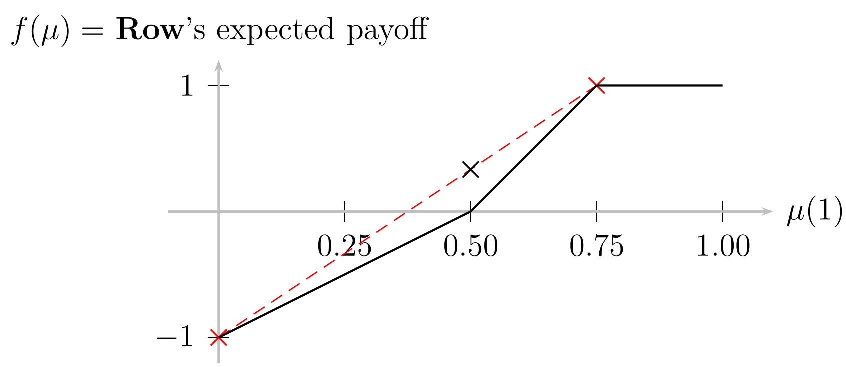

Figure 4.3: **Row**'s expected payoff as a function of $\mu(1) = \Pr[\theta = 1]$. The dashed line represents the optimal signaling scheme in the prior decomposition form.

In the poker game, depending on the posterior beliefs, there are three different outcomes of the game. This can be observed in the three sections of the piecewise linear function in Figure 4.3. In order to decompose the prior optimally, there is no need for more than one signal per outcome, due to the fact that the function is linear in each section. This observation is crucial for our algorithm in Section 4.3.

We hope that these two examples illustrate the importance and complexity of optimal signaling in normal form games. We will now proceed to state our algorithmic and hardness results.

4.2 Summary of Results

In this chapter, we investigate the computational complexity of optimal signaling in Bayesian zero-sum games.

In Section 4.3, we develop the first (bi-criteria) quasipolynomial time approximation scheme (QPTAS) for signaling in normal form games. In other words, for every constant $\epsilon > 0$, we can in quasipolynomial time compute a near-optimal signaling scheme, losing an additive ϵ in the objective as well as in the equilibrium constraints (Theorem 4.1).

In Section 4.4, we first show that the relaxation in players' incentive constraints is necessary, otherwise the problem becomes NP-hard (Theorem 4.5). We then settle the complexity of the signaling problem with respect to NP-hardness by showing that it is NP-hard to obtain an additive FPTAS (Theorem 4.7). Finally, we show that assuming the planted clique conjecture (Conjecture 2.6), the QPTAS in Section 4.3 is essentially optimal (Theorem 4.10).

It is worth noting that our algorithm applies to general sum normal form games with any constant number of players, and an abstract class of objectives which includes the social welfare and weighted combinations of player utilities as a special case; while *all* of our hardness results hold for Bayesian *two-player zero-sum* games. Zero-sum games admit a canonical and tractable notion of equilibrium, which allows us to study the complexity of optimal signaling without equilibrium computation concerns.

The work presented in this section appeared previously as research papers. The algorithm in Section 4.3 appeared in [27], and the hardness results in Section 4.4 appeared in [15].

Previous and Recent Work

Dughmi [46] initiated the computational study of signaling in Bayesian zero-sum games, and obtained various hardness results. Specifically, it was shown that no FPTAS is possible for the signaling problem for zero sum games, assuming the planted clique conjecture. In Section 4.4.2, we strengthen the result of [46] by ruling out an FPTAS assuming P $\neq$ NP, and ruling out a PTAS based on the planted clique conjecture.

Recently, Rubinstein [86] showed that our QPTAS in Theorem 4.3 is essentially tight assuming the Exponential Time Hypothesis (ETH). Compared to our hardness results in Section 4.4.3, [86] replaced our average-case assumption of planted clique hardness with a more conventional worst-case assumption (ETH).

4.3 QPTAS for Signaling in Normal Form Games

In this section, we consider signaling in normal form games when the adopted solution concept is the ϵ-well-supported Nash equilibrium (ϵ-WSNE)[2], and give the first QPTAS for this problem. Our approach consists of two main steps.

1. *Construct a Small Dictionary of Equilibria*: This is a discrete family of objects which indexes the potential equilibria of a signaling scheme, with the property that they form an ϵ-cover of the space of all equilibria with respect to the space of signaling schemes and the design objective.
2. *Construct a near-optimal Signaling Scheme*: We then compute a near-optimal signaling scheme which induces subgames with equilibria from our dictionary. This involves solving a nontrivial optimization problem which optimally decomposes the prior distribution into posterior beliefs inducing equilibria in our dictionary.

Our dictionary is based on the work of Lipton et al. [74]. Specifically, [74] shows the existence of a quasipolynomial-size family of mixed strategy profiles which, simultaneously for all games and equilibria of those games, includes a profile which approximates the payoffs of the equilibrium to within an additive ϵ, and itself forms an ϵ-equilibrium. To combine these approximate equilibria into a signaling scheme, we make two observations: First, the space of posterior beliefs which induce a particular equilibrium forms a convex polytope; second, the optimization problem of optimally partitioning the prior belief into a quasipolynomial number of posterior

[2] Since every ϵ-WSNE is also an ϵ-approximate Nash equilibrium (ϵ-NE), our results apply to ϵ-NE as well.

beliefs, one in each polytope corresponding to an equilibrium, can be formulated via a linear program after an appropriate change of variables.

Formally, we prove the following bi-criteria result. Recall that $F(\varphi, \mathcal{X})$ is the principal's objective value for the signaling scheme φ and associated (approximate) equilibria $\mathcal{X}$; and $\text{OPT}(\mathcal{A}, \lambda)$ is the maximum reward for the principal over all possible signaling schemes and (exact) equilibria.

Theorem 4.1. *Fix $\epsilon > 0$. Given as input an explicitly-described Bayesian normal form game $(\mathcal{A}, \lambda)$ with $k = O(1)$ players, n actions, and M states of nature, and an objective $A_0 : [M] \times [n]^k \to [-1, 1]$, there is an algorithm with runtime* $\text{poly}(M, n^{\log n/\epsilon^2})$ *which outputs a signaling scheme φ and corresponding Bayesian ϵ-well-supported Nash equilibria $\mathcal{X}$ satisfying $F(\varphi, \mathcal{X}) \geq OPT(\mathcal{A}, \lambda) - \epsilon$.*

The proof of Theorem 4.1 hinges on three main lemmas. The first lemma (Lemma 4.2) allows us to restrict attention to equilibria with small support, which follows easily from the results of [74][3]. The second lemma (Lemma 4.3) states that the posterior beliefs implementing a particular approximate equilibrium form a simple polytope, in doing so reducing the signaling problem to optimization over convex decompositions of λ into a family of posteriors, each belonging to a given polytope. The third lemma (Lemma 4.4) shows that optimization over such convex decompositions reduces to a linear program.

Lemma 4.2. *Let tensors $A_1, \ldots, A_k : [n]^k \to [-1, 1]$ describe a k-player game of complete information with n pure strategies per player, and let $A_0 : [n]^k \to [-1, 1]$*

[3] Babichenko et al. [8] later improved the parameters of [74]. For a normal form game with k players and n strategies and for any constant $\epsilon > 0$, [8] proves there exists an ϵ-cover of Nash equilibria in which each player randomizes uniformly among a set of size $O(\log k + \log n)$; in contrast [74] requires a set of size $O(k^2(\log k + \log n))$. Since we are interested in the setting where the number of players $k = O(1)$, the two results are asymptotically the same for us.

be a tensor describing an objective function on pure strategy profiles. For each $\epsilon > 0$, there exists an integer $s = s(\epsilon) = O(k^2 \log(kn)/\epsilon)$, such that for every mixed strategy profile $x = (x_1, \ldots, x_k)$, there is a profile $\widetilde{x} = (\widetilde{x}_1, \ldots, \widetilde{x}_k)$ of s-uniform mixed strategies such that $|A_i(x) - A_i(\widetilde{x})| \leq \epsilon$ for all players i, $|A_0(x) - A_0(\widetilde{x})| \leq \epsilon$, and if x is a Nash equilibrium of $\mathcal{A} = \{A_1, \ldots, A_k\}$ then $\widetilde{x}$ is an ϵ-WSNE of $\mathcal{A}$.

Lemma 4.3. *Fix a normal form game of incomplete information $\{A_i^\theta \in [-1,1]^{n^k} : i \in [k], \theta \in [M]\}$ with k players, n actions, and M states of nature. Consider a mixed strategy profile $x = (x_1, \ldots, x_k)$ with $x_i \in \Delta_n$. For each $\epsilon \geq 0$, the set of posterior beliefs inducing x as an ϵ-WSNE is a convex polytope described by $L = \text{poly}(k, n)$ linear inequalities.*

Lemma 4.4. *Given a family of non-empty polytopes $\mathcal{P}_1, \ldots, \mathcal{P}_t \subseteq \Delta_M$ described by L inequalities each, a point $\lambda \in \Delta_M$, and linear objectives $w_1, \ldots, w_t \in \mathbb{R}^M$, the non-linear optimization problem (4.1) can be solved in $\text{poly}(t, L, M)$ time.*

$$\begin{aligned} \text{maximize} \quad & \textstyle\sum_{\sigma=1}^t p_\sigma w_\sigma \cdot \mu_\sigma \\ \text{subject to} \quad & \textstyle\sum_{\sigma=1}^t p_\sigma = 1 \\ & \textstyle\sum_{\sigma=1}^t p_\sigma \mu_\sigma = \lambda \\ & \mu_\sigma \in \mathcal{P}_\sigma, \qquad \text{for } \sigma = 1, \ldots, t. \end{aligned} \tag{4.1}$$

Before proving each of these lemmas, we first elaborate on how they imply Theorem 4.1.

Proof of Theorem 4.1. Given a signaling scheme φ with decomposition form (p, μ), and an (approximate) equilibrium x^σ for each subgame corresponding to σ, the principal's objective value is

$$F(\varphi, x) = \sum_{\sigma \in \Sigma} p_\sigma A_0(\mu_\sigma, x^\sigma)$$

where $A_0(\mu, x)$ denotes $\mathbb{E}_{\theta \sim \mu}[\mathbb{E}_{s \sim x}[A_0(\theta, s)]]$.

Let $\mathcal{N}_{n,\epsilon} \subseteq \Delta_n$ denote the set of all s-uniform mixed strategies. Lemma 4.2 implies that, in order to complete the proof of Theorem 4.1, it suffices to show how to exactly optimize, in the claimed time, over signaling schemes in which $x^\sigma \in (\mathcal{N}_{n,\epsilon})^k$ for each signal $\sigma \in \Sigma$. We may restrict attention to signaling scheme/equilibrium combinations for which each mixed strategy profile $x \in (\mathcal{N}_{n,\epsilon})^k$ is selected for at most one subgame: when x is the equilibrium for both the subgames A^{σ_1} and A^{σ_2}, we can "merge" the two signals σ_1 and σ_2 into a signal (σ_1, σ_2), giving rise to a new subgame $A^{(\sigma_1,\sigma_2)}$ with posterior belief $\mu_{(\sigma_1,\sigma_2)} = \frac{p_{\sigma_1}}{p_{\sigma_1}+p_{\sigma_2}}\mu_{\sigma_1} + \frac{p_{\sigma_2}}{p_{\sigma_1}+p_{\sigma_2}}\mu_{\sigma_2}$ and probability $p_{(\sigma_1,\sigma_2)} = p_{\sigma_1} + p_{\sigma_2}$. Lemma 4.3 implies that x remains an (approximate) equilibrium of the merged subgame. Moreover, the objective is unchanged because $A_0(\mu, x)$ is linear in its first argument.

We first discard strategy profiles in $(\mathcal{N}_{n,\epsilon})^k$ which cannot be induced as equilibria of any posterior belief. This can be done in polynomial time, by checking whether the corresponding polytope (as given by Lemma 4.3) is empty. Let $\mathcal{N} \subseteq (\mathcal{N}_{n,\epsilon})^k$ denote the set of s-uniform ϵ-equilibria that can be induced by some posterior belief. For notational convenience we assume that each $x \in \mathcal{N}$ is induced as an equilibrium of exactly one subgame, by allowing signals which occur with probability 0. Since the number of players k is a constant, we can index $\mathcal{N} \subseteq (\mathcal{N}_{n,\epsilon})^k$ as $\{x^1, \ldots, x^t\}$ for $t = |\mathcal{N}| \le |\mathcal{N}_{n,\epsilon}|^k = n^{O(\log n/\epsilon^2)}$, and we can write our optimization task as follows.

$$
\begin{aligned}
\text{maximize} \quad & \textstyle\sum_{\sigma=1}^{t} p_\sigma A_0(\mu_\sigma, x^\sigma) \\
\text{subject to} \quad & \textstyle\sum_{\sigma=1}^{t} p_\sigma = 1 \\
& \textstyle\sum_{\sigma=1}^{t} p_\sigma \mu_\sigma = \lambda \\
& x^\sigma \text{ is an equilibrium of } A^{\mu_\sigma}, \quad \text{for } \sigma = 1, \ldots, t.
\end{aligned} \tag{4.2}
$$

Lemma 4.3, and the linearity of $A_0(\mu, x)$ in its first argument, imply that optimization problem (4.2) is of the form given in (4.1) with $L = \text{poly}(k, n)$ constraints. Lemma 4.4, and our assumption that $k = O(1)$, imply that (4.2) can be solved in time $\text{poly}(M, n^{\log n/\epsilon^2})$. This completes the proof of Theorem 4.1. □

We now prove Lemmas 4.2, 4.3, and 4.4.

Proof of Lemma 4.2

We can think of the objective tensor $A_0^\mu : [n]^k \to [-1, 1]$ as describing the utility of an additional player (the principal) in the game with a trivial strategy set. The rest follows from [74, Theorem 2].

Proof of Lemma 4.3

For x to be an ϵ-WSNE of $A^\mu = \sum_{\theta=1}^{M} \mu(\theta) A^\theta$, the following set of inequalities must hold for $\mu \in \Delta_M$:

$$\textstyle\sum_{\theta=1}^{M} \mu(\theta) A_i^\theta(j, x_{-i}) \geq \sum_{\theta=1}^{M} \mu(\theta) A_i^\theta(k, x_{-i}) - \epsilon, \quad \text{for } i \in [k], j \in supp(x_i), k \in [n].$$

Since x is fixed, we have a system of $\text{poly}(k, n)$ linear inequalities in μ.

Proof of Lemma 4.4

We write an equivalent linear program via a change of variables. Specifically, we let $\nu_\sigma = p_\sigma \mu_\sigma$. Observe that after this change (4.1) becomes:

$$\begin{aligned} \text{maximize} \quad & \textstyle\sum_{\sigma=1}^{t} w_\sigma \cdot \nu_\sigma \\ \text{subject to} \quad & \textstyle\sum_{\sigma=1}^{t} p_\sigma = 1 \\ & \textstyle\sum_{\sigma=1}^{t} \nu_\sigma = \lambda \\ & \tfrac{\nu_\sigma}{p_\sigma} \in \mathcal{P}_\sigma, \qquad \text{for } \sigma = 1, \ldots, t. \end{aligned} \tag{4.3}$$

(4.3) is not yet a linear program. However, since $\mathcal{P}_\sigma$ is described by an explicit set of inequalities $C^\sigma y \preceq b^\sigma$, the non-linear inequalities $C^\sigma \nu_\sigma / p_\sigma \preceq b^\sigma$ can be rewritten as $C^\sigma \nu_\sigma \preceq p_\sigma b^\sigma$. Moreover, note that $\nu_\sigma / p_\sigma \in \mathcal{P}_\sigma \subseteq \Delta_M$, so $p_\sigma = \sum_\theta \nu_\sigma(\theta)$ holds for every feasible solution. This results in an equivalent linear program with variables $\nu_1, \ldots, \nu_t \in \mathbb{R}_+^M$, from which $p_\sigma = \sum_\theta \nu_\sigma(\theta)$ and $\mu_\sigma = \nu_\sigma / p_\sigma$ can be recovered efficiently.

Remarks

Zero-sum games When applied to two-player zero-sum games with the objective to maximize the row-player's payoff, our signaling scheme provides a stronger guarantee. In such settings, both players retain the same payoff in any exact Nash equilibrium. Also, any ϵ-equilibria give a payoff that is ϵ-close to the payoff of any exact equilibrium. Thus, the signaling scheme provided in Theorem 4.1 can be directly compared with the quality of the optimal signaling scheme without worrying about equilibrium selection, instead of a bi-criteria guarantee.

Reducing the number of signals Although the signaling scheme provided in Theorem 4.1 might use a quasipolynomial number of signals, we can reduce the number of signals to $M+1$. Let f_λ be the objective value of the signaling scheme, and consider the set of t signals used $\{\mu_1, \ldots, \mu_t\}$ and their corresponding objective values $\{f_1, \ldots, f_t\}$. Observe that the $(M+1)$-dimensional point (f_λ, λ) is a convex combination of the set of points $P = \{(w_1, \mu_1), \ldots, (w_t, \mu_t)\}$. Since f_λ is the objective value of the best signaling scheme that uses only the posteriors $\{\mu_1, \ldots, \mu_t\}$, (f_λ, λ) belongs to some facet of the convex hull of P. Hence by Carathéodory's theorem, (f_λ, λ) can be written as a convex combination of $M+1$ points from P, and such a decomposition can be computed in time polynomial in the size of P. This decomposition gives a valid signaling scheme with the same objective value, using only $M+1$ signals.

Extending the bicriteria guarantee Our algorithm can extend beyond exact Nash equilibria. For every $0 \leq \delta < \epsilon$, we can compute a signaling scheme with ϵ-equilibria that are competitive with the optimal signaling scheme that uses δ-equilibria. Formally, let $\mathrm{OPT}^\delta(\mathcal{A}, \lambda)$ denote the maximum reward for the principal over all possible signaling schemes and δ-equilibria. We can compute a signaling scheme φ and corresponding Bayesian $(\epsilon+\delta)$-equilibria $\mathcal{X}$ in time $\mathrm{poly}(M, n^{\log n/\epsilon^2})$, and the value of the signaling scheme satisfies $F(\varphi, \mathcal{X}) \geq \mathrm{OPT}^\delta(\mathcal{A}, \lambda) - \epsilon$. Theorem 4.1 is a special case of this result with $\delta = 0$.

Stackelberg games Our result can be extended to *Stackelberg games* which often arise in security games. Recall that in a Stackelberg game [91], one player (the leader) first commits to a (mixed) strategy, and then all other players (followers) simultaneously play their strategies upon learning the leader's strategy. Our result can be readily extended to Bayesian Stackelberg games when the objective of the signaling scheme is to maximize the leader's payoff. In this case, we can simply drop the constraints regarding the leader in the polytopes defined in Lemma 4.3, and only require the followers to play an approximate equilibrium in our algorithm presented in Theorem 4.1.

Equilibrium selection rules Our algorithm computes a signaling scheme as well as the associated (approximate) equilibria for the subgames. We assume that the principal can implement any equilibrium (i.e., the best equilibrium she can compute) in each subgame. It remains open whether one can find a near-optimal signaling scheme independent of equilibrium selection. For example, when the (real) players always choose the worst equilibrium for the principal after a signal is revealed.

4.4 Hardness Results for Signaling in Normal Form Games

In this section, we prove hardness results for approximately optimal signaling in normal form games.

In Section 4.4.1, we show that relaxing the incentive constraints is necessary if the principal's objective can be a tensor over the state of play (Theorem 4.5). In Section 4.4.2, we show that it is NP-hard to obtain an additive FPTAS (Theorem 4.7). In Section 4.4.3, we show that assuming the planted clique conjecture (Conjecture 2.6), there is no PTAS for signaling in zero-sum games (Theorem 4.10).

4.4.1 NP-hardness of Signaling with Exact Equilibria

Our bicriteria QPTAS in Theorem 4.1 allows the principal's payoff to depend on the specific strategies the players take. We show that the relaxation in players' incentive constraints is necessary if we want general objective functions, even for signaling in zero-sum games. More specifically, we show that it is NP-hard to distinguish whether the optimal signaling scheme has value 0 or at least 1/2.

Theorem 4.5. *Given a Bayesian zero-sum game* $\left(\Theta, \{A^\theta\}_{\theta\in\Theta}, \lambda\right)$ *and a principal objective tensor* A_0*, it is NP-hard to distinguish whether the optimal signaling scheme has value* 0 *or at least* $\frac{1}{2}$*.*

The NP-hardness proof uses a reduction from the *balanced vertex cover* (BVC) problem proposed by Conitzer and Sandholm [32]. In BVC, we are given a graph $G = (V, E)$, and we want to know if G has a vertex cover of size $\frac{|V|}{2}$. Given an instance of BVC with n nodes, we construct the following Bayesian zero-sum game where the states of nature correspond to nodes of G (i.e., $\Theta = V$) and the prior is uniform, i.e., $\lambda = \mathbb{1}_n/n$.

We use **Row** and **Col** to denote the row player and the column player respectively. **Row**'s pure strategies correspond to picking a node $u \in V$, and **Col**'s strategies correspond to either picking a vertex v, an edge e, or a special strategy s. The payoff of **Col** when she plays

$$v \text{ is } \begin{cases} \frac{n}{n-2} & \text{if } v \notin \{\theta, u\}, \\ 0 & \text{otherwise.} \end{cases} \quad e \text{ is } \begin{cases} \frac{n}{n-2} & \text{if } e \text{ is not incident with } \theta, \\ 0 & \text{otherwise.} \end{cases} \quad s \text{ is } 1.$$

The principal is only interested in getting **Col** to play the strategy s, that is, $A_0^\theta(v, s) = 1$ for all $\theta, v \in V$; all other entries of A_0 are 0.

The idea behind our construction is the following: nature and **Row** pick two nodes $\theta, u \in V$ to "protect", but only nature "protects" all the edges incident to θ. **Col** can choose to "attack" either a node v, an edge e, or to "give up" and play the strategy s. The principal wants **Col** to give up, so he has to coordinate the state of nature and **Row**'s strategy to protect different nodes. Because we do not relax the incentive constraints, the principal must find a vertex cover of size $n/2$. which becomes NP-hard.

Lemma 4.6. *The Bayesian zero-sum game defined above has a signaling scheme of value at least $\frac{1}{2}$ if and only if G has a vertex cover of size $\frac{n}{2}$.*

Proof. First, suppose G has a vertex cover C with $|C| = \frac{n}{2}$. The principal simply signals if $\theta \in C$ or not. That is, λ is decomposed as $(\mu_1 + \mu_2)/2$, where $\mu_1(v) = \frac{2}{n}$ for all $v \in C$ (and 0 otherwise), and $\mu_2(v) = \frac{2}{n}$ for all $v \notin C$. For posterior μ_1, there is a Nash equilibrium where **Row** picks $u \in V \setminus C$ uniformly at random and **Col** chooses strategy s; thus, the principal gets a value of 1. This is because every node is protected with probability $\frac{2}{n}$, and every edge is protected with probability at least $\frac{2}{n}$; so the payoff of **Col** for a pure strategy v or e is at most $\frac{n}{n-2}\left(1 - \frac{2}{n}\right) = 1$. Since μ_1 is announced with probability $\frac{1}{2}$, this signaling scheme achieves value at least $\frac{1}{2}$.

On the other hand, we show that if μ is a posterior with $f(\mu) > 0$, then G has a BVC solution. Recall that $f(\mu)$ is the principal's objective value under the posterior μ. Let (x, y) be a Nash equilibrium that attains value $f(\mu)$, that is, $f(\mu) = x^T \left(\sum_\theta \mu(\theta) A_0^\theta\right) y$. Since $f(\mu) > 0$, we must have $y_s > 0$, so every node in V must be protected with probability at least $\frac{2}{n}$. This requires nature and **Row** to be perfectly negatively correlated and never protect the same node.

Formally, we must have $\frac{n}{n-2}(1 - x(u))(1 - \mu(u)) \leq 1$ for all $u \in V$, and summing up over all nodes we have $\sum_u (1 - x(u))(1 - \mu(u)) \leq n - 2$, which implies

$\sum_u x(u)\mu(u) \le 0$. So it must be that for all $u \in V$, exactly one of $\mu(u)$ and $x(u)$ is equal to $\frac{2}{n}$, and this induces a natural bisection of the graph. Let $C = \{v : \mu_v > 0\}$; we know $|C| = \frac{n}{2}$. The payoff of **Col** for an edge $e = (s,t)$ is $\frac{n}{n-2}(1 - \mu_s - \mu_t)$, which must be at most 1, so we have $\mu_s + \mu_t \ge \frac{2}{n}$ for all edges (s,t). It follows that C is a vertex cover of G. □

In light of the hardness result with general objective functions (Theorem 4.5), for the rest of this chapter, we focus on signaling in two-player zero-sum games and the simplest principal's objective function — maximizing the row player's expected payoff at equilibrium.

In zero-sum games, both players retain the same payoff in any exact Nash equilibrium, and any ϵ-equilibrium gives a payoff that is ϵ-close to the playoff of any exact equilibrium. Thus, for signaling in zero-sum games with the objective to maximize the row-player's payoff, we can absorb the loss in the incentive constraints into the loss in objectives. Signaling schemes with bicriteria guarantees (e.g., our QPTAS in Section 4.3) can be directly compared with the quality of the optimal signaling scheme. In this setting, we can study the complexity of optimal signaling without worrying about equilibrium selection or bicriteria/single-criteria guarantees.

4.4.2 NP-hardness of an FPTAS

In this section, we prove signaling in normal form games does not admit an FPTAS unless P = NP (Theorem 4.7).

Theorem 4.7. *There is no FPTAS for the signaling problem, even for network security games, unless P = NP.*

Theorem 4.7 follows by considering the signaling problem from a *dual* perspective. The signaling problem can be written as a mathematical program (P) with linear objective and constraints, but an infinite number of variables. Ignoring this issue, we can consider the dual problem (D). Motivated by the separation problem for the dual, we consider the *posterior selection problem* (Definition 2.5).

Our key insight is that the posterior selection problem is a useful tool for deriving hardness results. This usefulness stems from the equivalence of separation and optimization [61], which shows that an algorithm for the separation problem can be used to solve the optimization problem and vice versa. We exploit and build upon this equivalence. We prove that this equivalence holds despite the infinite-dimensionality of (P), and furthermore, is approximation preserving: an FPTAS for signaling yields an FPTAS for the posterior selection problem (Theorem 4.8).

Whereas, typically, an (approximate) separation or membership oracle is used to (approximately) solve the optimization problem, we exploit this equivalence in an unorthodox fashion by leveraging the *hardness* of the optimization problem to prove *hardness* results for the membership problem. We show that it is NP-hard to obtain an FPTAS for the posterior selection problem in normal form games (Lemma 4.9), and thus it is NP-hard to obtain an FPTAS for optimal signaling in normal form games. It is worth noting that we obtain our NP-hardness result with minimal effort, a fact that underscores the benefits of moving to the posterior selection problem.

Theorem 4.7 follows immediately from Theorem 4.8 and Lemma 4.9.

Theorem 4.8. *An FPTAS for the signaling problem yields an FPTAS for the posterior selection problem.*

Proof. Recall that $f : \Delta_M \to [-1, 1]$ maps a posterior distribution to the principal's objective value, and f^+ is the concave envelope of f (Definition 2.1). Observe that f is decided by the Bayesian game, and does not depend on the prior distribution λ. The optimal signaling scheme has value $f^+(\lambda)$ for a given prior λ.

Consider the hypograph $\mathcal{P} \subseteq \mathbb{R}^{M+1}$ of f^+. An algorithm $\mathcal{B}$ for (approximately) optimal signaling gives a membership oracle for $\mathcal{P}$: a point $(\mu, \eta) \in \mathbb{R}^{M+1}$ belongs to $\mathcal{P}$ if and only if $\eta \leq f^+(\mu)$. For the posterior selection problem we want to compute $\max_\mu f(\mu) = \max_\mu f^+(\mu)$. Let $w = (0, \ldots, 0, 1) \in \mathbb{R}^{M+1}$. The posterior selection problem can be viewed as maximizing a linear function over $\mathcal{P}$: $\max_{(\mu,\eta)\in\mathcal{P}} \eta = \max_{x\in\mathcal{P}} w^T x$. At a high level, the theorem statement can be interpreted as *membership oracle is sufficient for optimization* [61, 81], and we need to quantify to what extent it is approximation preserving.

Formally, we show that if we have a polynomial time $\frac{\epsilon}{\text{poly}(M)}$-approximation algorithm $\mathcal{B}$ for optimal signaling, then we can convert it into a polynomial time ϵ-approximation algorithm for the posterior selection problem. Given a posterior selection instance and a precision parameter $\epsilon > 0$, we invoke part (ii) of Lemma 2.8 with $X = \Delta_M$, $h(\cdot) = f^+(\cdot)$, $K = 2$, and $\mathcal{B}$ as the imperfect value oracle with precision $\delta = \frac{\epsilon}{2p(M)}$. Note that $p(\cdot)$ is the polynomial given in Lemma 2.8, and f^+ is concave as needed. Lemma 2.8 will return a point $x^* \in \Delta_M$ in polynomial time with $f^+(x^*) \geq \max_{\mu\in\Delta_M} f^+(\mu) - 2p(M)\delta = \max_{\mu\in\Delta_M} f(\mu) - \epsilon$, an ϵ-optimal solution for the posterior selection problem. ∎

Lemma 4.9. *There is no FPTAS for the posterior selection problem, even for network security games, unless P = NP.*

Proof. The proof follows via a reduction from the *balanced complete bipartite subgraph* (BCBS) problem. In BCBS, given as input a bipartite graph $G = (V \cup W, E)$

and an integer $r \geq 0$, we want to determine if G contains an $r \times r$ biclique. Garey and Johnson [58] showed that the BCBS problem is NP-complete.

Given a BCBS instance, we set $\epsilon = \frac{1}{n^8}$ where $n = |V| + |W|$. We create a Bayesian network security game on G (defined in Section 2.1.2) and set $\rho = 2rn\epsilon$. This means that states of nature correspond to nodes of G ($\Theta = V \cup W$) and the payoff matrix for a distribution $\mu \in \Delta_\Theta$ is given by Equation (2.1) where B is the adjacency matrix of G and $\overline{A} = D^T = -\rho I_{n\times n}$. Intuitively, the principal and the row player want to be adjacent to each other, while at the same time they are forced to randomize because of the large penalty term ρ if the column player catches either of them. We show that solving this posterior selection instance to precision ϵ would decide the BCBS-instance.

We first show that if G has a $r \times r$ biclique $V' \times W'$, then there exists some μ with $f(\mu) \geq 1 - 2n\epsilon$. Set $\mu(v) = 1/r$ for all $v \in V'$ and $x(v) = 1/r$ for all $v \in W'$. Then, by Equation (2.2), we have $f(\mu) \geq x^T B\mu - \rho\|\mu + x\|_\infty = 1 - \rho/r = 1 - 2n\epsilon$, where $x^T B\mu = 1$ because V', W' form a biclique.

On the other hand, if there exists $\mu \in \Delta_M$ with $f(\mu) \geq 1 - (2n+2)\epsilon$, then G contains an $r \times r$ biclique. Let x be the row player's mixed strategy at equilibrium, so $f(\mu) = x^T B\mu - \rho\|\mu + x\|_\infty$. Let $V' \stackrel{\text{def}}{=} \{v \in V \cup W : \mu(v) \geq 1/n^3\}$ and $W' \stackrel{\text{def}}{=} \{v \in V \cup W : x(v) \geq 1/n^3\}$. Every vertex in V' must be adjacent to every vertex in W', otherwise $x^T B\mu \leq 1 - 1/n^6 < 1 - (2n+2)\epsilon$. Thus, V' and W' must be in different partitions and form a biclique. It remains to show that $|V'| \geq r$ and $|W'| \geq r$. By the definition of V' we have $\sum_{v \in V'} \mu(v) = 1 - \sum_{v \notin V'} \mu(v) > 1 - 1/n^2$. Since $\|\mu + x\|_\infty = \frac{x^T B\mu - f(\mu)}{\rho} \leq \frac{(2n+2)\epsilon}{\rho} = \frac{1+1/n}{r}$, $|V'| \geq \frac{\sum_{v \in V'} \mu(v)}{(1+1/n)/r} > r\frac{1-1/n^2}{(1+1/n)} = r(1 - 1/n)$. Hence $|V'| \geq r$, and similarly $|W'| \geq r$. □

To our best knowledge, it remains open whether Theorem 4.8 can be strengthened to show that an ϵ-approximation for signaling yields an $O(\epsilon)$-approximation for posterior selection, so that a PTAS for signaling yields a PTAS for posterior selection. We leave this as an intriguing open question (Problem 7.2). Below, we rule out a PTAS for signaling under an orthogonal hardness assumption.

4.4.3 Planted-Clique Hardness of a PTAS

In this section, we rule out a PTAS for signaling in normal form games assuming the planted-clique conjecture.

Theorem 4.10. *There is a constant $\epsilon_0 > 0$ such that, assuming the planted-clique conjecture (Conjecture 2.6), there is no ϵ_0-approximation for the signaling problem in Bayesian zero-sum games.*

We follow the intuition behind the proof of Lemma 4.9 and construct a game where the principal needs to identify dense subgraphs. The main differences are: (1) the error parameter ϵ_0 in this section is a constant, which is too large for detecting a few missing edges in a clique; and (2) Theorem 4.8 does not hold for translating PTAS hardness results. We resolve (1) by reducing from gap/promise problems (planted clique), where the densest large subgraphs either have density 1 or close to 1/2; and we handle (2) by using a "direct" reduction from the *planted clique cover* problem (Definition 2.7). Intuitively, a clique corresponds to a good posterior distribution; and optimal signaling decomposes the prior distribution into good posteriors, which corresponds to partitioning an input graph into dense subgraphs.

The proof of Theorem 4.10 combines and strengthens techniques from [6, 46]. The idea is to set up a Bayesian zero-sum game where both the principal and the row player must randomize over $\Omega(\log n)$-size high-density node sets for the signaling scheme to achieve large value; recovering these large-density sets from a near-optimal signaling scheme allows one to solve the planted-clique cover problem. Dughmi [46] used payoffs of magnitude $\Omega(\log^2 n)$ to enforce the above property, which is insufficient to rule out a PTAS. We instead leverage a device by Althöfer [6] to ensure the above "large-spreading" property. This device is used to show planted-clique hardness for computing the ϵ-best Nash equilibrium by Hazan and Krauthgamer [63]; and also used to show that ϵ-approximate equilibrium requires $\Omega(\log n)$ support size by Feder et al. [54] (both results are for constant ϵ).

One crucial technical issue is that we need to strengthen the planted-clique recovery result in [46]. To recover a specific planted clique S of size $k = \omega(\log^2 n)$ with high probability (in the presence of other such planted cliques), [46] requires a set $R \subseteq S$ with $|R| = \omega(\log^2 n)$, whereas we only require a set $R \subseteq S$ with $|R| = \Omega(\log n)$ (which is asymptotically tight). This is important because we can only ensure that spreading takes place over $\Omega(\log n)$-size sets.

We reduce from the planted clique cover problem with $k = k(n)$ such that $k = \omega(\log n)$ and $k = o(\sqrt{n})$, and $r = \frac{5n}{k}$, which we omit for the rest of the section. We use G^- and G^+ to denote the background edges and the clique edges added in steps (1) and (2) of Definition 2.7 respectively. Note that G^- and G^+ may contain the same edges. We use $\text{bi-density}_G(S,T)$ to denote the density of the bipartite graph $S \times T$ in G:

$$\text{bi-density}_G(S,T) \stackrel{\text{def}}{=} \frac{|\{(u,v) \in S \times T : \{u,v\} \in E\}|}{|S||T|}.$$

We require the planted clique instance to satisfy Lemma 4.11.

Lemma 4.11. *A graph $G \sim \mathcal{G}\Big(n, \frac{1}{2}, k, r\Big)$ with planted cliques $\{S_1, \ldots, S_r\}$ satisfies the following properties with high probability (for sufficiently large n),*

(i) All large bipartite subgraphs have density about $\frac{1}{2}$ before planting the cliques: For all $S, T \subseteq V$ with $|S|, |T| \geq c_2 \log n$, bi-density$_{G^-}(S,T) \leq \frac{1}{2} + \frac{1}{20}$.

(ii) Almost all nodes are in some clique: Let $\overline{S} \stackrel{\text{def}}{=} V \setminus \bigcup_i S_i$. We have $|\overline{S}| \leq e^{-4.9} n$.

(iii) All cliques are robustly recoverable: For every planted clique S_i and every subset $R \subseteq S_i$ with $|R| = c_3 \log n$, there is a polynomial time algorithm that recovers S_i from G given R.

Theorem 4.10 follows immediately from Lemmas 4.11 and 4.12.

Lemma 4.12. *Let $G \sim \mathcal{G}\Big(n, \frac{1}{2}, k, r\Big)$ be a planted clique cover instance that satisfies Lemma 4.11. There is a polynomial-time randomized reduction that takes the graph G as input and outputs a Bayesian zero-sum game such that the following hold with high probability.*

(Completeness) *There is a signaling scheme of value at least* 0.99*.*

(Soundness) *Given a signaling scheme of value at least* 0.97*, one can recover a constant fraction of the cliques planted in G.*

It is worth pointing out that the Bayesian zero-sum game we construct always admits a signaling scheme of large value; however *finding* a near-optimal signaling scheme in polynomial time would refute the planted-clique conjecture.

In the rest of this section, we prove Lemma 4.12. We use the following parameters.

$$\epsilon = 0.03, \quad Z = 20, \quad c_3 = 10^3, \quad c_2 = 10^5, \quad c_1 = c_2 \log(4Z/3) + 2, \quad N = n^{c_1}.$$

To keep the presentation simple, we give a construction where $A^\theta_{i,j} \in [-Z, Z]$ (as opposed to $[-1, 1]$). Let A_G denote the $(n \times n)$ adjacency matrix of $G = (V, E)$, and let A^-_G and A^+_G denote the adjacency matrices of G^- and G^+ respectively. We use **Row** and **Col** to denote the row and column players respectively. The states of nature and **Row**'s strategies correspond to the nodes of G. The prior λ is $\mathbb{1}_n/n$, i.e., each state of nature (each vertex) is equally likely to occur. For every $\theta \in \Theta = V$, the payoff matrix $A^\theta \in [-Z, Z]^{n\times(2N+1)}$ is given by $[a^\theta \quad B \quad \mathbb{1}_n(d^\theta)^T]$, which are defined as follows:

(1) a^θ is the θ-th column of the adjacency matrix A_G, so $a^\theta_i = 1$ if $(i, \theta) \in E$ and is 0 otherwise.

(2) B is an $n \times N$ matrix where each $B_{i,j}$ is set independently to $2 - Z$ with probability $\frac{3}{4Z}$, and 2 otherwise.

(3) $d^\theta \in [-Z, Z]^N$ is the θ-th row of B. Equivalently, if we use D to denote the $n \times N$ matrix having rows $(d^\theta)^T$ for $\theta \in \Theta$, we have $D = B$.

To gain some intuition, observe that for a posterior μ and **Row**'s mixed strategy x, the row vector $x^T A^\mu$ yielding **Col**'s payoffs is $[x^T A_G\mu \quad x^T B \quad \mu^T D]$. Thus, if **Col** plays action 1 (with probability 1), the expected payoff of **Row** is equal to $x^T A_G\mu$. If μ and x are uniform over $S, T \subseteq V$, the expected payoff is exactly bi-density$_G(S, T)$. The remaining $2N$ pure strategies of **Col** are used to force the principal and **Row** to choose a posterior μ and mixed strategy x respectively that are "well spread out".

The average of the entries in any column of B and D is $\frac{5}{4} > \max_i a^\theta_i$. Exploiting this, part (i) of Lemma 4.13 implies that if x and μ both randomize uniformly over a large set of vertices, **Col** plays column 1. The completeness proof follows from (roughly speaking) choosing posteriors and **Row**'s strategies that randomize uniformly over the planted cliques. For the completeness proof, if **Row**'s strategy x (in some subgame) has support of size at most $c_2 \log n$, then part (ii) of Lemma 4.13 implies that there exists a column of B that **Col** can play to make $f(\mu)$ negative. Similarly, for a posterior μ with small support, **Col** can play some column of D to make $f(\mu)$ negative. Thus, in order to obtain value close to 1, both μ and **Row** have to randomize over $\Omega(\log n)$-size sets of nodes. Using this, one can carefully extract a collection of node-sets which can then be used to recover the planted cliques.

Intuitively, B is used to force **Row** (x) to randomize over a large set, and D is used to force Nature (μ) to randomize. Formally, we prove the following lemma about the matrix B (and D).

Lemma 4.13. *For the $n \times N$ matrix $B_{i,j}$ where each $B_{i,j}$ is independently set to $2-Z$ with probability $\frac{3}{4Z}$ and set to 2 otherwise, the following hold with high probability.*

(i) Randomizing over a large set is good: For a fixed set $R \subseteq V$ with $|R| = \omega(\log n)$, we have $\frac{1}{|R|}\sum_{i\in R} B_{i,j} > 1$ for every $j \in [N]$.

(ii) Any distribution supported on a small set is bad: For every $R \subseteq V$ with $|R| \le c_2 \log n$, there exists some $j \in [N]$ such that $B_{i,j} = 2 - Z$ for all $i \in R$.

Proof. We first prove (i). The proof is a standard application of Chernoff bounds. Fix a column $j \in [N]$. We have $\mathbb{E}\left[\frac{\sum_{i\in R} B_{i,j}}{|R|}\right] = \frac{5}{4}$, where the expectation is over the random construction of B. Since $|R| = \omega(\log n)$, the size of R is large enough so that Chernoff bounds imply that $\Pr\left[\frac{\sum_{i\in R} B_{i,j}}{|R|} < \frac{9}{8}\right] \le \frac{1}{2N\,\mathrm{poly}(n)}$. The union bound over all N columns yields the claim.

We now prove (ii). It is sufficient to show the claim for all $R \subseteq V$ with $|R| = c_2 \log n$. Fix some R with $|R| = c_2 \log n$. For a given $j \in [N]$, we have $\Pr[\exists i \in R \text{ s.t. } B_{i,j} \neq 2 - Z] = 1 - \left(\frac{3}{4Z}\right)^{|R|}$. So

$$\Pr\left[\forall j \in [N], \exists i \in R \text{ s.t. } B_{i,j} \neq 2 - Z\right] = \left[1 - \left(\frac{3}{4Z}\right)^{|R|}\right]^N$$
$$\leq \exp\left(-N\left(\frac{3}{4Z}\right)^{|R|}\right) = \exp(-n^2).$$

In other words, the probability that B fails to "catch" a small set R is exponentially small. Taking the union bound over all $R \subseteq V$ with $|R| = c_2 \log n$, we obtain

$$\Pr\Big[\exists R \subseteq V \text{ with } |R| = c_2 \log n \text{ s.t. no } j \in [N] \text{ satisfies } B_{i,j} = 2 - Z \text{ for all } i \in R\Big]$$
$$\leq \binom{n}{c_2 \log n} \exp(-n^2) \leq \exp\Big(c_2 \log^2 n - n^2\Big) \leq \frac{1}{\exp(n)}. \qquad \square$$

Therefore, the probability that there exists a $j \in [N]$ to "catch" every R with $|R| = c_2 \log n$ is at least $1 - \frac{1}{\exp(n)}$ as claimed.

Completeness proof in Lemma 4.12

We use a deterministic signaling scheme that groups together states of nature in the same planted clique. We first partition the graph into disjoint large cliques and a small number of remaining nodes. Let $S'_i = S_i \setminus \bigcup_{1 \leq j < i} S_j$ for $i \in [r]$ be the set of vertices in S_i that do not appear in earlier cliques. Define $\overline{S} \stackrel{\text{def}}{=} V \setminus \bigcup_j S_j$ as the remaining vertices. Finally, let $S'_0 \stackrel{\text{def}}{=} \overline{S} \cup \left\{v \in S'_i : |S'_i| < \frac{k}{10^4}\right\}$.

Our signaling scheme is (Σ, p, μ) where the set of signals is $\Sigma = \{0\} \cup \left\{i \in [r] : |S'_i| \geq \frac{k}{10^4}\right\}$. For each signal σ, $p_\sigma = \frac{|S'_\sigma|}{n}$ and μ_σ is the uniform distribution over S'_σ. For posterior μ_σ, where $\sigma \neq 0$, consider the strategy x_σ where **Row** plays the uniform distribution on S'_σ. Part (i) of Lemma 4.13 implies that **Col**'s best response to x_σ is to play column 1, with high probability over the randomness in B. Therefore, $f(\mu_\sigma) \geq \text{bi-density}(S'_\sigma, S'_\sigma) = 1 - \frac{1}{|S'_\sigma|} \geq 1 - \frac{10^4}{k}$.

Recall that G is a good planted clique instance, thus part (ii) of Lemma 4.11 guarantees that $|\overline{S}| \leq e^{-4.9}n$. With $r = \frac{5n}{k}$, we have $|S'_0| \leq |\overline{S}| + \frac{5n}{k} \cdot \frac{k}{10^4} \leq e^{-4.7}n$. So the signaling scheme has value at least

$$\sum_{\sigma \in (\Sigma \cap [r])} p_\sigma f(\mu_\sigma) \geq \sum_{\sigma \in (\Sigma \cap [r])} p_\sigma \Big(1 - \frac{10^4}{k}\Big) \geq (1 - e^{-4.7})\Big(1 - \frac{10^4}{k}\Big) \geq 0.99.$$

Soundness proof in Lemma 4.12

For a signal $\sigma \in \Sigma$ with corresponding posterior μ_σ, let x_σ denote **Row**'s equilibrium strategy for A^{μ_σ}. We first filter out the set of "useful" signals with relatively high values. Let $\Sigma_1 = \{\sigma \in \Sigma : f(\mu_\sigma) \geq 1 - \sqrt{\epsilon}\}$. The value of the signaling scheme is $\sum_{\sigma \in \Sigma} p_\sigma f(\mu_\sigma) \geq 1 - \epsilon$. Noting that $f(\mu) \leq 1$ for all μ, by a simple counting argument, we have $p_{\Sigma_1} \geq 1 - \sqrt{\epsilon}$. We show that for all $\sigma \in \Sigma_1$, μ_σ and x_σ place a significant mass over a large set of nodes, and use this insight to extract clusters.

Recall that $\epsilon = 0.03$ and $Z = 20$. For every signal $\sigma \in \Sigma_1$, define $T_\sigma = \Big\{i : e_i^T A_G \mu_\sigma \geq 1 - \frac{Z\sqrt{\epsilon}}{Z-2}\Big\}$, and let $\widetilde{x}_\sigma$ be the uniform distribution on T_σ. Intuitively, T_σ is the set of good strategies for **Row** under the signal σ. Let $\mathcal{T} = \{T_\sigma : \sigma \in \Sigma_1\}$ denote the collection of these node-sets. As we shall see, $\mathcal{T}$ is going to play an important role for recovering a constant fraction of the planted cliques.

Fix $\sigma \in \Sigma_1$ with $1 - \sqrt{\epsilon} \leq f(\mu_\sigma) \leq 1$. We first show that T_σ cannot be too small. Otherwise, **Col** can punish **Row** for concentrating on a few strategies. By the definition of T_σ, $f(\mu_\sigma) \geq 1 - \sqrt{\epsilon}$ and a simple counting argument, we have $x_\sigma(T_\sigma) \geq \frac{2}{Z}$. It follows that $|T_\sigma| > c_2 \log n$, because every $R \subseteq V$ with $|R| \leq c_2 \log n$ must satisfy $x_\sigma(R) < \frac{2}{Z}$. Otherwise, suppose $x_\sigma(R) \geq \frac{2}{Z}$, then by part (ii) of Lemma 4.13, there exists a column j of B having $B_{i,j} = 2 - Z$ for all $i \in R$. We have $\sum_{i \in [n]}(x_\sigma(i))B_{i,j} \leq (2 - Z)x_\sigma(R) + 2\Big(1 - x_\sigma(R)\Big) \leq 0$, which implies that $f(\mu_\sigma) \leq 0$, a contradiction.

We now switch from x_σ to $\widetilde{x}_\sigma$ in order to relate the value of the signaling scheme to bi-density. As before, G^- are the background edges and G^+ are the clique edges, and A_G^- and A_G^+ are the corresponding adjacency matrices. Let A_G^i be the adjacency matrix of the clique S_i. Note that $A_G \leq A_G^- + A_G^+ \leq A_G^- + \sum_{i=1}^r A_G^i$ (since the planted cliques contain existing edges of G^- and may overlap).

Let R denote the $c_2 \log n$ largest entries in $\widetilde{x}_\sigma^T A_G^-$, and let $\widetilde{\mu}_\sigma$ be the uniform distribution on R. Recall that G is a good planted clique instance that satisfies Lemma 4.11. Since both T_σ and R have size at least $c_2 \log n$, and $\widetilde{\mu}_\sigma$ and $\widetilde{x}_\sigma$ are uniform distributions over them, part (i) of Lemma 4.11 guarantees that $\widetilde{x}_\sigma^T A_G^- \widetilde{\mu}_\sigma = \text{bi-density}(T_\sigma, R) \leq \frac{11}{20}$. Moreover, we have $\mu_\sigma(R) < \frac{2}{Z} = \frac{1}{10}$ because otherwise **Col** can choose a strategy in D to punish the concentration in μ_σ. Since the maximum entry of $\widetilde{x}_\sigma^T A_G^-$ outside of R is at most the average entry in R, we have $\widetilde{x}_\sigma^T A_G^- \mu_\sigma \leq \frac{1}{10} + \frac{9}{10} \cdot \widetilde{x}_\sigma^T A_G^- \widetilde{\mu}_\sigma \leq \frac{1}{10} + \frac{9}{10} \cdot \frac{11}{20} < \frac{3}{5}$. This tells us that the background density does not contribute enough to make $f(\mu_\sigma)$ close to 1. Therefore, the reason behind the high value of μ_σ must be that $T_\sigma \times R$ overlaps with some of the planted cliques.

Because $\sum_{\sigma\in\Sigma_1} p_\sigma(\widetilde{x}_\sigma^T A_G \mu_\sigma) \geq (1-\sqrt{\epsilon})\Big(1 - \frac{Z\sqrt{\epsilon}}{Z-2}\Big) > \frac{2}{3}$ (substituting in our choice of $\epsilon = 0.03$ and $Z = 20$), we have

$$\begin{aligned}
\frac{1}{15} = \frac{2}{3} - \frac{3}{5} &< \sum_{\sigma\in\Sigma_1} p_\sigma(\widetilde{x}_\sigma^T A_G \mu_\sigma) - \max_{\sigma\in\Sigma_1} \widetilde{x}_\sigma^T A_G^- \mu_\sigma \\
&\leq \sum_{\sigma\in\Sigma_1} p_\sigma \widetilde{x}_\sigma^T (A_G - A_G^-)\mu_\sigma \\
&\leq \sum_{\sigma\in\Sigma_1} p_\sigma \sum_{i=1}^{r} \widetilde{x}_\sigma^T A_G^i \mu_\sigma \\
&= \sum_{i=1}^{r} \sum_{\sigma\in\Sigma_1} p_\sigma \mu_\sigma(S_i) \frac{|T_\sigma \cap S_i|}{|T_\sigma|} \\
&\leq \sum_{i=1}^{r} \Big(\sum_{\sigma\in\Sigma_1} p_\sigma \mu_\sigma(S_i) \Big) \left(\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|} \right) \\
&\overset{(*)}{\leq} \sum_{i=1}^{r} \frac{|S_i|}{n} \left(\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|} \right) \\
&= \frac{5}{r} \sum_{i=1}^{r} \left(\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|} \right).
\end{aligned}$$

Inequality $(*)$ follows since for every $v \in \Theta$, we have $\sum_{\sigma\in\Sigma_1} p_\sigma(\mu_\sigma)_v$ is at most $\sum_{\sigma\in\Sigma} p_\sigma(\mu_\sigma)_v = \lambda_v = \frac{1}{n}$.

We have $\frac{1}{r}\sum_{i=1}^{r}\Big(\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|}\Big) \geq \frac{1}{75}$. By a simple counting argument, at least a $\frac{1}{297}$-fraction of $S_1, \ldots, S_r$ satisfy $\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|} \geq \frac{1}{100}$. Therefore, to recover a constant fraction of the cliques, it is sufficient to show that we can recover any $S = S_i$ from a set T with $|T| \geq c_2 \log n$ and $|S \cap T| \geq \frac{|T|}{100} \geq c_3 \log n$. Observe that we can assume without loss of generality that $|T| \leq 2c_2 \log n$. Otherwise we can partition T into disjoint subsets of size between $c_2 \log n$ and $2c_2 \log n$, and one of these subsets T' would have $\frac{|S\cap T'|}{|T'|} \geq \frac{|S\cap T|}{|T|}$. Part (iii) of Lemma 4.11 states that we can, in polynomial time, recover any planted clique S_i given an arbitrary subset $R \subseteq S_i$ with size $|R| \geq c_3 \log n$. With $|T| \leq 2c_2 \log n$ and $|S \cap T| \geq c_3 \log n$, we can simply enumerate all subsets $R \subseteq T$ of size $c_3 \log n$ and run the clique recovery algorithm in

Lemma 4.11. Moreover, the enumeration can be done in time $\binom{2c_2 \log n}{c_3 \log n} = \text{poly}(n)$. So by iterating through every $T \in \mathcal{T}$, partitioning and running the clique recovery algorithm, we can recover all the S_i's that satisfy $\max_{T\in\mathcal{T}} \frac{|T\cap S_i|}{|T|} \geq \frac{1}{100}$, which is at least a constant fraction of all the planted cliques. This concludes the soundness proof.

A Tighter Amplification Lemma for Planted Clique

The rest of this chapter is devoted to prove Lemma 4.11. Lemma 4.11 gives three properties that a planted clique cover instance $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ should satisfy with high probability. Recall that $k = k(n)$ satisfies $k = \omega(\log n)$ and $k = o(\sqrt{n})$, $r = \Theta(n/k)$, and $c_3 = 10^3$. As before, we use G^- to denote the background edges and G^+ to denote the clique edges (see Definition 2.7).

For part (i), we need to show that, with high probability, all large bipartite subgraphs have density close to $\frac{1}{2}$ in $G^- \sim \mathcal{G}(n, \frac{1}{2})$. This is a direct corollary of Lemma 4.14 with $c = c_2 = 10^5$ and $\epsilon = 0.1$. Lemma 4.14 follows from a standard application of the Chernoff bound and the union bound.

Lemma 4.14 (Proposition B.2 in [46] quantified)**.** *Let* $0 < \epsilon < 1$ *and* $c \geq 50 \cdot \frac{1+\epsilon}{\epsilon^2}$. *For all* $n \geq 2$*, we have*

$$\Pr\Big[\exists S, T \subseteq V \textit{ with } |S|, |T| \geq c \ln n,\ \textit{bi-density}_{G^-}(S,T) > \frac{1+\epsilon}{2}\Big] \leq \frac{1}{n^4}.$$

Part (ii) claims that all except a small constant fraction of the nodes are covered by some clique. Recall that $r = \frac{5n}{k}$, and $\overline{S} \overset{\text{def}}{=} V \setminus \bigcup_j S_j$ is the set of uncovered nodes. We have $\mathbb{E}\Big[|\overline{S}|\Big] = n \cdot \Pr\Big[v \in \overline{S}\Big] = n(1 - \frac{k}{n})^r \leq e^{-5}n$, and $|S| \leq$

$e^{0.1} \cdot \mathbb{E}[|S|] \leq e^{-4.9}n$ with high probability due to standard Chernoff bounds (since the events $\{v \in |S|\}_{v \in V}$ are negatively correlated).

Our main technical contribution in Lemma 4.11 is part (iii). The claim is the following: For $G \sim \mathcal{G}(n, p, k, r)$ with planted cliques $\{S_1, \ldots, S_r\}$, with high probabiltiy over the randomness in G, every planted clique S_i can be recovered in polynomial time given an arbitrary subset $R \subseteq S_i$ with $|R| \geq c_3 \ln n$.

It is well known that for a planted clique instance $G \sim \mathcal{G}(n, \frac{1}{2}, k)$, with high probability over the randomness in G, one can recover the planted clique S given $R \subseteq S$ with $|R| \geq c_3 \ln n$. We generalize this result and show that, despite the presence of $\Theta(\frac{n}{k})$ other planted cliques, every clique S_i can still be recovered from $c_3 \ln n$ nodes.[4] It is important that our recovery algorithm works for any $R \subseteq S$ (rather than with high probability for a fixed R). This is because the set R is obtained from a near-optimal signaling scheme φ. Since φ is produced by an algorithm *after* examining the planted clique cover instance G, the choice of R can depend on the realization of G.

Fix some $i \in [r]$ and let $S = S_i$. We use the following algorithm to recover S:

(1) Let S' be all the common neighbors of R.

(2) Let $\widehat{S}$ be the vertices in S' with at least $k - 1$ neighbors in S'.

Part (iii) follows immediately from Lemmas 4.15 and 4.16, and a union bound over all $i \in [r]$. We use E, E^- and E^+ to denote the edges of G, G^- and G^+ respectively.

---

[4] We tighten the recovery algorithm in [46], which requires $|R| = \omega(\log^2 n)$ nodes from the planted clique. The difference in the magnitude of $|R|$ poses certain challenges and necessitates some key changes to the analysis in [46].

Lemma 4.15. *With high probability, we have* $|S'| \leq |S| + c_3 \ln n$*.*

Lemma 4.16. *With high probability, for all* $v \notin S$ *we have* $|E(v, S)| \leq 0.7|S|$*.*

We first elaborate how they imply part (iii) of Lemma 4.11. Observe that all the nodes in S will survive Step (1) and (2), so $S \subseteq \widehat{S}$. We show that, with high probability, no other vertices survive Step (1) and (2). Lemma 4.15 states that at most $c_3 \ln n$ nodes outside of S survive Step (1). Lemma 4.16, together with the assumption that $|S| = \omega(\log n)$, implies that $E(v, S') \leq E(v, S) + (|S'| - |S|) \leq 0.7|S| + c_3 \ln n < 0.71|S| < k - 1$ for all $v \in S' \setminus S$ (and for sufficiently large n). Therefore, all nodes in $S' \setminus S$ gets filtered in Step (2), and we have $\widehat{S} = S$ as claimed.

Before we continue to prove Lemmas 4.15 and 4.16, we state the following lemma which is crucial for our analysis. Lemma 4.17 is the main reason why our clique recovery result is asymptotically tight and better than that of [46].

Lemma 4.17. *With high probability, we have* $|E^+(v, S)| \leq 12 \ln n$ *for all* $v \notin S$*.*

Proof of Lemma 4.15

We first look at bad vertices in S' that is due to the background edges. Let $A = \{v \notin R : |E^-(v, R)| \geq 0.8|R|\}$. Then, bi-density$_{G^-}(R, A) \geq 0.8$. Because this density is much higher than $\frac{1}{2}$, it cannot be the case that both A and R are large. Formally, Lemma 4.14 holds for $c = c_3 = 10^3$ and $\epsilon = 0.6$ with high probability. Since $|R| \geq c_3 \ln n$, we have $|A| < c_3 \ln n$.

We next show that the clique edges do not introduce extra vertices to S'. By Claim 4.17, for every $v \notin S$, the clique edges will not increase $E(v, S)$ too much. For all but at most $c_3 \ln n$ nodes $v \notin S$, we have

$$|E(v,R)| = |E^-(v,R)|+|E^+(v,R)| \leq 0.8|R|+|E^+(v,S)| \leq 0.8|R|+12\ln n \leq 0.82|R|.$$

Hence, with high probability, at most $c_3 \ln n$ nodes outside of S survive Step (1).

Proof of Lemma 4.16

We prove the claim for a fixed $v \notin S$ and then take a union bound over all $v \notin S$. We upper bound $|E(v, S)|$ by inspecting the edges in E^- and E^+ separately. With high probability, we have $|E^-(v, S)| \leq 0.6|S|$ for all $v \notin S$. This is a standard application of the Chernoff bound and the union bound, since $|S| = \omega(\log n)$.

For edges in $E^+(v, S)$ we use Lemma 4.17. With high probability, for all $v \notin S$, we have $|E^+(v, S)| \leq 12 \ln n = o(|S|)$, and therefore for all $v \notin S$,

$$|E(v,S)| = |E^-(v,S)| + |E^+(v,S)| \leq 0.6|S| + o(|S|) \leq 0.7|S|.$$

Proof of Lemma 4.17

Lemma 4.17 bounds the number of edges between a set S (of size k) and a vertex $v \notin S$, when the graph is exactly the union of $\Theta(\frac{n}{k})$ random k-cliques. The expected number of times v gets covered by these cliques is $\Theta(1)$, so with high probability, v is covered $O(\log n)$ times. On the other hand, the expected size of the overlap between some clique S_i and S is $\frac{k^2}{n} = o(1)$, so with high probability, every

S_i overlaps with S on $O(\log n)$ vertices. If we simply combined these two bounds (as in [46]), we get a weaker version of Lemma 4.17 with $E^+(v,S) \le O(\log^2 n)$.

The key observation (by David Kempe) is that we can use the principle of deferred decisions to improve this analysis. We ask all the cliques first to decide whether to include v or not, and defer their choices on other nodes. Roughly speaking, there are $O(\log n)$ cliques that include v, and they contain $O(k \log n)$ random vertices; The expected total size of the overlap between these cliques and S is at most $O(k \log n \cdot \frac{k}{n}) = o(1)$, so tail bounds and the union bound imply that $E^+(v,S) \le O(\log n)$.

The following lemma will be useful in proving Lemma 4.17. Intuitively it says that, to upper bound $E^+(v,S)$, we can pretend the $O(k \log n)$ vertices (from the $O(\log n)$ cliques that contain v) are chosen independently and uniformly at random.

Lemma 4.18 (see Ex. 1.13 in [45], Lemma 1.19 in [44])**.** *Let $X_1, \ldots, X_n$ be arbitrary binary random variables. Suppose for every i, and every $x_1, \ldots, x_{i-1} \in \{0,1\}$, we have $\Pr[X_i = 1 \,|\, X_1 = x_1, X_2 = x_2, \ldots, X_{i-1} = x_{i-1}] \le p_i$. Let $Y_1, \ldots, Y_n$ be independent binary random variables with $\Pr[Y_i = 1] = p_i$ for all $i \in [n]$. Then, for any M, we can upper bound $\Pr[\sum_{i=1}^n X_i > M]$ using the upper-tail Chernoff bound for $\Pr[\sum_{i=1}^n Y_i > M]$.*

In particular, for any $\epsilon \in (0,1)$ and $\mu \ge \sum_{i=1}^n p_i$, we have $\Pr[\sum_{i=1}^n X_i > (1+\epsilon)\mu] \le e^{-\epsilon^2 \mu/3}$.

Proof of Lemma 4.17. Fix $v \notin S$, and let X denote the random variable $|E^+(v,S)|$. Let $S_1, \ldots, S_{r-1}$ be the planted cliques other than S. Let I be the random index-set of cliques that contain v; that is, $I \subseteq [r-1]$ is such that $v \in S_i$ for all $i \in I$, and

$v \notin S_i$ for all $i \notin I$. Notice that the events $\{i \in I\}$ for $i \in [r-1]$ are independent Bernoulli trials with probability $\frac{k}{n}$. So we have $\Pr[|I| > 6\log n] \leq \frac{1}{n^2}$.

Fix an index set $J \subseteq [r-1]$ with $|J| \leq 6\log n$ and consider $\Pr[X > 12\log n \mid I = J]$. We use $\Pr'$ to denote probabilities conditioned on the event $I = J$. Conditioned on $I = J$, we have $X \leq \sum_{i \in J, u \in S} Y_{i,u}$, where $Y_{i,u}$ is the random variable indicating if $u \in S_i$. Fix an ordering of the $Y_{i,u}$ random variables. If we consider the random variable $Y_{i,u}$, and any realization σ of the random variables appearing before $Y_{i,u}$, we have $\Pr'[Y_{i,u} = 1 \mid \text{realization } \sigma \text{ of the variables before } Y_{i,u}] \leq \frac{k}{n}$. Since $\frac{|J|k^2}{n} < 6\log n$, we can now use Lemma 4.18 and infer that $\Pr'[X > 12\log n] \leq e^{-\frac{6\log n}{3}}$.

Finally, we have

$$\begin{aligned}
\Pr[X > 12\log n] &= \sum_{J \subseteq [r-1]} \Pr[I = J] \cdot \Pr[X > 12\log n \mid I = J] \\
&\leq \sum_{\substack{J \subseteq [r-1]: \\ |J| > 6\log n}} \Pr[I = J] + \sum_{\substack{J \subseteq [r-1]: \\ |J| \leq 6\log n}} \Pr[I = J] \cdot \Pr[X > 12\log n \mid I = J] \\
&\leq \Pr[|I| > 6\log n] + \sum_{\substack{J \subseteq [r-1]: \\ |J| \leq 6\log n}} \Pr[I = J] \cdot \frac{1}{n^2} \leq \frac{2}{n^2}.
\end{aligned}$$

□

Chapter 5

Mixture Selection: An Algorithmic Framework

In Chapter 3 and Chapter 4, we settled the computational complexity of optimal signaling in Bayesian network routing games and normal form games. The best approximation algorithm we had for these two classes of games are fairly different. In network routing games we simply reveal full information; and in normal form games we use one signal for each approximate equilibrium and solve a quasipolynomial size linear program. There are many other interesting game-theoretic applications that involves the design of information structures. Do we need to come up with a different approximation algorithm for every new class of game we encounter, or is there a building block that many of these signaling problems have in common?

In this chapter, we identify two parameters that seem to dictate the complexity of optimal signaling, and present an algorithmic framework that (approximately) solves the optimal signaling problem in a number of different Bayesian games. We pose and study a fundamental algorithmic problem which we term *mixture selection*, a problem that arises naturally in the design of optimal information structures:

Definition 5.1 (Mixture Selection)**.** *For a function $g : [-1, 1]^n \to [-1, 1]$ and a positive integer M, M-dimensional mixture selection for g is the following optimization problem: Given an $n \times M$ matrix A with entries in $[-1, 1]$, find x in the M-dimensional simplex Δ_M maximizing $f(x) \stackrel{\text{def}}{=} g(Ax)$.*

The mixture selection problem is closely related to the posterior selection problem (Definition 2.5) and the optimal signaling problem. At a high level, mixture selection and posterior selection ask for the best posterior distribution for the principal's objective, while the optimal signaling problem asks for the optimal way to decompose a prior into "good" posteriors that maximizes the principal's expected objective value.

Recall that $f(\mu)$ denotes the principal's objective value under the posterior μ. To connect mixture selection to the posterior selection problem, we consider signaling problems where $f(\mu)$ can be written as $g(A\mu)$ for a function $g : [-1, 1]^n \to [-1, 1]$ and a matrix $A \in [-1, 1]^{n \times M}$. The mixture selection problem is more general, and it captures the posterior selection problem with a fixed g and arbitrary A. Mixture selection focuses on how the complexity of g affects the complexity of posterior selection. In Section 4.4.2, we have already seen that the posterior selection problem can be useful when proving hardness results of signaling. In this chapter, we present a meta-algorithm that works for both signaling and posterior selection, where the running time of our algorithm depends only on the "complexity" of g.

The work presented in this chapter appeared in [26].

5.1 Summary of Results

We investigate how the complexity of mixture selection (and optimal signaling) depends on the complexity of the function g. We identify two "smoothness" parameters of the function g which tightly control the complexity of mixture selection. The first smoothness quantity is a familiar one, namely Lipschitz continuity in the L^∞ metric. The second quantity, which we define and term *noise stability*

(Definition 5.3), borrows ideas from related definitions of stability in other contexts (e.g., [68, 79]), though it is importantly different. Informally, noise stability controls the degree to which low-probability — and possibly correlated — errors in the inputs of g can impact its output.

The approximation guarantee of our algorithm degrades gracefully as a function of the Lipschitz continuity and noise stability of g (Theorem 5.6). Moreover, the same conditions — noise stability and Lipschitz continuity — on the function g also lead to a similar approximation scheme for the corresponding signaling problem (Theorem 5.10). In particular, when g is both $O(1)$-Lipschitz continuous and $O(1)$-stable, we obtain an (additive) polynomial-time approximation scheme (PTAS) for mixture selection and optimal signaling. We also show that neither assumption suffices by itself for an additive PTAS (Theorems 5.18 and 5.19).

Our results for mixture selection can be viewed as generalizing the main insights of Lipton et al. [74]. First, we show that when g is noise stable and Lipschitz continuous, and $x \in \Delta_M$ is arbitrary, there is a sparse vector $\widetilde{x}$ for which $g(A\widetilde{x})$ is not much smaller than $g(Ax)$. The proof of this fact proceeds by sampling from x and letting $\widetilde{x}$ be the empirical distribution, as in [74]. However, when g is sufficiently noise stable and Lipschitz continuous, we obtain a better tradeoff between the number of samples required and the error introduced into the objective than does [74], and this is crucial for our applications. Our analysis bounds the expected difference between $g(Ax)$ and $g(A\widetilde{x})$ as the sum of two terms: The first term represents the error in the output of g caused by the low-probability "large errors" in its n inputs, and the second term represents the error in the output of g introduced by the higher-probability "small errors" in its n inputs. The first term is bounded using noise stability, and the second is bounded using Lipschitz continuity.

Second, we instantiate the above insight algorithmically, as does [74]. Specifically, our algorithm enumerates vectors $\widetilde{x}$ of the desired sparsity in order to find an approximately optimal solution to our mixture selection problem. We note that our guarantees are all parametrized by the Lipschitz continuity c and the noise stability β of the function g. Most notably, we obtain an additive polynomial-time approximation scheme (PTAS) whenever both β and c are constants.

Despite the simplicity of our framework, we find that it has powerful implications for optimal signaling in games. Notably, we find that we resolve or make progress on a number of known open problems, and some new ones, using one unified algorithmic framework.

1. **Optimal signaling in Bayesian normal form games (defined in Section 2.1.2)**: In Section 5.3.1, we derive a new QPTAS for this problem using the mixture selection framework. We use the fact that every function is $O(n)$-stable, and the fact that the function measuring the quality of equilibria satisfies a bi-criteria notion of Lipschitz continuity which we define.

2. **Revenue-maximizing signaling in probabilistic second-price auctions (defined in Section 2.1.5)**: A PTAS for this problem follows easily from our framework. We use the fact that the function $\max2(\cdot)$, the second largest entry of a vector, is Lipschitz continuous and noise stable.

3. **Persuasion in voting (defined in Section 2.1.6)**: We design a multi-criteria PTAS for this problem using our framework, using the fact that the function $g^{(\text{vote-sum})}(t) = \frac{1}{n}|\{i : t_i \geq 0\}|$ is noise stable and Lipschitz continuous in a bi-criteria sense.

We present the results for auctions and voting together in the next chapter (Chapter 6), where we systematically explore optimal signaling in anonymous games.

5.2 Noise Stability and Lipschitz Continuity

In this section, we present our notion of *noise stability*, and derive approximation algorithms for this problem when the function g is simultaneously noise stable and Lipschitz continuous with respect to the L^∞ metric.

Our approximation guarantees will be additive — i.e., an ϵ-approximation algorithm for mixture selection outputs $x \in \Delta_M$ with $f(x) \geq \max_{y \in \Delta_M} f(y) - \epsilon$. To illustrate our techniques, we use the following function $g^{(\text{mid})} : [-1,1]^n \to [-1,1]$, which averages all but the top and bottom quartiles of its inputs, as a running example.

$$g^{(\text{mid})}(t) = \frac{1}{\lceil 3n/4 \rceil - \lfloor n/4 \rfloor} \sum_{i=\lfloor n/4 \rfloor + 1}^{\lceil 3n/4 \rceil} t_{[i]},$$

where $t_{[i]}$ denotes the i^{th} largest entry of t. Throughout this chapter, we use t_i to denote the i^{th} entry of t, and use $t_{[i]}$ to denote the i^{th} largest entry of t.

Though we present our framework for functions $g : [-1,1]^n \to [-1,1]$, we define mixture selection similarly for functions $g : [0,1]^n \to [0,1]$. The two definitions are equivalent up to normalization, and it is easy to verify that all our results and bounds for mixture selection carry through unchanged to either definition.

Our main result applies to functions g which are both noise stable and Lipschitz continuous with respect to the L^∞ metric. We now formalize these two conditions.

Lipschitz Continuity

A function $g : [-1,1]^n \to [-1,1]$ is *c-Lipschitz continuous in* L^∞ — or c-Lipschitz for short — if and only if for all t, t' in the domain of g, $|g(t) - g(t')| \leq c||t - t'||_\infty$. To illustrate, our example function $g^{(\mathrm{mid})}$ is 1-Lipschitz. We note that Lipschitz continuity in L^∞ is a stronger assumption than in any other L^p norm.

Noise Stability

Our notion of noise stability captures the following desirable property of a function $g : [-1,1]^n \to [-1,1]$: if a random process corrupts (i.e., modifies arbitrarily) some of the inputs to g, with no individual input disproportionately likely to be corrupted, then the output of g does not decrease by much in expectation. Such random corruption patterns are captured by our notion of a *light distribution* over subsets of $[n]$, defined below.

Definition 5.2 (Light Distribution)**.** *Let $\mathcal{D}$ be a distribution supported on subsets of $[n]$. For $\alpha \in (0,1]$, we say $\mathcal{D}$ is α-light if and only if* $\Pr_{R \sim \mathcal{D}}[i \in R] \leq \alpha$ *for all $i \in [n]$.*

In other words, a light distribution bounds the *marginal probability* of any individual element of $[n]$. When corrupted inputs follow a light distribution, no individual input is too likely to be corrupted. However, we note that our notion of light distribution allows arbitrary correlations between the corruption events of various inputs. We define a noise stable function as one which is robust, in an average sense, to corrupting a subset R of its n inputs when R follows a light distribution $\mathcal{D}$. Our notion of robustness is one-sided: we only require that our function's output not decrease substantially in expectation. This one-sided guarantee suffices for all our applications, and is necessitated by some. We note that the light distribution $\mathcal{D}$, as well as the (corrupted) inputs, are chosen adversarially. We make use of the following notation in our definition: Given vectors $t, t' \in [-1,1]^n$ and a set $R \subseteq [n]$, we say $t' \underset{R}{\approx} t$ if $t_i = t'_i$ for all $i \notin R$. In other words, if $t' \underset{R}{\approx} t$, then t' is a result of corrupting only the entries of t corresponding to R.

Definition 5.3 (Noise Stability)**.** *Given a function $g : [-1,1]^n \to [-1,1]$ and a real number $\beta \geq 0$, we say g is β-stable if and only if the following holds for all $t \in [-1,1]^n$, $\alpha \in (0,1]$, and α-light distributions $\mathcal{D}$ over subsets of $[n]$:*

$$\mathop{\mathbb{E}}_{R \sim \mathcal{D}}\Big[\min\{g(t') : t' \underset{R}{\approx} t\}\Big] \geq g(t) - \alpha\beta.$$

To illustrate this definition, we show that our example function $g^{(\mathrm{mid})}$ is 4-stable. To see this, observe that changing k entries of the input to $g^{(\mathrm{mid})}$ can decrease its output by at most $\frac{4k}{n}$. This is because each of the k entries can go from 1 to -1 in the worst case, causing a change of $2k$, and then we normalize by $n/2$. When R is drawn from an α-light distribution and t is an arbitrary input, 4-stability therefore follows from the linearity of expectations:

$$\mathop{\mathbb{E}}_{R \sim \mathcal{D}}\Big[\min\{g^{(\mathrm{mid})}(t') : t' \underset{R}{\approx} t\}\Big] \geq \mathop{\mathbb{E}}_{R \sim \mathcal{D}}\left[g^{(\mathrm{mid})}(t) - \frac{4|R|}{n}\right] \geq g^{(\mathrm{mid})}(t) - 4\alpha.$$

We note that every function $g : [-1,1]^n \to [-1,1]$ is $2n$-stable, which follows from the union bound.

As a useful building block for proving some of our functions stable, we show that stable functions can be combined to yield other stable functions if composed with a convex, nondecreasing, and Lipschitz continuous function.

Proposition 5.4. *Fix $\beta, c \geq 0$, and let $g_1, g_2, \ldots, g_k : [-1,1]^n \to [-1,1]$ be β-stable functions. For every convex function $h : [-1,1]^k \to [-1,1]$ which is nondecreasing in each of its arguments and c-Lipschitz continuous in L^∞, the function $g(t) \stackrel{\text{def}}{=} h(g_1(t), \ldots, g_k(t))$ is (βc)-stable.*

Proof. For all $t \in [-1,1]^n$ and all α-light distributions $\mathcal{D}$,

$$
\begin{aligned}
&\mathop{\mathbb{E}}_{R\sim\mathcal{D}}\left[\min_{t' \underset{R}{\approx} t} g(t')\right] \\
&= \mathop{\mathbb{E}}_{R\sim\mathcal{D}}\left[\min_{t' \underset{R}{\approx} t} h(g_1(t'), \ldots, g_k(t'))\right] \\
&\geq \mathop{\mathbb{E}}_{R\sim\mathcal{D}}\left[h(\min_{t' \underset{R}{\approx} t} g_1(t'), \ldots, \min_{t' \underset{R}{\approx} t} g_k(t'))\right] && \text{(Since } h \text{ is nondecreasing)} \\
&\geq h(\mathop{\mathbb{E}}_{R\sim\mathcal{D}}\left[\min_{t' \underset{R}{\approx} t} g_1(t')\right], \ldots, \mathop{\mathbb{E}}_{R\sim\mathcal{D}}\left[\min_{t' \underset{R}{\approx} t} g_k(t')\right]) && \text{(Jensen's inequality)} \\
&\geq h(g_1(t) - \alpha\beta, \ldots, g_k(t) - \alpha\beta) && \text{(Stability of each } g_i\text{)} \\
&\geq h(g_1(t), \ldots, g_k(t)) - \alpha\beta c && \text{(Lipschitz continuity of } h\text{)} \\
&= g(t) - \alpha\beta c. && \square
\end{aligned}
$$

As a consequence of the above proposition, a convex combination of β-stable functions is β-stable, and the point-wise maximum of β-stable functions is β-stable.

5.2.1 Consequences of Stability and Continuity

We now state the two main results of our framework. Both results apply to functions $g : [-1,1]^n \to [-1,1]$ which are simultaneously Lipschitz continuous and noise stable, and $n \times M$ matrices A with entries in $[-1,1]$. Given a vector $x \in \Delta_M$

and integer $s > 0$, we view x as a probability distribution over $[M]$, and use the random variable $\widetilde{x} \in \Delta_M$ to denote the empirical distribution of s i.i.d. samples from x. Formally, $\widetilde{x} = \frac{1}{s}\sum_{i=1}^{s} e_{k_i}$, where $k_1, \ldots, k_s \in [M]$ are drawn i.i.d. according to x.

Our first result shows that when the number of samples s is chosen as a suitable function of the Lipschitz continuity and noise stability parameters, $g(A\widetilde{x})$ is not much smaller than $g(Ax)$ in expectation over $\widetilde{x}$. At a high level, we bound this difference as a sum of two error terms: one accounts for the effect of low-probability large errors in the inputs $\widetilde{t} = A\widetilde{x}$ to g, and the other accounts for effect of higher-probability small errors in the inputs $\widetilde{t}$. The former error term is bounded using noise stability, and the latter error term is bounded using Lipschitz continuity.

Theorem 5.5. *Let $g : [-1,1]^n \to [-1,1]$ be β-stable and c-Lipschitz in L^∞, let A be an $n \times M$ matrix with entries in $[-1,1]$, let $\alpha, \delta > 0$, and let $s \geq 2\ln(\frac{2}{\alpha})/\delta^2$ be an integer. Fix a vector $x \in \Delta_M$, and let the random variable $\widetilde{x}$ denote the empirical distribution of s i.i.d. samples from the probability distribution x. The following then holds: $\mathbb{E}[g(A\widetilde{x}))] \geq g(Ax) - \alpha\beta - c\delta$.*

Proof. Denote $t = Ax$ and $\widetilde{t} = A\widetilde{x}$. Note that $\widetilde{t}$ is a random variable. Also note that t_i and $\widetilde{t}_i$ can be viewed as the mean and empirical mean, respectively, of a distribution supported on $A_{i,1}, \ldots, A_{i,M} \in [-1,1]$. We say the i^{th} entry of t is *approximately preserved* if $|t_i - \widetilde{t}_i| \leq \delta$, and we say it is *corrupted* otherwise. Let $R \subseteq [n]$ denote the set of corrupted entries. Hoeffding's inequality, and our choice of the number of samples s, imply that R follows an α-light distribution.

Let t' be such that (1) $t'_i = \widetilde{t}_i$ for $i \in R$, and (2) $t'_i = t_i$ otherwise. Observe that $t' \underset{R}{\approx} t$, and $||t' - \widetilde{t}||_\infty \leq \delta$. We can now bound the expected difference between $g(t)$ and $g(\widetilde{t})$ as a sum of the error introduced by corrupted entries and the error introduced by the approximately preserved entries of t:

$$g(t) - \mathbb{E}\Big[g(\widetilde{t})\Big] = \mathbb{E}[g(t) - g(t')] + \mathbb{E}\Big[g(t') - g(\widetilde{t})\Big] \leq \alpha\beta + c\delta. \qquad \square$$

Notice that if we fix the desired approximation error ϵ, the minimum required number of samples s in Theorem 5.5 to guarantee that $\mathbb{E}[g(A\widetilde{x}))] \geq g(Ax) - \epsilon$ is obtained by minimizing $\lceil 2\ln(\frac{2}{\alpha})/\delta^2 \rceil$ over $\alpha, \delta > 0$ satisfying $\alpha\beta + \delta c \leq \epsilon$. Therefore, the required number of samples depends only on the error term ϵ, the noise stability parameter β, and the Lipschitz continuity parameter c; in particular, it is independent of n and M.

As a corollary of Theorem 5.5, we derive the following algorithmic result.

Theorem 5.6. *Let $g : [-1,1]^n \to [-1,1]$ be β-stable and c-Lipschitz, and let $M > 0$ be an integer. For every $\delta, \alpha > 0$, the M-dimensional mixture selection problem for g admits an $(\alpha\beta + c\delta)$-approximation algorithm in the additive sense, with runtime $n \cdot M^{O(\delta^{-2}\log(1/\alpha))} \cdot T$, where T denotes the time needed to evaluate g on a single input.*

Proof. Let $s \geq 2\ln(2/\alpha)/\delta^2$ be an integer. Our algorithm simply enumerates all s-uniform distributions, and outputs the one maximizing $g(Ax)$. This takes time $n \cdot M^{O(s)} \cdot T$. The approximation guarantee follows from Theorem 5.5 and the probabilistic method. $\square$

As a consequence of Theorem 5.6, the mixture selection problem for $g^{(\mathrm{mid})}$ admits a polynomial-time approximation scheme (PTAS) in the additive sense. The same holds for every function g which is $O(1)$-stable and $O(1)$-Lipschitz continuous. Specifically, by setting $\alpha = \frac{\epsilon}{2\beta}$ and $\delta = \frac{\epsilon}{2c}$, an ϵ-approximation algorithm runs in time $n \cdot m^{O(c^2\log(\beta/\epsilon)/\epsilon^2)} \cdot T$. Interestingly, neither noise stability nor Lipschitz continuity alone suffices for such a PTAS, as we argue in Section 5.4.

5.2.1.1 A Bi-criteria Extension of the Framework

Motivated by two of our applications, namely *Optimal signaling in normal form games* and *Persuasion in voting*, we extend our framework to the design of approximation algorithms for mixture selection with a *bi-criteria guarantee* when the function in question is stable but not Lipschitz continuous. We first define a (δ, ρ)-relaxation of a function.

Definition 5.7. *Given two functions $g, h : [-1,1]^n \to [-1,1]$ and parameters $\delta, \rho \geq 0$, we say h is a (δ, ρ)-relaxation of g if for all $t_1, t_2 \in [-1,1]^n$ with $||t_1 - t_2||_\infty \leq \delta$, $h(t_2) \geq g(t_1) - \rho$.*

Note that Lipschitz continuous functions are their own relaxations. In lieu of the Lipschitz continuity condition, we prove our bounds for a relaxation of the function.

Theorem 5.8. *Let $g : [-1,1]^n \to [-1,1]$ be β-stable, let A be an $n \times M$ matrix with entries in $[-1,1]$, let $\alpha > 0$ and $\delta, \rho \geq 0$, and let $s \geq 2\ln(\frac{2}{\alpha})/\delta^2$ be an integer. Fix a vector $x \in \Delta_M$, and let the random variable $\widetilde{x}$ denote the empirical distribution of s i.i.d. samples from probability distribution x. The following then holds for any (δ, ρ)-relaxation h of g,*

$$\mathbb{E}[h(A\widetilde{x}))] \geq g(Ax) - \alpha\beta - \rho.$$

Proof. Because the proof is almost identical to the proof of Theorem 5.5, we just mention the necessary modifications. Again, let $t = Ax$, let $\widetilde{t} = A\widetilde{x}$, let $R \subseteq [n]$ denote the set of corrupted inputs, and let t' be such that $t'_i = \widetilde{t}_i$ for $i \in R$ and $t'_i = t_i$ otherwise. Then

$$
\begin{aligned}
g(t) - \mathbb{E}\Big[h(\widetilde{t})\Big] &= \mathbb{E}[g(t) - g(t')] + \mathbb{E}\Big[g(t') - h(\widetilde{t})\Big] \\
&\leq \alpha\beta + \mathbb{E}\Big[g(t') - h(\widetilde{t})\Big] \\
&\leq \alpha\beta + \rho,
\end{aligned}
$$

where the first inequality follows by noise stability of g, and the last inequality follows from the fact that h is a (δ, ρ)-relaxation of g. □

Having replaced Theorem 5.5 by Theorem 5.8, a similar computational result as Theorem 5.6 can be inferred in the bi-criteria sense.

5.3 A Meta-Algorithm for Signaling

In this section, we use our framework to define an abstract signaling problem and characterize its approximation complexity. This abstract problem captures all of the signaling problems considered in this thesis.

To connect to our mixture selection framework, we consider signaling problems in which the principal's utility $f(\mu)$ from a posterior distribution $\mu \in \Delta_M$ can be written as $g(A\mu)$ for a function $g : [-1, 1]^n \to [-1, 1]$ and a matrix $A \in [-1, 1]^{n \times M}$. As described in Section 2.1.1, a signaling scheme φ with signals Σ corresponds to a family of probability-posterior pairs $\{(p_\sigma, \mu_\sigma)\}_{\sigma \in \Sigma}$ decomposing the prior $\lambda \in \Delta_M$ into a convex combination of posterior distributions (one per signal): $\lambda = \sum_{\sigma \in \Sigma} p_\sigma \mu_\sigma$. The objective of our signaling problem is then

$$
F(\varphi) = \sum_{\sigma \in \Sigma} p_\sigma f(\mu_\sigma) = \sum_{\sigma \in \Sigma} p_\sigma g(A\mu_\sigma).
$$

We note that this signaling problem can alternatively be written as an (infinite-dimensional) linear program which searches over probability measures supported on Δ_M with expectation λ. The separation oracle for the dual of this linear program is a mixture selection problem. Whereas we do not use this infinite-dimensional formulation or its dual directly, we nevertheless show that the same conditions — noise stability and Lipschitz continuity — on the function g which lead to an approximation scheme for mixture selection also lead to a similar approximation scheme for our signaling problem with $f(\mu) = g(A\mu)$.

Lemma 5.9. *If g is β-stable and c-Lipschitz, then for any constants $\alpha, \delta > 0$, and for any integer $s \geq 2\delta^{-2} \ln(2/\alpha)$, there exists a signaling scheme $\widetilde{\varphi}$ for which every posterior distribution is s-uniform, and $F(\widetilde{\varphi}) \geq OPT - (\alpha\beta + c\delta)$ where OPT denotes the value of the optimal signaling scheme.*

Proof. Let $s \geq 2\delta^{-2} \ln(2/\alpha)$, and let $\tau \in [M^s]$ index all s-uniform posteriors, with $\widetilde{\mu}_\tau$ denoting the τ'th such posterior. For an arbitrary signaling scheme $\varphi = (\Sigma, \{(p_\sigma, \mu_\sigma)\}_{\sigma \in \Sigma})$, we show that each posterior μ_σ can be decomposed into s-uniform posteriors without degrading the objective by more than $\alpha\beta + c\delta$:

1. μ_σ can be expressed as a convex combination of s-uniform posteriors as follows.

$$\mu_\sigma = \sum_{\tau \in [M^s]} \widetilde{p}_{\sigma,\tau} \widetilde{\mu}_\tau \quad \text{with} \quad \widetilde{p}_\sigma \in \Delta_{M^s}. \tag{5.1}$$

2. The value of objective function, i.e., $g(A\mu_\sigma)$, is decreased by no more than $\alpha\beta + c\delta$ through this decomposition,

$$\sum_{\tau \in [M^s]} \widetilde{p}_{\sigma,\tau} \cdot g(A\widetilde{\mu}_\tau) \geq g(A\mu_\sigma) - (\alpha\beta + c\delta). \tag{5.2}$$

The existence of such a decomposition follows from Theorem 5.5: Fix σ, and let $\widetilde{\mu} \in \Delta_M$ be the empirical distribution of s i.i.d. samples from distribution $\mu_\sigma \in \Delta_M$. The vector $\widetilde{\mu}$ is itself a random variable supported on s-uniform posteriors, its expectation is μ_σ, and by Theorem 5.5 we have $\mathrm{E}[g(A\widetilde{\mu})] \geq g(A\mu_\sigma) - (\alpha\beta + c\delta)$. Therefore, by taking $\widetilde{p}_{\sigma,\tau} = \Pr[\widetilde{\mu} = \widetilde{\mu}_\tau]$ for each $\tau \in [M^s]$ we get the desired decomposition of μ_σ.

The lemma follows by composing the decomposition φ with the decompositions of the posterior beliefs μ_σ to yield a signaling scheme $\widetilde{\varphi}$ with only s-uniform posteriors and $F(\widetilde{\varphi}) \geq F(\varphi) - (\alpha\beta + c\delta)$. Specifically, the signals of $\widetilde{\varphi}$ are $\Sigma \times [M^s]$, where signal (σ, τ) has probability $p_\sigma \cdot \widetilde{p}_{\sigma,\tau}$ and induces the posterior $\widetilde{\mu}_\tau$.[1] Using Equations (5.1) and (5.2), it is easy to verify that this describes a valid signaling scheme with $F(\widetilde{\varphi}) \geq F(\varphi) - (\alpha\beta + c\delta)$. □

Lemma 5.9 permits us to restrict attention to s-uniform posteriors without much loss in our objective. Since there are only M^s such posteriors, a simple linear program with M^s variables computes an approximately optimal signaling scheme.

Theorem 5.10 (Polynomial-Time Signaling)**.** *If g is β-stable and c-Lipschitz, then for any constant $\alpha, \delta > 0$, there exists a deterministic algorithm that constructs a signaling scheme with objective value at least $OPT - (\alpha\beta + c\delta)$, where OPT is the value of the optimal signaling scheme. Moreover, the algorithm runs in time* $\mathrm{poly}(M^{\delta^{-2}\ln(1/\alpha)}) \cdot n \cdot T$, *where T is the time needed to evaluate g on a single input.*

Proof. Let s be an integer with $s \geq (2\delta^{-2}\ln(2/\alpha))$, and let $\tau \in [M^s]$ index all s-uniform posteriors. Lemma 5.9 shows that restricting to s-uniform posteriors only introduces an $\alpha\beta + c\delta$ additive loss in the objective. Thus it suffices to compute the

[1] Note, however, that we can also "merge" all signals with the same posterior $\widetilde{\mu}_\tau$ without loss.

optimal signaling scheme supported only on s-uniform posteriors. This can be done using the following linear program:

$$\begin{aligned} \text{maximize} \quad & \textstyle\sum_{\tau \in [M]} p_\tau \cdot g(A\mu_\tau) \\ \text{subject to} \quad & \textstyle\sum_{\tau \in [M]} p_\tau \mu_\tau = \lambda \\ & p \in \Delta_M \end{aligned} \tag{5.3}$$

Note μ_τ is the τ'th s-uniform posterior — the only variables in this LP are $p_1, \ldots, p_{M^s}$. □

Our proofs can be adapted to obtain a bi-criteria guarantee in the absence of Lipschitz continuity, as in Section 5.2. The following theorem follows easily, and we omit the details.

Theorem 5.11 (Polynomial-Time Signaling (Bi-criteria))**.** *Let $g, h : [-1,1]^n \to [-1,1]$ be such that g is β-stable and h is a (δ,ρ)-relaxation of g, and let $\alpha > 0$ be a parameter. There exists a deterministic algorithm which, when given as input a matrix $A \in [-1,1]^{n\times m}$ and a prior distribution $\lambda \in \Delta_M$, constructs a signaling scheme $\varphi = \{(p_\sigma, \mu_\sigma)\}_{\sigma\in\Sigma}$ such that*

$$\sum_{\sigma\in\Sigma} p_\sigma h(A\mu_\sigma) \geq OPT - \alpha\beta - \rho,$$

where OPT is the maximum value of $F(\varphi^*) = \sum_{\sigma\in\Sigma^*} p^*_\sigma g(A\mu^*_\sigma)$ over signaling schemes $\varphi^* = \{(p^*_\sigma, \mu^*_\sigma)\}_{\sigma\in\Sigma^*}$. Moreover, the algorithm runs in time $\text{poly}(M^{\delta^{-2}\ln(1/\alpha)}) \cdot n \cdot T$*, where T denotes the time needed to evaluate h on a single input.*

Remarks We note that our proof suggests an extension of the result in Theorem 5.10 to cases in which f is given by a "black box" oracle, so long as we are promised that it is of the form $f(\mu) = g(A\mu)$. In this model the runtime of our algorithm does not depend on n, but instead depends on the cost of querying f. We also point out that even though we precompute the quality of all M^s posteriors, we can guarantee that our output signaling scheme uses at most $M + 1$ signals; this is because LP (5.3) has only $M + 1$ constraints, and therefore admits an optimal solution where at most $M + 1$ variables are non-zero.

5.3.1 A New QPTAS for Signaling in Normal Form Games

In this section, we present an approach different from the one in Section 4.3, which also gives a quasipolynomial-time approximation scheme for the problem of optimal signaling in Bayesian normal form games. We prove the following bi-criteria result.

Theorem 5.12. *Let $\epsilon > 0$ denote an approximation parameter, let $(\mathcal{A}, \lambda)$ be a Bayesian normal form game with $k = O(1)$ players, n actions, and M states of nature, and let $A_0 : [M] \times [n]^k \to [-1, 1]$ be an objective function given as a tensor. There is an algorithm with runtime* $\text{poly}(M^{\frac{\ln(n/\epsilon)}{\epsilon^2}}, n^{\frac{\ln n}{\epsilon^2}})$ *which outputs a signaling scheme φ and corresponding Bayesian ϵ-equilibria $\mathcal{X}$ satisfying $F(\varphi, \mathcal{X}) \geq OPT(\mathcal{A}, \lambda, A_0) - \epsilon$. This holds for both approximate NE and approximate WSNE.*

In other words, when the number of players is a constant we can in quasipolynomial time approximate the optimal reward from signaling while losing an additive ϵ in the objective as well as in the incentive constraints. Compared to our result in Section 4.3, the running time is slightly worse for general sum games, but not directly comparable for zero-sum games (depending on which one of n and M is larger). More specifically, for zero-sum games and constant $\epsilon > 0$, the QPTAS in Section 4.3 builds an ϵ-cover over Nash equilibria and runs in time $n^{O(\log n)}$, while the QPTAS in this section builds an ϵ-cover over posterior beliefs and runs in time $M^{O(\log n)}$.

Fix $\epsilon > 0$. To prove this theorem, we define functions g and g_ϵ which each take as input a k-player n-action game of complete information B, given as payoff tensors $B_1 \ldots, B_k : [n]^k \to [-1, 1]$, and an objective tensor $B_0 : [n]^k \to [-1, 1]$, and output a number in $[-1, 1]$. Specifically, $g(B, B_0) = \max\{B_0(x) : x \in EQ(B)\}$ and $g_\epsilon(B, B_0) = \max\{B_0(x) : x \in EQ_\epsilon(B)\}$, where $EQ(B)$ denotes the set of Nash equilibria of the game B, and $EQ_\epsilon(B)$ denotes the (non-empty) set of $\lceil s(\epsilon/4) \rceil$-uniform ϵ-Nash equilibria (or ϵ-WSNE) for s as given in Lemma 4.2. Recall that $B_0(x)$ denotes evaluating the multilinear map described by tensor B_0 at the mixed strategy profile $x \in \Delta_n^k$.

Now suppose we fix a Bayesian game $(\mathcal{A}, \lambda)$ and objective tensor A_0 as in the statement of Theorem 5.12. For a subgame with a posterior distribution $\mu \in \Delta_M$ over states of nature, the principal's expected utility at the "best" Nash equilibrium of this subgame can be written as $g(A^\mu, A_0^\mu)$. Similarly, the principal's expected utility at the "best" $\lceil s(\epsilon/4) \rceil$-uniform ϵ-NE (or ϵ-WSNE) can be written as $g_\epsilon(A^\mu, A_0^\mu)$. Observe that the input to both g and g_ϵ is a linear function of μ, as needed to apply the results in Section 5.3. For a signaling scheme φ corresponding to a decomposition $\lambda = \sum_{\sigma \in \Sigma} p_\sigma \mu_\sigma$ of the prior distribution λ into posterior distributions (see Section 2.1.1), we can write the principal's expected utility as $F(\varphi) = \sum_{\sigma \in \Sigma} p_\sigma g(F^{\mu_\sigma}, A^{\mu_\sigma})$ assuming that the players reach the "best" Nash equilibrium in each subgame, and $F_\epsilon(\varphi) = \sum_{\sigma \in \Sigma} p_\sigma g_\epsilon(F^{\mu_\sigma}, A^{\mu_\sigma})$ assuming that the players reach the "best" $\lceil s(\epsilon/4) \rceil$-uniform ϵ-equilibria. We use OPT to denote the maximum value of F over all signaling schemes.

We prove Theorem 5.12 by exhibiting an algorithm for computing a signaling scheme φ such that $F_\epsilon(\varphi) \geq \text{OPT} - \epsilon$. The proof hinges on two main lemmas.

Lemma 5.13. *The function g is $2(k+1)n^k$-stable.*

Proof. As noted in Section 5.2, any function mapping a hypercube $[-1,1]^N$ to the interval $[-1,1]$ is $2N$ stable. The function g is such a function with $N = (k+1)n^k$. □

Lemma 5.14. *The function g_ϵ is an $(\epsilon/4, \epsilon/2)$-relaxation of g.*

Proof. Consider tensors $B_0, \widetilde{B_0} : [n]^k \to [-1,1]$ with $|B_0(s) - \widetilde{B_0}(s)| \leq \epsilon/4$ for all $s \in [n]^k$, and two k-player n-action games $B = (B_1, \ldots, B_k)$ and $\widetilde{B} = (\widetilde{B}_1, \ldots, \widetilde{B}_k)$ with $|B_i(s) - \widetilde{B}_i(s)| \leq \epsilon/4$ for all $s \in [n]^k$. It suffices to show that $g_\epsilon(\widetilde{B}, \widetilde{B_0}) \geq g(B, B_0) - \epsilon/2$. Let $x \in \Delta_n^k$ be the Bayesian equilibrium of B for which $B_0(x) = g(B, B_0)$. By Lemma 4.2, there is a profile $\widetilde{x}$ of $\lceil s(\epsilon/4) \rceil$-uniform mixed strategies such that $\widetilde{x}$ is an $\epsilon/4$-equilibrium of B, and $B_0(\widetilde{x}) \geq B_0(x) - \epsilon/4$. Since $\widetilde{B}$ differs from B by at most $\epsilon/4$ everywhere, it follows that $\widetilde{x}$ is an ϵ-equilibrium of $\widetilde{B}$, i.e., $\widetilde{x} \in EQ_\epsilon(\widetilde{B})$. Similarly, since $\widetilde{B_0}$ differs from B_0 by at most $\epsilon/4$ everywhere, it follows that $\widetilde{B_0}(\widetilde{x}) \geq B_0(\widetilde{x}) - \epsilon/4 \geq B_0(x) - \epsilon/2$. We conclude that $g_\epsilon(\widetilde{B}, \widetilde{B_0}) \geq \widetilde{B_0}(\widetilde{x}) \geq g(B, B_0) - \epsilon/2$. □

We now complete the proof of Theorem 5.12 by instantiating Theorem 5.11 with g, $h = g_\epsilon$, and $\alpha = \frac{\epsilon}{4(k+1)n^k}$. The runtime is $\text{poly}(M^{\frac{\ln(1/\alpha)}{\epsilon^2}}, (k+1)n^k, T)$, where T is the time needed to evaluate g_ϵ (and compute the corresponding $\lceil s(\epsilon/4) \rceil$-uniform

ϵ-equilibrium) on a given input. Recall that $k = O(1)$ and $\alpha = \frac{\epsilon}{\text{poly}(n)}$. Moreover, using brute-force enumeration of all $\lceil s(\epsilon/4) \rceil$-uniform mixed strategy profiles we conclude that T is bounded by a polynomial in $n^{\frac{\ln n}{\epsilon^2}}$. Therefore our total runtime is $\text{poly}(M^{\frac{\ln(n/\epsilon)}{\epsilon^2}}, n^{\frac{\ln n}{\epsilon^2}})$, as needed.

Remarks Similar to our results in Section 4.3, in the special case of two-player zero-sum games and a principal interested in maximizing one player's utility, our techniques lead to a more efficient approximation scheme and a uni-criteria guarantee. This is because the principal's payoff tensor B_0 equals the payoff tensor B of one of the players (say, player 1), and consequently the function $g(B, B_0) = g(B, B) = \max_x \min_y x^T B y$ is n^2-stable and 2-Lipschitz. Its Lipschitz continuity follows from the fact that an ϵ-equilibrium of a zero-sum game leads to utilities within ϵ of the equilibrium utilities. Moreover, evaluating g now takes time $T = \text{poly}(M, n)$. Theorem 5.10 instantiated with $\alpha = \frac{\epsilon}{4n^2}$ and $\delta = \epsilon/4$, leads to an algorithm with runtime $\text{poly}(M^{\frac{\ln(n/\epsilon)}{\epsilon^2}}, n)$, which outputs a signaling scheme φ and corresponding Bayesian (exact) Nash-equilibria $\mathcal{X}$ satisfying $F(\varphi, \mathcal{X}) \geq \text{OPT}(\mathcal{A}, \lambda, A_0) - \epsilon$.

5.4 Hardness Results for Mixture Selection

We now present evidence that both our assumptions — Noise stability and Lipschitz continuity — appear necessary for general positive results along the lines of those in Theorem 5.6.

Noise stability alone is not sufficient for a PTAS. In Section 5.4.1, we define a function $g^{(\text{slope})} : [0,1]^n \to [0,1]$ which is 1-stable. Furthermore, $g^{(\text{slope})}$ is $O(1)$-Lipschitz with respect to the L^1 metric, which is a weaker property than

Lipschitz continuity with respect to L^∞. We show in Theorem 5.18 that there is a polynomial-time reduction from the maximum independent set problem on n-node graphs to the n-dimensional mixture selection for $g^{(\text{slope})}$. The reduction precludes a polynomial-time (additive) ϵ-approximation algorithm for some constant $\epsilon > 0$.

Lipschitz continuity alone is not sufficient for a PTAS. One might hope to prove NP-hardness of mixture selection in the absence of stability. However, since every function $g : [-1, 1]^n \to [-1, 1]$ is $2n$-stable, Theorem 5.6 implies a quasipolynomial-time approximation scheme in the additive sense whenever g is $O(1)$-Lipschitz. Nevertheless, we prove hardness of approximation assuming the *planted clique conjecture* ([66] and [72]). More specifically, in Section 5.4.2, we exhibit a reduction from the planted k-clique problem to mixture selection for the 3-Lipschitz function $g_k^{(\text{clique})}(t) = t_{[k]} - t_{[k+1]} + t_{[n]}$. When $k = \omega(\log^2 n)$ and A is the adjacency matrix of an n-node undirected graph G, we show that $\max_x g_k^{(\text{clique})}(Ax) \approx \frac{1}{2}$ with high probability if $G \sim \mathcal{G}(n, \frac{1}{2})$, and $\max_x g_k^{(\text{clique})}(Ax) \approx 1$ with high probability if $G \sim \mathcal{G}(n, \frac{1}{2}, k)$ (defined in Section 2.3).

5.4.1 NP-hardness in the Absence of Lipschitz Continuity

We now show that stability alone does not suffice for an additive PTAS for mixture selection, in general. First, we show that mixture selection for a 1-stable function $g^{(\text{vote-sum})}(t)$ does not admit a (uni-criteria) additive PTAS unless P = NP. $g^{(\text{vote-sum})}$ is motivated by the application of persuading voters presented in Section 6.2, and simply returns the fraction of nonnegative entries of $t = Ax$, i.e.,

$$g^{(\text{vote-sum})}(t) \stackrel{\text{def}}{=} \sum_{i \in [n]} \frac{1}{n} I[t_i \geq 0].$$

In addition, since $g^{(\text{vote-sum})}$ is not continuous in any metric, we exhibit a "smoothed" function $g^{(\text{slope})}$ which is 1-stable and $O(1)$-Lipschitz with respect to L^1, but not $O(1)$-Lipschitz with respect to L^∞, and show that mixture selection for $g^{(\text{slope})}$ still does not admit an additive PTAS unless P = NP.

Both NP-hardness results share a similar reduction from the *maximum independent set* problem. We use a consequence of the result by [71], namely that there exists a constant ϵ such that it is NP-hard to approximate maximum independent set to within an additive error of ϵn, where n denotes the number of vertices.

Given an n-node undirected graph G, let $\text{OPT}_{\text{IS}} = \text{OPT}_{\text{IS}}(G)$ be the size of its largest independent set. We define the $n \times n$ matrix $A = A(G)$ as follows:

- Diagonal entries of A are all $\frac{1}{2}$ ($A_{i,i} = \frac{1}{2}$ for all $1 \le i \le n$).
- When vertices i and j share an edge in G, both $A_{i,j}$ and $A_{j,i}$ are -1.
- All other entries of A, namely $A_{i,j}$ for non-adjacent distinct vertices i and j, are $-\frac{1}{4n}$.

We relate OPT_{IS} to convex combinations of the columns of A as follows.

Observation 5.15. *Let $\mathcal{I}$ be an independent set of G with $|\mathcal{I}| = k$. There exists $x \in \Delta_n$ such that k entries of Ax are at least $\frac{1}{4n}$, and all remaining entries are strictly negative.*

Proof. Let $x \in \Delta_n$ be the normalized indicator vector of $\mathcal{I}$ — i.e., $x_i = \frac{1}{k}$ if $i \in \mathcal{I}$ and $x_i = 0$ otherwise. By construction $(Ax)_i = \frac{1}{k}(\frac{1}{2} - (k-1)\frac{1}{4n}) \ge \frac{1}{4n}$ whenever $i \in \mathcal{I}$, and $(Ax)_i \le -\frac{1}{4n}$ otherwise. □

Observation 5.16. *For any $x \in \Delta_n$, nonnegative entries of Ax correspond to an independent set of G. Consequently, Ax can have at most OPT_{IS} nonnegative entries.*

Proof. Let $t = Ax$. Consider an edge $\{i, j\}$ of graph G. We have $t_i \leq \frac{x_i}{2} - x_j - \frac{1}{4n}(1 - x_i - x_j)$ and a similar inequality for t_j, so

$$t_i + t_j \leq -\frac{x_i + x_j}{2} - \frac{1 - x_i - x_j}{2n} < 0.$$

Therefore, t_i and t_j cannot be both nonnegative. We conclude that the nonnegative coordinates of t correspond to an independent set of G. □

Observations 5.15 and 5.16 imply that $\max_{x \in \Delta_n} g^{(\text{vote-sum})}(Ax) = \frac{\text{OPT}_{\text{IS}}}{n}$. Combined with the fact that obtaining an additive PTAS for the maximum independent set problem is NP-hard, we get the following theorem.

Theorem 5.17. *Mixture selection for the 1-stable function $g^{(\text{vote-sum})}$ admits no additive PTAS unless P = NP.*

Noting that $g^{(\text{vote-sum})}$ is a discontinuous function, for emphasis we exhibit a function $g^{(\text{slope})}$ which is Lipschitz continuous in L^1 (but not in L^∞) and 1-noise stable, but for which the same impossibility result holds by an identical reduction. Informally, $g^{(\text{slope})}$ "smoothes" the threshold behavior of $g^{(\text{vote-sum})}$ as follows: each input t_i contributes 0 to $g^{(\text{slope})}(t)$ when $t_i \leq 0$, contributes $\frac{1}{n}$ when $t_i \geq \frac{1}{4n}$, and the contribution is a linear function of t_i increasing from 0 to $\frac{1}{n}$ for $t_i \in [0, \frac{1}{4n}]$. Formally, we define $g^{(\text{slope})}(t) = \sum_{i=1}^{n} \min \left\{ 4 \max\{0, t_i\}, \frac{1}{n} \right\}$. Since each entry of t contributes at most $\frac{1}{n}$ to $g^{(\text{slope})}(t)$, it is easy to verify that $g^{(\text{slope})}$ is 1-stable. Moreover, since the partial derivatives of $g^{(\text{slope})}(t)$ are upper-bounded by 4, it is

4-Lipschitz continuous with respect to the L^1 metric. Observations 5.15 and 5.16 imply that $\max_{x \in \Delta_n} g^{(\text{slope})}(Ax) = \frac{\text{OPT}_{\text{IS}}}{n}$, ruling out an additive PTAS for mixture selection for $g^{(\text{slope})}$.

Theorem 5.18. *The function $g^{(slope)}$ is 1-stable and $O(1)$-Lipschitz with respect to L^1, and yet mixture selection for $g^{(slope)}$ admits no additive PTAS unless P = NP.*

5.4.2 Planted Clique Hardness in the Absence of Stability

We now present evidence that Lipschitz continuity alone does not suffice for a PTAS for mixture selection. Recalling that a quasipolynomial time algorithm follows from our framework whenever a function is $O(1)$-Lipschitz, we reduce from the *planted clique problem*—for which a quasipolynomial time algorithm exists, and yet a polynomial-time algorithm is conjectured not to exist—rather than from an NP-hard problem.

We let $k = k(n)$ be as in Conjecture 2.6, and consider mixture selection for the function $g_k^{(\text{clique})} : [0,1]^n \to [0,1]$ with $g_k^{(\text{clique})}(t) = t_{[k]} - t_{[k+1]} + t_{[n]}$, where $t_{[i]}$ denotes the i'th largest entry of the vector t. It is easy to verify that $g_k^{(\text{clique})}$ is 3-Lipschitz with respect to the L^∞ metric, yet is not $O(1)$-stable. We prove the following theorem.

Theorem 5.19. *Conjecture 2.6 implies that there is no additive PTAS for mixture selection for $g_k^{(clique)}$.*

To prove Theorem 5.19, we show that $\max_{x \in \Delta_n} g_k^{(\text{clique})}(Ax)$ is arbitrarily close to 1 with high probability when A is the adjacency matrix of $G \sim \mathcal{G}(n, \frac{1}{2}, k)$, and is bounded away from 1 with high probability when A is the adjacency matrix of $G \sim \mathcal{G}(n, \frac{1}{2})$. For convenience, and without loss of generality, we assume that both random graphs include each self-loop with probability $\frac{1}{2}$ — i.e., diagonal entries of the adjacency matrix A are independent uniform draws from $\{0,1\}$ in both cases. Our argument is captured by the following two lemmas.

Lemma 5.20. *Fix a constant $\epsilon > 0$. Let $G \sim \mathcal{G}(n, \frac{1}{2}, k)$, and let A be its adjacency matrix. With probability $1 - o(1)$, there exists an $x \in \Delta_n$ such that $g_k^{(clique)}(Ax) \geq 1 - \epsilon$.*

Proof. Let $\mathcal{C}$ denote the vertices of the planted k-clique. We set $x_i = \frac{1}{k}$ if $i \in \mathcal{C}$ and 0 otherwise. Let $t = Ax$. For $i \in \mathcal{C}$, $t_i \geq 1 - \frac{1}{k}$. On the other hand, all other entries of t concentrate around $\frac{1}{2}$ with high probability. For $i \notin \mathcal{C}$, t_i is simply the average of k independent Bernoulli random variables by definition of $\mathcal{G}(n, \frac{1}{2}, k)$; using Hoeffding's inequality, we bound the probability that t_i deviates from its expectation by more than a constant $\delta > 0$, to be chosen later:

$$\Pr\left[\left|t_i - \frac{1}{2}\right| > \delta\right] \leq 2e^{-2\delta^2 k}.$$

By the union bound, $t_i \in [\frac{1}{2} - \delta, \frac{1}{2} + \delta]$ simultaneously for all $i \notin \mathcal{C}$ with probability at least $1 - n2^{-\Omega(k)} = 1 - o(1)$. Thus $t_{[k+1]} - t_{[n]} \leq 2\delta$ and $g_k^{(\text{clique})}(t) = t_{[k]} - (t_{[k+1]} - t_{[n]}) \geq 1 - \frac{1}{k} - 2\delta$ with probability $1 - o(1)$. Choosing $\delta = \epsilon/3$, we conclude that $g_k^{(\text{clique})}(t) \geq 1 - \epsilon$ with probability $1 - o(1)$. □

Lemma 5.21. *Fix a constant $\epsilon > 0$. Let $G \sim \mathcal{G}(n, \frac{1}{2})$, and let A be its adjacency matrix. With probability $1 - o(1)$, $g_k^{(clique)}(Ax) \leq \frac{3}{4} + \epsilon$ for all $x \in \Delta_n$.*

Proof. Recall that $g_k^{(\text{clique})}$ is $O(1)$-Lipschitz and — like any other function from the hypercube to the bounded interval — $O(n)$-stable. If there exists x^* such that $g_k^{(\text{clique})}(Ax^*) \geq \frac{3}{4} + \epsilon$, then Theorem 5.5 implies that there is an integer $s = O(\log n)$

and an s-uniform vector $\widetilde{x}$ such that $g_k^{(\text{clique})}(A\widetilde{x}) > \frac{3}{4}$. There are n^s such vectors. We next show that for an arbitrary fixed vector $x \in \Delta_n$ the probability that $g_k^{(\text{clique})}(Ax) > \frac{3}{4}$ is at most $2^{-\Omega(k)}$. This will complete the proof by the union bound, since $1 - n^s \cdot 2^{-\Omega(k)} = 1 - o(1)$.

Fix $x \in \Delta_n$, and let $t = Ax$. Define $\mathcal{D}$ as the distribution supported on $[0, 1]$ which is sampled as follows: draw a uniformly from $\{0, 1\}^n$, and output $a \cdot x$. Since A is the adjacency matrix of $G \sim \mathcal{G}(n, \frac{1}{2})$, each entry t_i of t can be viewed as an independent draw from $\mathcal{D}$. We exploit a key property of $\mathcal{D}$ in our proof, namely the fact that $\mathcal{D}$ is *symmetric* about $\frac{1}{2}$. Formally we mean that $\Pr_{\mathcal{D}}[r] = \Pr_{\mathcal{D}}[1 - r]$ for all $r \in [0, 1]$, and this follows easily from the definition of $\mathcal{D}$.

Symmetry of $\mathcal{D}$ implies that $\Pr_{r\sim\mathcal{D}}[r \geq \frac{1}{2}] = \Pr_{r\sim\mathcal{D}}[r \leq \frac{1}{2}] \geq \frac{1}{2}$. Recalling that $k = o(n)$ and that entries of t are independent draws from $\mathcal{D}$, the Chernoff bound implies that the following holds with probability at least $1 - 2^{-\Omega(n)}$:

$$t_{[n]} \leq \frac{1}{2} \leq t_{[k+1]}. \tag{5.4}$$

If $g_k^{(\text{clique})}(t) > \frac{3}{4}$, then the following two conditions must hold:

1. $t_{[k]} > \frac{3}{4}$, and
2. $t_{[k+1]} - t_{[n]} < \frac{1}{4}$.

Condition 1 implies that the k largest entries of t are all at least $\frac{3}{4}$. Furthermore, unless Inequality (5.4) is violated — which happens with probability $2^{-\Omega(n)}$ — Condition 2 implies that the remaining entries of t are all strictly between $\frac{1}{4}$ and $\frac{3}{4}$. Let p denote $\Pr_{r\sim\mathcal{D}}[r \leq \frac{1}{4}]$, also equal to $\Pr_{r\sim\mathcal{D}}[r \geq \frac{3}{4}]$ by symmetry of $\mathcal{D}$. The probability that k entries of t are at least $\frac{3}{4}$ and all remaining entries are in $(\frac{1}{4}, \frac{3}{4})$ is given by

$\binom{n}{k}p^k(1-2p)^{n-k}$, which is maximized at $p = \frac{k}{2n}$, with maximum value $2^{-\Omega(k)}$. In summary, the probability that $g_k^{(\text{clique})}(Ax) > \frac{3}{4}$ is at most $2^{-\Omega(k)} + 2^{-\Omega(n)} = 2^{-\Omega(k)}$. $\square$

ETH-Hardness in the Absence of Stability

For mixture selection in the absence of noise stability, we can also show that a QPTAS is the best-possible approximation scheme, assuming the Exponential Time Hypothesis (ETH) [15]. Our proof follows from a clean reduction from the best-Nash problem, for which Braverman et al. [20] showed that a QPTAS is essentially optimal. We choose to present the planted clique hardness result in this thesis because it is more elementary, and gives a simple function $g_k^{(\text{clique})}$.

Chapter 6

Signaling in Anonymous Games

Anonymous games are multiplayer games in which the utility of each player depends on her own strategy, as well as the number (as opposed to the identity) of other players who play each of the strategies. Anonymous games comprise an important class of succinct games — well-studied in the economics literature (see, e.g., [17, 18, 76]) — capturing a wide range of phenomena that frequently arise in practice, including auctions, voting systems, and congestion games.

In this chapter, we study the complexity of optimal signaling in anonymous games. We start with two special cases: probabilistic second price auctions, and majority voting with uncertainty. We give the first polynomial time approximation schemes (PTAS) for both problems (Theorem 6.2 and 6.5), which follow from the powerful mixture selection framework presented in Chapter 5.

We then take a slight detour to present the currently (asymptotically) best algorithm for computing Nash equilibria in anonymous games (Theorem 6.6); and we also present some evidence suggesting our algorithm might be essentially tight (Theorem 6.7). Anonymous games have a unique property compared to all other games we study in this thesis (e.g., network routing games, normal form games, second price auctions and majority voting); the computational complexity of (approximate) Nash equilibria in complete-information anonymous games is still open.

The work presented in this chapter appeared in [26] and [28].

6.1 Signaling in Second-Price Auctions

In this section, we examine signaling in probabilistic second-price auctions as defined in Section 2.1.5. Recall that in this setting, a probabilistic item is being auctioned, and the instantiation of the item is known to the auctioneer but not to the bidders. This is particularly relevant in advertising auctions, where items are impressions associated with demographics that are a priori unknown to the advertisers bidding in the auction.

We consider the algorithmic problem faced by an auctioneer, who seeks to reveal partial information to maximize the expected revenue before subsequently running a second-price auction. It was shown in [21, 52] that polynomial-time algorithms exist for several special cases of this problem. However, the general problem was shown to be NP-hard even with 3 bidders — specifically, no additive FPTAS exists unless P = NP. In this section, we resolve the approximation complexity of this basic signaling problem by giving an additive PTAS. We note that variations of this problem were considered in [59, 62], with different constraints on the signaling scheme — the results in these works are not directly relevant to our model.

6.1.1 PTAS from Mixture Selection: Revenue is Stable

Given a probabilistic auction with valuation distribution $\mathcal{D}$, and a signaling scheme φ expressed as a decomposition $\{p_\sigma, \mu_\sigma\}_{\sigma \in \Sigma}$ of the prior distribution λ, we can express the auctioneer's expected revenue as

$$\sum_{\sigma \in \Sigma} p_\sigma \mathop{\mathbb{E}}_{V \sim \mathcal{D}} [\max 2(V \mu_\sigma)] \,,$$

where the function max2 returns the second largest entry of a given vector. To apply our main theorem, we need to show that the revenue in a subgame with posterior distribution $\mu \in \Delta_M$ — namely $\mathbb{E}_{V\sim\mathcal{D}}[\max2(V\mu)]$ — can be written in the form $g(W\mu)$ for a matrix W. To facilitate our discussion we assume that the valuation distribution $\mathcal{D}$ has finite support size C, though this is without loss of generality. Imagine we form a large matrix W by stacking matrices in the support of $\mathcal{D}$ on top of each other. Formally, $W = [V_1^T, V_2^T, \dots, V_C^T]^T$ where V_i is the ith matrix in the support of $\mathcal{D}$. When matrix V_i is drawn from $\mathcal{D}$, we take the second-highest bid from the rows of W corresponding to V_i (rows $(i-1)\cdot n+1$ to $i \cdot n$, where n is the number players). For $S \subseteq [nC]$ and $t \in [0,1]^{nC}$, let $\max2_S(t)$ denote the second-highest value among entries of t indexed by S. Then we can write the auctioneer's expected revenue as

$$g^{(\mathrm{rev})}(W\mu) = \mathop{\mathbb{E}}_{V\sim\mathcal{D}}\Big[\max2_{S(V)}(W\mu)\Big]$$

where $S(V)$ is the set of rows in W corresponding to V.

Lemma 6.1 (Smooth and Stable Revenue)**.** *The function* $g^{(rev)}(t) = \mathbb{E}_{V\sim\mathcal{D}}\Big[\max2_{S(V)}(t)\Big]$ *is 1-Lipschitz and 2-stable.*

Proof. Because $\max2_S$ is 1-Lipschitz for a fixed set of indices S, it follows that $g^{(\mathrm{rev})}$, which is a convex combination of these 1-Lipschitz functions, is also 1-Lipschitz.

To show that $g^{(\mathrm{rev})}$ is stable, we first show that the function $\max2 : [0,1]^n \to [0,1]$ is stable. Given $t \in [0,1]^n$ and a random set $R \subseteq [n]$ drawn from an α-light distribution $\mathcal{D}$, the union bound implies that R includes neither of the two largest entries of t with probability at least $1-2\alpha$. In this case, the value of max2 is not affected by corruption of the entries indexed by R. Hence

$$\mathop{\mathbb{E}}_{R\sim\mathcal{D}}\Big[\min\{\max2(t') : t' \underset{R}{\approx} t\}\Big] \geq (1-2\alpha)\cdot \max2(t) + 2\alpha \cdot 0 \geq \max2(t) - 2\alpha.$$

Therefore max2 is 2-stable, which implies that $\max2_S : [0,1]^{nC} \to [0,1]$ is also 2-stable for any fixed set of indices S. The function $g^{(\text{rev})}$ is a convex combination of functions of the form $\max2_S$, and is therefore also 2-stable by Proposition 5.4. ☐

Theorem 6.2. *The revenue-maximizing signaling problem in second-price auctions admits an additive PTAS when the valuation distribution is given explicitly, and an additive PRAS when the valuation distribution is given by a sampling oracle.*

Proof. Lemma 6.1 shows that the function $g^{(\text{rev})}$ is 2-stable and 1-Lipschitz. If the valuation distribution $\mathcal{D}$ is explicitly given with support size C, the function $g^{(\text{rev})}$ can be evaluated in $\text{poly}(n, M, C)$ time. Then for any $\epsilon > 0$, it follows from Theorem 5.10 by setting $\alpha = \epsilon/4$ and $\delta = \epsilon/2$ that there is a deterministic algorithm that computes a signaling scheme with expected revenue $(\text{OPT} - \epsilon)$, in time $\text{poly}(n, M^{\epsilon^{-2}\ln(1/\epsilon)}, C)$.

If $\mathcal{D}$ is given via a sampling oracle, standard tail bounds and the union bound imply that $C = \Theta((s\log m + \log(\gamma^{-1}))/\epsilon^2)$ samples from $\mathcal{D}$ suffice to estimate to within $O(\epsilon)$ the revenue associated with every s-uniform posterior in Δ_M, with success probability $1-\gamma$. Since revenue is $O(1)$-stable and $O(1)$-Lipschitz, Lemma 5.9 implies that we can restrict attention to signaling schemes with s-uniform posteriors for $s = \text{poly}(\frac{1}{\epsilon})$. Proceeding as in Theorem 5.10, using the revenue estimates from Monte-Carlo sampling in lieu of exact values, we can construct a signaling scheme with revenue $(\text{OPT} - \epsilon)$ in time $\text{poly}(n, M^{\epsilon^{-2}\ln(1/\epsilon)}, \log(\frac{1}{\gamma}))$, with success probability $1-\gamma$. ☐

6.1.2 NP-hardness of an Additive FPTAS

Emek et al. [52] proved that revenue-maximizing signaling in probabilistic second price auctions is NP-hard, via a reduction from MAX-CUT. More specifically, given a graph G with n nodes and m edges, they can construct a Bayesian second price auction such that the value of the optimal signaling scheme is roughly $\frac{m+C^*}{\text{poly}(n)}$, where C^* is the size of the maximum cut of G. Since MAX-CUT is APX-hard and their reduction is gap preserving up to a multiplicative factor of poly(n), Emek et al. [52] implicitly ruled out an additive FPTAS for this problem.

6.2 Persuasion in Voting

In this section, we study persuasion problem in voting as defined in Section 2.1.6. Recall that we have an election with two possible outcomes. Each voter casts a 'Yes'/'No' vote, and the ballot measure is passed if the fraction of 'Yes' votes exceeds a certain pre-specified threshold. As in [5], we focus on the scenario in which voters have uncertainty regarding their utilities for the two possible outcomes. We consider a principal looking to influence the outcome of the election by signaling, who wants to maximize the probability of the measure passing.

For our approximation algorithms, we also allow implementation in approximately dominant strategies — i.e., we sometimes assume a voter votes 'Yes' if his utility $u(i, \mu)$ is at least $-\delta$ for a small parameter δ.[1] We assume that the state of nature $\theta \in \Theta$ is drawn from a common prior $\lambda \in \Delta_M$, and a principal with access to θ reveals a public signal σ prior to voters casting their votes. As usual, we adopt the

[1] Such relaxations seem necessary for our results. Moreover, depending on the context, modes of intervention for shifting the votes of voters who are close to being indifferent may be realistic.

perspective of a principal looking to commit to a public signaling scheme $\varphi : \Theta \to \Sigma$, for some set of signals Σ.

Alonso and Câmara [5] consider a principal interested in maximizing the probability that at least 50% (or some given threshold) of the voters vote 'Yes', in expectation over states of nature. Theirs is the natural objective when the election employs a majority (or threshold) voting rule, and the principal is interested in influencing the outcome of the vote. Approximating this objective requires nontrivial modifications to our framework, and therefore we begin this section by examining a different, yet also natural, objective: the expected number of 'Yes' votes. We design a bi-criteria approximation scheme for this objective, then describe the necessary modifications for the threshold function objective of [5].

6.2.1 Maximizing Expected Number of Votes

We now examine bi-criteria approximation algorithms for maximizing the expected number of 'Yes' votes. For our benchmark, we use the function $g^{(\text{vote-sum})}(t) \stackrel{\text{def}}{=} \sum_{i\in[n]} \frac{1}{n} I[t_i \geq 0]$, where $I[\mathcal{E}]$ denotes the indicator function for event $\mathcal{E}$. Assuming voters vote 'Yes' precisely when their posterior expected utility for a 'Yes' outcome is nonnegative, the number of 'Yes' votes when voters have preferences $U \in [-1,1]^{n\times m}$ and posterior belief $\mu \in \Delta_M$ equals $g^{(\text{vote-sum})}(U\mu)$. When the state of nature is distributed according to a common prior λ, and voters are informed according to a signaling scheme $\varphi = \{\mu_\sigma, p_\sigma\}_{\sigma\in\Sigma}$, the expected number of 'Yes' votes equals $F^{(\text{vote-sum})}(\varphi, U, \lambda) \stackrel{\text{def}}{=} \sum_{\sigma\in\Sigma} p_\sigma g^{(\text{vote-sum})}(U\mu_\sigma)$. We use $\text{OPT}^{(\text{vote-sum})}(U, \lambda)$ to denote the maximum value of $F^{(\text{vote-sum})}(\varphi, U, \lambda)$ over public signaling schemes φ.

As the first step to apply our framework, we prove that $g^{(\text{vote-sum})}$ is stable.

Lemma 6.3. *The function $g^{(vote\text{-}sum)}$ is 1-stable.*

Proof. For each voter $i \in [n]$, let $g_i : [-1,1]^n \to \{0,1\}$ be the function indicating whether voter i prefers the 'Yes' outcome, i.e., $g_i(t) = I[t_i \geq 0]$. Each individual g_i is 1-stable, because as long as the i'th input t_i is not corrupted the output of g_i does not change. Therefore $g^{(\text{vote-sum})}(t) = \frac{1}{n}\sum_{i=1}^{n} g_i(t)$, being a convex combination of 1-stable functions, is 1-stable by Proposition 5.4. ☐

Unfortunately, $g^{(\text{vote-sum})}$ is not $O(1)$-Lipschitz. We therefore employ the bi-criteria extension to our framework from Definition 5.7. Specifically, for a parameter $\delta > 0$, we assume a voter votes 'Yes' as long as his expected utility from a 'Yes' outcome is at least $-\delta$. Correspondingly, we define the relaxed function $g_\delta^{(\text{vote-sum})}(t) \stackrel{\text{def}}{=} \sum_{i\in[n]} \frac{1}{n} I[t_i \geq -\delta]$; the expected number of 'Yes' votes from a signaling scheme $\varphi = \{\mu_\sigma, p_\sigma\}_{\sigma\in\Sigma}$ can analogously be written as $F_\delta^{(\text{vote-sum})}(\varphi, U, \lambda) \stackrel{\text{def}}{=} \sum_{\sigma\in\Sigma} p_\sigma g_\delta^{(\text{vote-sum})}(U\mu_\sigma)$.

We can verify that $g_\delta^{(\text{vote-sum})}$ is a $(\delta, 0)$-relaxation of $g^{(\text{vote-sum})}$; combining this fact with Theorem 5.11 yields a bi-criteria approximation scheme for the problem of maximizing the expected number of 'Yes' votes.

Theorem 6.4. *Let $\epsilon, \delta > 0$ be parameters, let $U \in [-1,1]^{n\times M}$ describe the preferences of n voters in M states of nature, and let $\lambda \in \Delta_M$ be the prior of states of nature. There is an algorithm with runtime* $\text{poly}(M^{\delta^{-2}\ln(1/\epsilon)}, n)$ *for computing a signaling scheme φ such that $F_\delta^{(vote\text{-}sum)}(\varphi, U, \lambda) \geq OPT^{(vote\text{-}sum)}(U, \lambda) - \epsilon$.*

6.2.2 Maximizing the Probability of a Majority Vote

We now sketch the necessary modifications when the principal is interested in maximizing the probability of a 'Yes' outcome, assuming a majority voting rule. We make two relaxations: we assume a voter votes 'Yes' as long as his expected utility from a 'Yes' outcome is at least $-\delta$, and assume that the 'Yes' outcome is attained when at least a $(0.5 - \delta)$ fraction of voters vote 'Yes'. Our benchmark will be the maximum probability of a 'Yes' outcome in the absence of these two relaxations.

We note that [5] do not require these relaxations. They focus on characterizing the structures of the optimal signaling scheme, and fall short at providing an algorithm for (approximately) optimal signaling. In their analysis for persuading multiple voters, they make use of (the convex hull of) the set of posteriors that induce more than 50% of the voters to vote 'Yes'; this set is in general non-convex and may have exponentially many disconnected regions, making it difficult to convert their insights into efficient algorithms.

We define our benchmark using the function $g^{(\text{vote-thresh})}(t) = I[g^{(\text{vote-sum})}(t) \geq 0.5]$ which evaluates to 1 if at least half of its n inputs are nonnegative, and to 0 otherwise. This function is not $O(1)$-stable, so we work with a more stringent benchmark which is. Specifically, for a parameter $\delta > 0$, we use the function $g_\delta^{(\text{vote-smooth-thresh})}$ which is pointwise greater than or equal to $g^{(\text{vote-thresh})}$, defined as follows:

$$g_\delta^{(\text{vote-smooth-thresh})}(t) = \begin{cases} \frac{1}{\delta}\left(g^{(\text{vote-sum})}(t) - 0.5 + \delta\right) & \text{if } g^{(\text{vote-sum})}(t) \in [0.5 - \delta, 0.5] \\ g^{(\text{vote-thresh})}(t) & \text{otherwise.} \end{cases}$$

Observe that $g_\delta^{(\text{vote-smooth-thresh})}$ applies a continuous piecewise-linear function to the output of $g^{(\text{vote-sum})}$; it is easy to verify that $g_\delta^{(\text{vote-smooth-thresh})}$ is $\frac{1}{\delta}$-stable, and upper bounds $g^{(\text{vote-thresh})}$.

Finally, to measure the quality of our output we define the relaxed function $g_\delta^{(\text{vote-thresh})} : [-1,1]^n \to \{0,1\}$, which outputs 1 if at least a $(0.5-\delta)$ fraction of its inputs exceed $-\delta$, and outputs 0 otherwise. By Definition 5.7, $g_\delta^{(\text{vote-thresh})}$ is a $(\delta,0)$-relaxation of $g_\delta^{(\text{vote-smooth-thresh})}$ (and, consequently, also of $g^{(\text{vote-thresh})}$).

Let $F^{(\text{vote-thresh})}(\varphi,U,\lambda)$ and $F_\delta^{(\text{vote-thresh})}(\varphi,U,\lambda)$ denote the functions which evaluate the quality of a signaling scheme φ using $g^{(\text{vote-thresh})}$ and $g_\delta^{(\text{vote-thresh})}$, respectively. Moreover, let $\text{OPT}^{(\text{vote-thresh})}(U,\lambda)$ be the maximum value of $F^{(\text{vote-thresh})}(\varphi,U,\lambda)$ over signaling schemes φ. We apply Theorem 5.11 to $g_\delta^{(\text{vote-thresh})}$ and $g^{(\text{vote-smooth-thresh})}$, setting $\alpha=\epsilon\delta$, and use the fact that $g^{(\text{vote-smooth-thresh})}$ upper-bounds our true benchmark $g^{(\text{vote-thresh})}$, to conclude the following.

Theorem 6.5. *Let $\epsilon,\delta>0$ be parameters, let $U \in [-1,1]^{n\times M}$ describe the preferences of n voters in M states of nature, and let $\lambda \in \Delta_M$ be the prior of states of nature. There is an algorithm with runtime $\text{poly}(n, M^{\delta^{-2}\ln(1/\epsilon\delta)})$ for computing a signaling scheme φ such that $F_\delta^{(\textit{vote-thresh})}(\varphi,U,\lambda) \geq OPT^{(\textit{vote-thresh})}(U,\lambda)-\epsilon$.*

Connection to Maximum Feasible Subsystem of Linear Inequalities

Turning our attention away from signaling, we note that $g^{(\text{vote-sum})}(Ax)$ simply counts the number of satisfied inequalities in the system $Ax \succeq 0$. Mixture selection for $g^{(\text{vote-sum})}$ is therefore the problem of maximizing the number of satisfied inequalities over the simplex. Using our framework from Section 5.2, we obtain a bi-criteria PTAS for this problem. Moreover, using Monte-Carlo sampling, our bi-criteria PTAS extends to the model in which A is given implicitly; specifically, the rows of A correspond to the sample space of a distribution $\mathcal{D}$ over $[-1, 1]^m$, and are weighted accordingly. In this implicit model, we can think of mixture selection for $g^{(\text{vote-sum})}$ as the problem of finding $x \in \Delta_M$ which maximizes the probability that $a \cdot x \geq 0$ for $a \sim \mathcal{D}$.

6.2.3 Hardness Results for Persuading Voters

In Section 5.4.1, we showed that the posterior selection problem for $g^{(\text{vote-sum})}$ does not admit a (uni-criteria) additive PTAS unless P = NP. Inspired by our reduction, Dughmi and Xu [51] ruled out a (uni-criteria) PTAS for the problem of signaling to maximize the expected number of votes. Both reductions construct a Bayesian voting instance $\mathcal{I}$ from a graph G. At a high level, we showed that a good posterior of $\mathcal{I}$ corresponds to a large independent set of G; and [51] showed that a near-optimal signaling scheme of $\mathcal{I}$ corresponds to covering G using large independent sets. Conceptually, the idea of switching from maximum independent set to graph-coloring is equivalent to moving from planted clique to planted clique cover.

6.3 Computing Equilibria in Anonymous Games

The complexity and efficient approximation of Nash equilibria have been studied intensively during the past decade, and much progress has been made (e.g., see [1, 9, 10, 11, 12, 23, 24, 25, 33, 35, 39, 53, 70, 74, 75, 87, 88, 89]). Despite much effort, the computational complexity of approximate Nash equilibria in anonymous games remains open.

In recent years, equilibrium computation in anonymous games has attracted significant attention in TCS [25, 33, 36, 37, 38, 39, 41, 60]. Consider the family of anonymous games where the number of players, n, is large and the number of strategies, k, is bounded. It was recently shown by Chen et al. [25] that computing an ϵ-approximate Nash equilibrium of such games is PPAD-complete when ϵ is exponentially small, even for anonymous games with 5 strategies[2].

On the algorithmic side, Daskalakis and Papadimitriou [36, 37] presented the first polynomial-time approximation scheme (PTAS) for this problem with running time $n^{(1/\epsilon)^{\Omega(k)}}$. For the case of 2-strategies, this bound was improved [34, 38, 39] to $\mathrm{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$, and subsequently sharpened to $\mathrm{poly}(n) \cdot (1/\epsilon)^{O(\log(1/\epsilon))}$ in [42].

In recent work, Daskalakis *et al.* [33] and Diakonikolas *et al.* [41] generalized the aforementioned results [39, 42] to any fixed number k of strategies, obtaining algorithms for computing ϵ-well-supported Nash equilibria (see Definition 2.3) with runtime of the form $n^{\mathrm{poly}(k)} \cdot (1/\epsilon)^{k \log(1/\epsilon)^{O(k)}}$. That is, the problem of computing approximate Nash equilibria in anonymous games with a fixed number of strategies admits an *efficient* polynomial-time approximation scheme (EPTAS). Moreover, the dependence of the running time on the parameter $1/\epsilon$ is *quasipolynomial* — as opposed to exponential.

We note that all the aforementioned algorithmic results are obtained by exploiting a connection between Nash equilibria in anonymous games and Poisson multinomial distributions (PMDs). This connection — formalized in [36, 37] — translates constructive upper bounds on ϵ-covers for PMDs to upper bounds on

[2] [25] showed that computing an equilibrium of 7-strategy anonymous games is PPAD-complete, but 3 of the 7 strategies in their construction can be merged, resulting in a 5-strategy anonymous game.

computing ϵ-Nash equilibria in anonymous games (see Section 2.1.7 for formal definitions). Unfortunately, as shown in [33, 41], this "cover-based" approach cannot lead to qualitatively faster algorithms, due to a matching existential lower bound on the size of the corresponding ϵ-covers. In related algorithmic work, Goldberg and Turchetta [60] studied two-strategy anonymous games ($k = 2$) and designed a polynomial-time algorithm that computes an ϵ-approximate Nash equilibrium for $\epsilon = \Omega(n^{-1/4})$.

The aforementioned discussion prompts the following natural question: *What is the precise approximability of computing Nash equilibria in anonymous games?* In this chapter, we make progress on this question by establishing the following result: For any $\delta > 0$, and any n-player anonymous game with a constant number of strategies, there exists a $\text{poly}_\delta(n)$ time algorithm that computes an ϵ-approximate Nash equilibrium of the game, for $\epsilon = 1/n^{1-\delta}$.[3] Moreover, we show that the existence of a polynomial-time algorithm that computes an ϵ-approximate Nash equilibrium for $\epsilon = 1/n^{1+\delta}$, for any small constant $\delta > 0$ — i.e., slightly better than the approximation guarantee of our algorithm — would imply the existence of a fully polynomial-time approximation scheme (FPTAS) for the problem. That is, we essentially show that the value $\epsilon = 1/n$ is the threshold for the polynomial-time approximability of Nash equilibria in anonymous games, unless there is an FPTAS. In the following subsection, we describe our results in detail and provide an overview of our techniques.

---

[3] The runtime of our algorithm depends exponentially on $1/\delta$. We remind the reader that the algorithms of [33, 41] run in quasipolynomial time for any value of ϵ inverse polynomial in n.

6.3.1 Summary of Results and Techniques

We study the following question:

For n-player k-strategy anonymous games, how small can ϵ be (as a function of n), so that an ϵ-approximate Nash equilibrium can be computed in polynomial time?

Upper Bounds. We present a polynomial time algorithm that computes ϵ-approximate equilibria in anonymous games for an inverse polynomial ϵ above a certain threshold.

Theorem 6.6. *For any $\delta > 0$, and any n-player k-strategy anonymous game, there is a $\text{poly}_{\delta,k}(n)$ time algorithm that computes a $(1/n^{1-\delta})$-approximate equilibrium of the game.*

This is the first polynomial time ϵ-approximation for some $k > 2$ strategies and some inverse polynomial ϵ.

Overview of Techniques. The high-level idea of our approach is this: If the desired accuracy ϵ is above a certain threshold, we do not need to enumerate over an ϵ-cover for the set of all PMDs. Our approach is in part inspired by [60], who design an algorithm (for $k = 2$ and $\epsilon = \Omega(n^{-1/4})$) in which all players use one of the two pre-selected mixed strategies. The [60] algorithm can be equivalently interpreted as guessing a PBD from an appropriately small set. One reason this idea succeeds is the following: If every player randomizes, then the variance of the resulting PBD must be relatively high, and (as a result) the corresponding subset of PBDs has a smaller cover.

Our quantitative improvement for the $k = 2$ case is obtained as follows: Instead of forcing players to selected specific mixed strategies — as in [60] — we show that there always exists an ϵ-approximate equilibrium where the associated PBD has variance at least $\Theta(n\epsilon)$. When $\epsilon = n^{-c}$ for some $c < 1$, the variance is polynomial in n. We then construct a polynomial-size ϵ-cover for the subset of PBDs with variance at least this much, which leads to a polynomial-time algorithm for computing ϵ-approximate equilibria in 2-strategy anonymous games.

The idea for the general case of $k > 2$ is similar, but the details are more elaborate, since the structure of PMDs is more complicated for $k > 2$. We proceed as follows: we start by showing that there is an ϵ-approximate equilibrium whose corresponding PMD has large variance in each direction. Our main structural result is a robust moment-matching lemma (Lemma 6.11), which states that the closeness in low-degree moments of two PMDs, with large variance in each direction, implies their closeness in total variation distance. The proof of this lemma uses Fourier analytic techniques, building on and strengthening previous work [41]. As a consequence of our moment-matching lemma, we can construct a polynomial-size $(\epsilon/5)$-cover for PMDs with such large variance. We then iterate through this cover to find an ϵ-approximate equilibrium, using a dynamic programming approach similar to the one in [39].

We now provide a brief intuition of our moment-matching lemma. Intuitively, if the two PMDs in question are both very close to discrete Gaussians, then closeness in the first *two* moments is sufficient. Lemma 6.11 can be viewed as a generalization of this intuition, which gives a quantitative tradeoff between the number of moments we need to approximately match and the size of the variance. The proof of Lemma 6.11 exploits the sparsity of the Fourier transform of our PMDs, and the fact that higher variance allows us to take fewer terms in the Taylor expansion when we use moments to approximate the logarithmic Fourier transform. This completes the proof sketch of Theorem 6.6.

Lower Bounds. When $\epsilon = 1/n$, we can show that there is an ϵ-approximate equilibrium where the associated PMD has a variance at least $1/k$ in every direction. Unfortunately, the PMDs in the explicit quasipolynomial-size lower bounds given in [33, 41] satisfy this property. Thus, we need a different approach to get a polynomial-time algorithm for $\epsilon = 1/n$ or smaller.

In fact, we prove the following result, which states that even a slight improvement of our upper bound in Theorem 6.6 would imply an FPTAS for computing Nash equilibria in anonymous games. It is important to note that Theorem 6.7 applies to all algorithms, not only the ones that leverage the structure of PMDs.

Theorem 6.7. *For n-player k-strategy anonymous games with $k = O(1)$, if we can compute an $O(n^{-c})$-approximate equilibrium in polynomial time for some constant $c > 1$, then there is an FPTAS for computing (well-supported) Nash equilibria of k-strategy anonymous games.*

Remark. As observed in [33], because there is a quasipolynomial time algorithm for computing an (n^{-c})-approximate equilibrium in anonymous games, the problem cannot be PPAD-complete unless PPAD $\subseteq$ Quasi-PTIME. On the other hand, we do not know how to improve the quasipolynomial-time upper bounds of [33, 41] when $\epsilon < 1/n$.

Recall that computing an ϵ-approximate equilibrium of a two-player general-sum $n \times n$ game (2-NASH) for constant ϵ also admits a quasipolynomial-time algorithm [74]. Very recently, Rubinstein [88] showed that, assuming the exponential time hypothesis (ETH) for PPAD, for some sufficiently small universal constant $\epsilon > 0$, quasipolynomial-time is necessary to compute an ϵ-approximate equilibrium of 2-NASH. It is a plausible conjecture that quasipolynomial-time is also required for ϵ-Nash equilibria in anonymous games, when $\epsilon = n^{-c}$ for some constant $c > 1$. In particular, this would imply that there is no FPTAS for computing approximate Nash equilibria in anonymous games, and consequently the upper bound of Theorem 6.6 is essentially tight.

6.3.2 Searching Fewer Moments

In this section, we present a polynomial-time algorithm that, for n-player anonymous games with a bounded number of strategies, computes an ϵ-approximate equilibrium with $\epsilon = n^{-c}$ for any constant $c < 1$ (Theorem 6.6). Theorem 6.6 applies to general k-strategy anonymous games for any constant $k \geq 2$. As a warm-up, we start by describing the simpler setting of two-strategy anonymous games ($k = 2$).

Lemma 6.8. *For an n-player k-strategy anonymous game, there always exists an ϵ-approximate equilibrium where every player plays each strategy with probability at least $\frac{\epsilon}{k-1}$.*

Proof. Given an anonymous game $G = (n, k, \{u_a^i\}_{i\in[n],a\in[k]})$, we smooth players' utility functions by requiring every player to randomize. Fix $\epsilon > 0$. We define an ϵ-perturbed game G_ϵ as follows. When a player plays some pure strategy $a \in [k]$ in G_ϵ, we map it back to the original game as if she played strategy j with probability $1 - \epsilon$, and played some other strategy $a' \neq a$ uniformly at random (i.e., she plays a' with probability $\frac{\epsilon}{k-1}$). Her payoff in G_ϵ also accounts for such perturbation, and is defined to be her expected payoff given that all the players (including herself) would deviate to other strategies uniformly at random with probability ϵ.

Formally, let $X_\epsilon(e_j)$ denote the k-CRV that takes value e_j with probability $1-\epsilon$, and takes value $e_{j'}$ with probability $\frac{\epsilon}{k-1}$ for each $j' \neq j$. The payoff structure of G_ϵ is given by

$$u'^i_a(x) \stackrel{\text{def}}{=} (1-\epsilon)\mathbb{E}\Big[u^i_a(M_\epsilon(x))\Big]+\frac{\epsilon}{k-1}\sum_{a'\neq a}\mathbb{E}\Big[u^i_{a'}(M_\epsilon(x))\Big], \quad \forall i \in [n], a \in [k], x \in \Pi^k_{n-1},$$

where $M_\epsilon(x) = \sum_{j\in[k]} x_j X_\epsilon(e_j)$ is an $(n-1,k)$-PMD that corresponds to the perturbed outcome of the partition $x \in \Pi^k_{n-1}$ of all other players.

Let $s' = (s'_1, \ldots, s'_n)$ denote any exact Nash equilibrium of G_ϵ. We can interpret this mixed strategy profile in G equivalently as $s = (s_1, \ldots, s_n)$, where $s_i = (1 - \frac{k\epsilon}{k-1})s'_i + \frac{\epsilon}{k-1}\mathbb{1}$, where $\mathbb{1} = (1, \ldots, 1)$. We know that under s each player has no incentive to deviate to the mixed strategies $X_\epsilon(e_j)$ for all $j \in [k]$, therefore a player can gain at most ϵ by deviating to pure strategies in G, so s is an ϵ-approximate equilibrium with $s_i(j) \geq \frac{\epsilon}{k-1}$ for all $i \in [n]$, $j \in [k]$. $\square$

Warm-up: The Case of $k = 2$ Strategies. For two-strategy anonymous games ($k = 2$), if all the players put at least ϵ probability mass on both strategies, the resulting PBD is going to have variance at least $n\epsilon(1-\epsilon)$. When $\epsilon = n^{-c}$ for some constant $c < 1$, the variance is at least $\Theta(n^{1-c}) = n^{\Theta(1)}$. We can now use the following lemma from [43], which states that if two PBDs P and Q are close in the first few moments, then P and Q are ϵ-close in total variation distance. Note that without any assumption on the variance of the PBDs, we would need to check the first $O(\log(1/\epsilon))$ moments, but when the variance is $n^{\Omega(1)}$, which is the case in our application, we only need the first constant number of moments to match.

Recall that an n-PBD is the sum of n independent Bernoulli random variables. An n-PBD P can be represented by its n parameters $p_1, \ldots, p_n$, where p_i is the probability of the i-th Bernoulli takes the value of 1. In the following lemma, for technical reasons, these parameters are partitioned into two sets with s and s' elements ($s + s' = n$), depending on whether they are greater than 1/2 or not.

Lemma 6.9 ([43]). *Let $\epsilon > 0$. Let P and Q be n-PBDs with P having parameters $p_1, \ldots, p_s \leq 1/2$ and $p'_1, \ldots, p'_{s'} > 1/2$, and Q having parameters $q_1, \ldots, q_s \leq 1/2$ and $q'_1, \ldots, q'_{s'} > 1/2$. Suppose that $V = \mathrm{Var}[P] + 1 = \Theta(\mathrm{Var}[Q] + 1)$ and let $C > 0$ be a sufficiently large constant. Suppose furthermore that the following holds for $A = C\sqrt{\log(1/\epsilon)/V}$ and for all positive integers ℓ,*

$$A^\ell \left(\left| \sum_{i=1}^{s} p_i^\ell - \sum_{i=1}^{s} q_i^\ell \right| + \left| \sum_{i=1}^{s'} (1 - p'_i)^\ell - \sum_{i=1}^{s'} (1 - q'_i)^\ell \right| \right) < \frac{\epsilon}{C \log(1/\epsilon)} \tag{6.1}$$

Then $\mathrm{d_{TV}}(P, Q) < \epsilon$.

Let $\epsilon = n^{-c}$. For Lemma 6.9 we have $V \geq n\epsilon(1 - \epsilon)$ and $A = \Theta\Big(\sqrt{\log(1/\epsilon)/V}\Big) = O\Big(\sqrt{\frac{\log n}{n^{1-c}}}\Big)$. The difference in the moments of parameters of P and Q in Equation (6.1) is bounded from above by n, so whenever $\ell > \frac{2+2c}{1-c}$, the condition in Lemma 6.9 is automatically satisfied for sufficiently large n because

$$A^\ell n = O\left(\frac{\log^{\ell/2} n}{n^{(1-c)\ell/2}} n \right) < \frac{1}{C \cdot n^c \cdot c \log n} = \frac{\epsilon}{C \log(1/\epsilon)}.$$

So it is enough to search over the first $\ell = \Theta\Big(\frac{1}{1-c}\Big)$ moments when each player puts probability at least $\Omega(n^{-c})$ on both strategies. The algorithm for finding such an ϵ-approximate equilibrium uses moment search and dynamic programming, and is given for the case of general k in the remainder of this section.

The General Case: k Strategies. We now present our algorithm for n-player anonymous games with $k > 2$ strategies and prove Theorem 6.6. The intuition of the $k = 2$ case carries over to the general case, but the details are more elaborate. First, we show (Claim 6.10) that there exists an ϵ-approximate equilibrium whose corresponding PMD has variance $(n\epsilon/k)$ in all directions orthogonal to the vector $\mathbf{1} = (1, \ldots, 1)$. Then, we prove (Lemma 6.11) that when two PMDs have such high variances, the closeness in their constant-degree parameter moments translates to their closeness in total variation distance. This structural result allows us to construct a polynomial-size $(\epsilon/5)$-cover for set subset of all PMDs with large variance. We then iterate through this cover to find an ϵ-approximate equilibrium (Algorithm 6.2).

We first prove that when all players put probability at least $\frac{\epsilon}{k-1}$ on each strategy, the covariance matrix of the resulting PMD has relatively large eigenvalues, except the zero eigenvalue associated with the all-one eigenvector. The all-one eigenvector has eigenvalue zero because the coordinates of X always sum to n.

Claim 6.10. *Let $X = \sum_{i=1}^{n} X_i$ be an (n, k)-PMD and let Σ be the covariance matrix of X. If $p_{i,j} = \Pr[X_i = e_j] \geq \frac{\epsilon}{k-1}$ for all $i \in [n]$ and $j \in [k]$, then all eigenvalues of Σ but one are at least $\frac{n\epsilon}{k-1}$.*

Proof. Fix any unit vector $v \in \mathbb{R}^k$ that is orthogonal to the all-one vector $\mathbb{1}$, i.e., $\sum_j v_j = 0$ and $\sum_j v_j^2 = 1$. Together with the assumption that $p_{i,j} \geq \frac{\epsilon}{k-1}$, we have

$$\begin{aligned}
\mathrm{Var}[v^T X_i] &= \mathbb{E}\Big[\Big(v^T X_i - \mathbb{E}\big[(v^T X_i)\big]\Big)^2\Big] \\
&= \sum_{j=1}^{n} p_{i,j} \left(v_j - \sum_{j'=1}^{n} p_{i,j'} v_{j'} \right)^2 \\
&\geq \min_j \{p_{i,j}\} \cdot \sum_{j=1}^{n} \left(v_j^2 + \left(\sum_{j'=1}^{n} p_{i,j'} v_{j'} \right)^2 - 2 v_j \left(\sum_{j'=1}^{n} p_{i,j'} v_{j'} \right) \right) \\
&= \min_j \{p_{i,j}\} \cdot \left(1 + n \left(\sum_{j'=1}^{n} p_{i,j'} v_{j'} \right)^2 \right) \\
&\geq \frac{\epsilon}{k-1}.
\end{aligned}$$

Therefore,

$$v^T \Sigma v = \mathrm{Var}[v^T X] = \sum_{i=1}^{n} \mathrm{Var}[v^T X_i] \geq \frac{n\epsilon}{k-1}.$$

So, for all eigenvectors v orthogonal to $\mathbb{1}$, we have $v^T \Sigma v = \lambda v^T v = \lambda \geq \frac{n\epsilon}{k-1}$ as claimed. □

We recall some of the notations for readability before we describe the construction of our ϵ-cover of high-variance PMDs. We use X to denote a generic (ℓ, k)-PMD for some $\ell \in [n]$, and we denote $p_{i,j} = \Pr[X_i = e_j]$. We use $A_t \subseteq [\ell]$ to denote the set of t-maximal CRVs in X, where a k-CRV is t-maximal if e_t is its most likely outcome, and we use $X^t = \sum_{i \in A_t} X_i$ to denote the t-maximal component PMD of X. For a vector $m = (m_1, \ldots, m_k) \in \mathbb{Z}_+^k$, we define m^{th} parameter moment of X^t to be $M_m(X^t) = \sum_{i \in A_t} \prod_{j=1}^{k} p_{i,j}^{m_j}$. We refer to $\|m\|_1 = \sum_{j=1}^{k} m_j$ as the *degree* of $M_m(X)$. We use $\mathcal{S}$ to denote the set of all k-CRVs whose probabilities are multiples of $\frac{\epsilon}{20kn}$.

The following robust moment-matching lemma provides a bound on how close degree-ℓ moments need to be so that two (n, k)-PMDs are ϵ-close to each other, under the assumption that $n \gg k$ (the anonymous game has many players and few strategies) and $p_{i,j} \geq \frac{\epsilon}{k-1}$ (every player randomizes). Lemma 6.11 allows us to build a polynomial-size $(\epsilon/5)$-cover for PMDs with high variance, and since we know that there is an ϵ-approximate equilibrium with a high variance, we are guaranteed to find one in our cover.

Lemma 6.11. *Fix $0 < c < 1$ and let $\epsilon = n^{-c}$. Assume that $n \geq k^{\Theta(k)}$ for some sufficiently large constant in the exponent. Let X, Y be (n, k)-PMDs with $X = \sum_{i=1}^k X^i$, $Y = \sum_{i=1}^k Y^i$ where each X^i, Y^i is an i-maximal PMD. Let Σ_X and Σ_Y denote the covariance matrices of X and Y respectively. Suppose all non-zero eigenvalues of Σ_X, Σ_Y are at least $\epsilon n/k$, and all the parameter moments m of degree $\ell \leq \frac{2+2c}{1-c}$ satisfy that*

$$\left|M_m(X^i) - M_m(Y^i)\right| \leq \epsilon.$$

Then, we have that $\mathrm{d_{TV}}(X, Y) \leq \epsilon$.

Lemma 6.11 follows from Proposition 6.12.

Proposition 6.12. *Let $\epsilon > 0$. Let X, Y be (n, k)-PMDs with $X = \sum_{i=1}^k X^i$, $Y = \sum_{i=1}^k Y^i$ where each X^i, Y^i is an i-maximal PMD. Let Σ_X and Σ_Y denote the covariance matrices of X and Y respectively, where all eigenvalues of Σ_X and Σ_Y but one are at least σ^2, where $\sigma \geq \mathrm{poly}(k \log(1/\epsilon))$. Suppose that for $1 \leq i \leq k$, $\ell \geq 1$, for all moments m of degree ℓ with $m_i = 0$, we have that*

$$\left|M_m(X^i) - M_m(Y^i)\right| \leq \frac{\epsilon \cdot \sigma^\ell}{C'^{k+\ell} \cdot k^{3\ell/2+1} \cdot \log^{k+\ell/2}(1/\epsilon)}$$

for a sufficiently large constant C'. Then $\mathrm{d_{TV}}(X, Y) \leq \epsilon$.

The proof of Proposition 6.12 exploits the sparsity of the continuous Fourier transform of our PMDs, as well as careful Taylor approximations of the logarithm of the Fourier transform. We defer the proof of Proposition 6.12 to the next section.

Proof of Lemma 6.11 from Proposition 6.12. To guarantee that $\mathrm{d_{TV}}(X, Y) \le \epsilon$, Proposition 6.12 requires the following condition to hold for a sufficiently large constant C':

$$\left|M_m(X^i) - M_m(Y^i)\right| \le \frac{\epsilon}{k(C' \log(1/\epsilon))^k} \cdot \left(\frac{\sqrt{\epsilon n/k}}{C' k^{3/2} \log^{1/2}(1/\epsilon)}\right)^{\ell}, \quad \forall i \in [k], \ell \ge 1. \tag{6.2}$$

To prove the lemma, we use the fact that $n \gg k$ and essentially ignore all the terms except polynomials of n. Formally, we first need to show that

$$\epsilon \le \frac{\epsilon}{k(C' \log(1/\epsilon))^k} \cdot \left(\frac{\sqrt{\epsilon n/k}}{C' k^{3/2} \log^{1/2}(1/\epsilon)}\right)^{\ell}, \quad \forall \ell \ge 1,$$

under the assumption that $c < 1$, $\epsilon = n^{-c}$ and $n \ge k^{O(k/(1-c))}$. After substituting $\epsilon = n^{-c}$, observe that $n^{1-c} \ge C'^2 k^4 \log n$, so the term inside the ℓ-th power is greater than 1. Thus, we only need to check this inequality for $\ell = 1$, which simplifies to $n^{1-c} \ge C'^{2k+2} k^6 (\log n)^{2k}$ and holds true.

In addition, we need to show that condition (6.2) holds automatically for $\ell > \frac{2+2c}{1-c}$. This follows from the fact that the difference in parameter moments is at most n and $n \gg k$,

$$\left|M_m(X^i) - M_m(Y^i)\right| \le n \le \frac{\epsilon}{k(C' \log(1/\epsilon))^k} \cdot \left(\frac{\sqrt{\epsilon n/k}}{C' k^{3/2} \log^{1/2}(1/\epsilon)}\right)^{\ell}, \quad \forall \ell > \frac{2+2c}{1-c}.$$

□

Lemma 6.11 states that the high-degree parameter moments match automatically, which allows us to impose an appropriate grid on the low-degree moments to cover the set of high-variance PMDs. The size of this cover can be bounded by a simple counting argument: We have at most $k^{O(\frac{1}{1-c})}$ moments with degree at most $O(\frac{1}{1-c})$, and we need to approximate these moments for each t-maximal component PMDs, so there are at most $k \cdot k^{O(\frac{1}{1-c})} = k^{O(\frac{1}{1-c})}$ moments $M_m(X^t)$ that we care about. We approximate these moments to precision $\epsilon = n^{-c}$, and the moments have value at most n, so the size of the cover is $\left(\frac{n}{n^{-c}}\right)^{k^{O(\frac{1}{1-c})}} = n^{k^{O(1/1-c)}}$.

We define this grid on low-degree moments formally in the following lemma. For every (ℓ, k)-PMD X with $\ell \in [n]$, we associate some *data* $D(X)$ with X, which is a vector of the approximate values of the low-degree moments $M_m(X^t)$ of X.

Lemma 6.13. *Fix $0 < c < 1$ and let $\epsilon = n^{-c}$. Assume that $n \geq k^{\Theta(k)}$ for some sufficiently large constant in the exponent. We define the data $D(W)$ of a k-CRV W as follows:*

$$D(W)_{m,t} = \begin{cases} M_m(W) \text{ rounded to the nearest integer multiple of } \epsilon n, & \text{if } W \text{ is } t\text{-maximal.} \\ 0, & \text{otherwise.} \end{cases}$$

For $\ell \in [n]$, we define the data of an (ℓ, k)-PMD $X = \sum_{i=1}^{\ell} X_i$ to be the sum of the data of its k-CRVs: $D(X) = \sum_{i=1}^{\ell} D(X_i)$. The data $D(X)$ satisfies two properties:

1. *(Representative) If $D(X) = D(Y)$ for two (n, k)-PMDs (or two $(n-1, k)$-PMDs) X and Y, then $\mathrm{d_{TV}}(X, Y) \leq \epsilon$.*

2. *(Extensible) For independent PMDs X and Y, we have that $D(X+Y) = D(X) + D(Y)$.*

Proof. The "extensible" property follows directly from the definition of $D(X)$. To see the "representative" property, note that we round $M_m(W)$ to the nearest integer multiple of ϵn, so the error in the moments of W is at most $\epsilon n/2$. When we add up the data of an (n, k)-PMD or $(n-1, k)$-PMD, the error in the moments of each t-maximal component PMDs is at most $\epsilon/2$. So if two PMDs X and Y have the same data, their low-degree moments differ by at most ϵ, and then by Lemma 6.11 we have $d_{TV}(X, Y) \leq \epsilon$. ☐

Algorithm 6.1: GenerateData

```
Input  : {S_i}_{i=1}^n, ε > 0.
Output: The set of all possible data 𝒟 of (n,k)-PMDs X = Σ_{i=1}^n X_i where
        X_i ∈ S_i.
𝒟_0 = {};
for ℓ = 1 ... n do
    forall the D ∈ 𝒟_{ℓ−1} do
        forall the W ∈ S_ℓ do
            Add D + D(W) to 𝒟_ℓ if it is not in 𝒟_ℓ already;
            Keep track of an (ℓ,k)-PMD whose data is D + D(W);
return 𝒟 = 𝒟_n;
```

Our algorithm (Algorithm 6.2) for computing approximate equilibria is similar to the approach used in [39] and [41]. We start by constructing a polynomial-size $(\epsilon/5)$-cover of high-variance PMDs (Algorithm 6.1), and then iterate over this cover. For each element in the cover, we compute the set of $(3\epsilon/5)$-best-responses for each player, and then run the cover construction algorithm again, but this time we only allow each player to choose from her $(3\epsilon/5)$-best-responses. If we can reconstruct a PMD whose moments are close enough to the one we started with, then we have found an ϵ-approximate Nash equilibrium.

Algorithm 6.2: Moment Search

Input : An n-player k-strategy anonymous game G, $\epsilon = n^{-c}$ for some $c < 1$.
Output: An ϵ-approximate Nash equilibrium of G.

$\mathcal{S} = \{$all k-CRVs whose probabilities are multiples of $\frac{\epsilon}{20kn}\}$;
$\mathcal{D}_n = \text{GenerateData}(\{\mathcal{S}_i = \mathcal{S}\}_{i=1}^{n}, \epsilon/5)$;
$\mathcal{D}_{n-1} = \text{GenerateData}(\{\mathcal{S}_i = \mathcal{S}\}_{i=1}^{n-1}, \epsilon/5)$;
forall the $D \in \mathcal{D}_n$ **do**
  Set $\mathcal{S}_i = \emptyset$ for all i;
  forall the $X_i \in \mathcal{S}$ **do**
    Let $D_{-i} = D - D(X_i)$;
    if $\exists Y_{D_{-i}} \in \mathcal{D}_{n-1}$ *with* $D(Y_{D_{-i}}) = D_{-i}$ *and* X_i *is a* $(3\epsilon/5)$-*best response to* $Y_{D_{-i}}$ **then**
      Add X_i to $\mathcal{S}_i$;
  $\mathcal{D}'_n = \text{GenerateData}(\{\mathcal{S}_i\}_{i=1}^{n}, \epsilon/5)$;
  if $D \in \mathcal{D}'_n$ **then**
    return $(X_1, \ldots, X_n)$ in $\mathcal{D}'_n$ with $D\left(\sum_{i=1}^{n} X_i\right) = D$

Recall that a mixed strategy profile for a k-strategy anonymous game can be represented as a list of k-CRVs $(X_1, \ldots, X_n)$, where X_i describes the mixed strategy of player i. Recall that $(X_1, \ldots, X_n)$ is an ϵ-approximate Nash equilibrium if for each player i we have $\mathbb{E}\left[u^i_{X_i}(X_{-i})\right] \geq \mathbb{E}[u^i_a(X_{-i})] - \epsilon$ for all $a \in [k]$, where $X_{-i} = \sum_{j \neq i} X_j$ is the distribution of the sum of other players strategies.

Lemma 6.14. *Fix an anonymous game $G = (n, k, \{u^i_a\}_{i \in [n], a \in [k]})$ with payoffs normalized to $[0,1]$. Let $(X_1, \ldots, X_n)$ and $(Y_1, \ldots, Y_n)$ be two lists of k-CRVs. If X_i is a δ-best response to X_{-i}, and $\mathrm{d_{TV}}(X_{-i}, Y_{-i}) \leq \epsilon$, then X_i is a $(\delta + 2\epsilon)$-best response to Y_{-i}. Moreover, if $(X_1, \ldots, X_n)$ is a δ-approximate equilibrium, and $\mathrm{d_{TV}}(X_i, Y_i) + \mathrm{d_{TV}}(X_{-i}, Y_{-i}) \leq \epsilon$ for all $i \in [n]$, then $(Y_1, \ldots, Y_n)$ is a $(\delta + 2\epsilon)$-approximate equilibrium.*

Proof. Since $u_a^i(x) \in [0,1]$ for all $a \in [k]$ and $x \in \Pi_{n-1}^k$, we have that

$$\left|\mathbb{E}\left[u_a^i(X_{-i})\right] - \mathbb{E}\left[u_a^i(Y_{-i})\right]\right| \leq \mathrm{d_{TV}}(X_{-i}, Y_{-i}), \quad \forall i \in [n], a \in [k].$$

Therefore, if $\mathrm{d_{TV}}(X_{-i}, Y_{-i}) \leq \epsilon$, and player i cannot deviate and gain more than δ when other players play X_{-i}, then she cannot gain more than $(\delta + 2\epsilon)$ when other players play Y_{-i} instead of X_{-i}. The second claim combines the inequality above with the fact that, if player i plays Y_i instead of X_i and the mixed strategies of other players remain the same, her payoff changes by at most $\mathrm{d_{TV}}(X_i, Y_i)$. Formally,

$$\left|\mathbb{E}\left[u_{X_i}^i(Z_{-i})\right] - \mathbb{E}\left[u_{Y_i}^i(Z_{-i})\right]\right| \leq \mathrm{d_{TV}}(X_i, Y_i), \quad \forall k\text{-CRV } X_i, Y_i, \forall (n-1,k)\text{-PMD } Z_{-i}.$$

□

The next lemma states that there exists an $(\epsilon/5)$-approximate equilibrium whose probabilities are all integer multiples of $\frac{\epsilon}{20kn}$.

Claim 6.15. *There is an $(\epsilon/5)$-approximate Nash equilibrium $(X_1, \ldots, X_n)$, such that for all $i \in [n]$ and $j \in [k]$, the probabilities $p_{i,j} = \Pr[X_i = e_j]$ are multiples of $\frac{\epsilon}{20kn}$, and also $p_{i,j} \geq \frac{\epsilon}{10k}$.*

Proof. We start with an $(\epsilon/10)$-approximate Nash equilibrium $(Y_1, \ldots, Y_n)$ from Lemma 6.8 with $p_{i,j} \geq \frac{\epsilon}{10k}$, and then round the probabilities to integer multiples of $\frac{\epsilon}{10kn}$. We construct X_i from Y_i as follows: for every $j < k$, we set $\Pr[X_i = e_j]$ to be $\Pr[Y_i = e_j]$ rounded down to a multiple of $\frac{\epsilon}{20kn}$ and we set $\Pr[X_i = e_k] = 1 - \sum_{j<k} \Pr[X_i = e_j]$ so the probabilities sum to 1. By triangle inequality of total variation distance, for every i we have $\mathrm{d_{TV}}(X_i, Y_i) \leq \frac{\epsilon}{20n}$ and $\mathrm{d_{TV}}(X_{-i}, Y_{-i}) \leq \frac{\epsilon(n-1)}{20n}$. An application of Lemma 6.14 shows that $(X_1, \ldots, X_n)$ is an $(\epsilon/5)$-approximate equilibrium. □

We are now ready to prove Theorem 6.6.

Proof of Theorem 6.6. We show that for any n-player k-strategy anonymous game, if both $c > 0$ and k are constants, then there is a poly(n) time algorithm that computes an ϵ-approximate equilibrium for $\epsilon = 1/n^{1-c}$. If $n = k^{O(k)} = O(1)$, we use the algorithm in [36] which runs in time $n^{(1/\epsilon)^{\Omega(k)}} = O(1)$. So for the rest of the proof, we assume that $n \geq k^{\Theta(k)}$ as required in Lemma 6.11 and 6.13, and prove that Algorithm 6.2 always outputs an ϵ-approximate Nash equilibrium, and bound the running time.

We first show that the output $(X_1, \ldots, X_n)$ is an ϵ-approximate equilibrium. Recall that $\mathcal{S}$ is the set of all k-CRVs whose probabilities are multiples of $\frac{\epsilon}{20kn}$, and $\mathcal{S}_i \subseteq \mathcal{S}$ is the set of approximate best-responses of player i. When we put X_i in $\mathcal{S}_i$, we checked that X_i is a $(3\epsilon/5)$-best response to $Y_{D_{-i}}$. Note that $D(Y_{D_{-i}}) = D - D(X_i) = D(X_{-i})$, so by Lemma 6.13 $\mathrm{d_{TV}}\left(X_{-i}, Y_{D_{-i}}\right) \leq \epsilon/5$ for all i. By Lemma 6.14, X_i is indeed an ϵ-best response to X_{-i} for all i.

Next we show the algorithm must always output something. By Claim 6.15 there exists an $(\epsilon/5)$-approximate equilibrium X'_i with each $X'_i \in \mathcal{S}$. If the algorithm does not terminate successfully first, it eventually considers $D(X')$. Because X'_{-i} is an $(n-1, k)$-PMD, the algorithm can find some $Y_{D_{-i}}$ with $D(Y_{D_{-i}}) = D(X') - D(X'_i) = D(X'_{-i})$, and by Lemma 6.13 we have $\mathrm{d_{TV}}\left(X'_{-i}, Y_{D_{-i}}\right) \leq \epsilon/5$ for all i. Since X'_i is an $(\epsilon/5)$-best response to X'_{-i}, Lemma 6.14 yields that X'_i is a $(3\epsilon/5)$-best response to $Y_{D_{-i}}$, so we would add each X'_i to $\mathcal{S}_i$. Then our cover construction algorithm is guaranteed to generate a set of data that includes $D(X')$, and Algorithm 6.2 would produce an output.

Finally, we bound the running time of Algorithm 6.2. Let $N = O\Big(n^{k^{O(1/1-c)}}\Big)$ denote the size of the $(\epsilon/5)$-cover for the high-variance PMDs. The cover can be constructed in time $O(n \cdot N \cdot |S|)$ as we try to add one k-CRV from S in each step. We iterate through the cover, and for each element in the cover, we need to find the subset $\mathcal{S}_i \subseteq \mathcal{S}$ of $(3\epsilon/5)$-best responses for player i, and then run the cover construction algorithm again using only the best responses $\{\mathcal{S}_i\}_{i=1}^n$. So the overall running time of the algorithm is $O(nN|S|) \cdot \Big(\mathrm{poly}(n^k)|S| + O(nN|S|)\Big) = n^{k^{O(1/1-c)}}$. When both $0 < c < 1$ and k are constants, the running time is polynomial in n. □

6.3.3 A New Moment Matching Lemma

This subsection is devoted to the proof of Proposition 6.12. For two (n, k)-PMDs with variance at least σ^2 in each direction, Proposition 6.12 gives a quantitative bound on how close degree-ℓ moments need to be (as a function of ϵ, σ, k and ℓ, but independent of n), in order for the two PMDs to be ϵ-close in total variation distance.

The proof of Proposition 6.12 exploits the sparsity of the continuous Fourier transforms of our PMDs, as well as careful Taylor approximations of the logarithm of the Fourier transform. The fact that our PMDs have large variance enables us to take fewer low-degree terms in the Taylor approximation. For technical reasons, we split our PMD as the sum of k independent component PMDs, $X = \sum_{i=1}^k X^i$, where all the k-CRVs in the component PMD X^i are i-maximal. Because the Fourier transform of X is the product of the Fourier transforms of X^i, we can just bound the pointwise difference between the logarithms of the Fourier transforms of each component PMD. One technicality is that since we have no assumption on the variances of the component PMDs X^i, their Fourier transforms may not be sparse, so it is crucial that we bound this difference only on the effective support of the Fourier transform of the entire PMD.

We start by considering a set S that includes the effective support of X (and Y when we show that the means are close):

Lemma 6.16 (Essentially Corollary 5.3 of [41])**.** *Let X be an (n,k)-PMD with mean μ and covariance matrix Σ, such that all the non-zero eigenvalues of Σ are at least σ^2 where $\sigma \geq \mathrm{poly}(1/\epsilon)$. Let S be the set of points $x \in \mathbb{Z}^k$ where $(x-\mu)^T \mathbb{1} = 0$ and*

$$(x-\mu)^T(\Sigma+I)^{-1}(x-\mu) \leq (Ck\log(1/\epsilon))\ ,$$

for some sufficiently large constant C. Then, $X \in S$ with probability at least $1-\epsilon/2$, and

$$|S| = \sqrt{\det(\Sigma+I)} \cdot O(\log(1/\epsilon))^{k/2}.$$

Proof. Applying Lemma 5.2 of [41], we have that $(X-\mu)^T(\Sigma+I)^{-1}(X-\mu) = O(k\log(k/\epsilon))$ with probability at least $1-\epsilon$. The set of integer coordinate points in this ellipsoid is the set S. Note that $|S|$ is equal to the volume of $S' = \left\{ y \in \mathbb{R}^k : \exists x \in S \text{ with } \|y-x\|_\infty \leq 1/2 \right\}$, because S' is the disjoint union of cubes of volume 1, one for each integer point. But S' is again contained in an ellipsoid with $(y-\mu)^T(\Sigma+I)^{-1}(y-\mu) = O(k\log(k/\epsilon))$, so $|S| = \mathrm{Vol}(S') = \sqrt{\det(\Sigma+I)} \cdot O(\log(1/\epsilon))^{k/2}$. □

Next we show that $\widehat{X}$, the Fourier transform of X, has a relatively small effective support. We fold the effective support onto $[0,1]^k$ to obtain the set T. We use $[x]$ to denote the additive distance of $x \in \mathbb{R}$ to the closest integer, i.e., $[x] = \min_{x' \in \mathbb{Z}} |x - x'|$.

Lemma 6.17. *Let X be an (n,k)-PMD with mean μ and covariance matrix Σ, such that all the non-zero eigenvalues of Σ are at least σ^2 where $\sigma \geq \mathrm{poly}(k\log(1/\epsilon))$. Let S be as above. Let $\widehat{X}$ be the Fourier transform of X. Let $T \stackrel{\mathrm{def}}{=} \left\{ \xi \in [0,1]^k : \exists \xi' \in \xi + \mathbb{Z}^k \text{ with } \xi'^T \Sigma \xi' \leq Ck\log(1/\epsilon) \right\}$, for some sufficiently large constant C. Then, we have that*

(i) *For $\xi \in T$, and for all $1 \leq i,j \leq k$, $[\xi_i - \xi_j] \leq 2\sqrt{Ck\log(1/\epsilon)}/\sigma$.*

(ii) $\mathrm{Vol}(T)|S| = O(C\log(1/\epsilon))^k$.

(iii) $\int_{[0,1]^k \setminus T} \left|\widehat{X}(\xi)\right| d\xi \leq \epsilon/(2|S|)$.

Lemma 6.17 is a technical generalization of Lemma 5.5 of [41]. This lemma establishes that the contribution to the Fourier transform $\widehat{X}$ coming from points outside of T is negligibly small. We then use the sparsity of the Fourier transform to show that, if two PMDs have Fourier transforms that are pointwise sufficiently close within the effective support T, then the two PMDs are close in total variation distance.

Lemma 6.18. *Let X, Y, S, T be as above. If $\left|\widehat{X}(\xi) - \widehat{Y}(\xi)\right| \leq \epsilon(C'\log(1/\epsilon))^{-k}$ for all $\xi \in T$ and a sufficiently large constant C', then $\mathrm{d_{TV}}(X,Y) \leq \epsilon$.*

Proof. For any $x \in \mathbb{Z}^k$, taking the inverse Fourier transform, we have that $\Pr[X = x] = \int_{\xi \in [0,1]^k} e(-\xi \cdot x)\widehat{X}(\xi) d\xi$ and similarly $\Pr[Y = x] = \int_{\xi \in [0,1]^k} e(-\xi \cdot x)\widehat{Y}(\xi) d\xi$. Thus,

$$\begin{aligned}
|\Pr[X = x] - \Pr[Y = x]| &= \left| \int_{\xi \in [0,1]^k} e(-\xi \cdot x) \left(\widehat{X}(\xi) - \widehat{Y}(\xi) \right) d\xi \right| \\
&\leq \int_{\xi \in [0,1]^k} \left| \widehat{X}(\xi) - \widehat{Y}(\xi) \right| d\xi \\
&= \int_{\xi \in T} \left| \widehat{X}(\xi) - \widehat{Y}(\xi) \right| d\xi + \int_{\xi \in [0,1]^k \setminus T} \left| \widehat{X}(\xi) - \widehat{Y}(\xi) \right| d\xi \\
&\leq \mathrm{Vol}(T) \cdot \epsilon (C' \log(1/\epsilon))^{-k} + \frac{\epsilon}{2|S|} \\
&\leq \frac{O(C \log(1/\epsilon))^k}{|S|} \cdot \epsilon (C' \log(1/\epsilon))^{-k} + \frac{\epsilon}{2|S|} \\
&\leq \frac{\epsilon}{|S|}.
\end{aligned}$$

Since X and Y are outside of S each with probability less than $\epsilon/2$, we have that $d_{\mathrm{TV}}(X, Y) \leq \epsilon/2 + \frac{1}{2} \sum_{x \in S} |\Pr[X = x] - \Pr[Y = x]| \leq \epsilon$. ☐

We now have all the ingredients to prove Proposition 6.12. For two PMDs X and Y that are close in their low-degree moments, we show that their Fourier transforms $\widehat{X}$ and $\widehat{Y}$ are pointwise close on T, and then by Lemma 6.18, X and Y are close in total variation distance.

Proof of Proposition 6.12. Let X, Y, S, T be as above. Given Lemma 6.18, we only need to show that $\forall \xi \in T$, $\left| \widehat{X}(\xi) - \widehat{Y}(\xi) \right| \leq \epsilon (C' \log(1/\epsilon))^{-k}$.

Fix $\xi \in T$. We first examine, without loss of generality, the Fourier transform $\widehat{X^k}$ of the k-maximal component PMD. Let $A_k \subseteq [n]$ denote the set of k-maximal CRVs.

$$
\begin{aligned}
\widehat{X^k}(\xi) &= \prod_{i \in A_k} \sum_{j=1}^{k} e(\xi_j) p_{i,j} \\
&= e(|A_k|\xi_k) \prod_{i \in A_k} \left(1 - \sum_{j=1}^{k-1} (1 - e(\xi_j - \xi_k)) p_{i,j} \right) \\
&= e(|A_k|\xi_k) \exp \left(\sum_{i \in A_k} \log \left(1 - \sum_{j=1}^{k-1} (1 - e(\xi_j - \xi_k)) p_{i,j} \right) \right) \\
&= e(|A_k|\xi_k) \exp \left(- \sum_{i \in A_k} \sum_{\ell=1}^{\infty} \frac{1}{\ell} \left(\sum_{j=1}^{k-1} (1 - e(\xi_j - \xi_k)) p_{i,j} \right) \right) \\
&= e(|A_k|\xi_k) \exp \left(- \sum_{m \in \mathbb{Z}_+^{k-1}} \binom{\|m\|_1}{m} \frac{1}{\|m\|_1} M_m(X^k) \prod_{j=1}^{k-1} (1 - e(\xi_j - \xi_k))^{m_j} \right)
\end{aligned}
\tag{6.3}
$$

For notational convenience, we use Ψ_X^k to denote the expression inside $\exp(\cdot)$ in Equation (6.3). A similar formula holds for the Fourier transforms $\widehat{X^i}$ and $\widehat{Y^i}$ of other i-maximal PMDs, and we use Ψ_X^i and Ψ_Y^i to denote the corresponding expressions inside $\exp(\cdot)$. Since the Fourier transform of a PMD is the product of the Fourier transform of its component PMDs, we have

$$
\begin{aligned}
\left| \widehat{X}(\xi) - \widehat{Y}(\xi) \right| &= \left| \prod_{t=1}^{k} \widehat{X^t}(\xi) - \prod_{t=1}^{k} \widehat{Y^t}(\xi) \right| \\
&= \left| e\left(\sum_{t=1}^{k} |A_t| \xi_t \right) \prod_{t=1}^{k} \left(\exp\left(\Psi_X^t\right) - \exp\left(\Psi_Y^t\right) \right) \right| \\
&\leq 2\pi \sum_{t=1}^{k} \left| \Psi_X^t - \Psi_Y^t \right| ,
\end{aligned}
$$

where the last inequality is due to $e(\sum_{t=1}^{k}|A_t|\xi_t)=1$, and $|\exp(a)-\exp(b)|\le|a-b|$ if the real parts of a and b satisfy $\mathrm{Re}(a),\mathrm{Re}(b)\le 0$.

So to prove that $\widehat{X}(\xi)$ and $\widehat{Y}(\xi)$ are pointwise close for all $\xi\in T$, it is enough to bound from above $2\pi\sum_{t=1}^{k}|\Psi_X^t-\Psi_Y^t|$. We use the fact that $|1-e(\xi_j-\xi_k)|=O([\xi_j-\xi_k])$, and recall that $[\xi_i-\xi_j]\le 2\sqrt{Ck\log(1/\epsilon)}/\sigma$ by Lemma 6.17. We also use the multinomial identity $\sum_{m\in\mathbb{Z}_+^{k-1},\|m\|_1=\ell}\binom{\ell}{m}=(k-1)^\ell$. When C' is a sufficiently large constant, we have

$$
\begin{aligned}
&\left|\widehat{X}(\xi)-\widehat{Y}(\xi)\right|\\
&\le 2\pi\sum_{t=1}^{k}\left|\Psi_X^t-\Psi_Y^t\right|\\
&=2\pi\sum_{t=1}^{k}\sum_{m\in\mathbb{Z}_+^{k-1}}\binom{\|m\|_1}{m}\frac{1}{\|m\|_1}\left|M_m(X^t)-M_m(Y^t)\right|\prod_{j=1}^{k-1}(1-e(\xi_j-\xi_k))^{m_j}\\
&\le 2\pi\sum_{\ell=1}^{\infty}\frac{(k-1)^\ell}{\ell}\left(O\left(\frac{\sqrt{k\log(1/\epsilon)}}{\sigma}\right)\right)^\ell\sum_{t=1}^{k}\max_{m\in\mathbb{Z}_+^{k-1},\|m\|_1=\ell}\left|M_m(X^t)-M_m(Y^t)\right|\\
&\le\sum_{\ell=1}^{\infty}k^\ell\left(\frac{C'\sqrt{k\log(1/\epsilon)}}{2\sigma}\right)^\ell k\cdot\frac{\epsilon\sigma^\ell}{C'^{k+\ell}\cdot k^{3\ell/2+1}\cdot\log^{k+\ell/2}(1/\epsilon)}\\
&=\sum_{\ell=1}^{\infty}2^{-\ell}\epsilon(C'\log(1/\epsilon))^{-k}\\
&=\epsilon(C'\log(1/\epsilon))^{-k}.
\end{aligned}
$$

□

6.3.4 Slight Improvement Gives FPTAS

In this section, we show that even a slight improvement of our upper bound would imply an FPTAS for computing (well-supported) Nash equilibria in anonymous games (Theorem 6.7). It is a plausible conjecture that assuming the ETH for

PPAD, there is no such FPTAS, in which case our upper bound (Theorem 6.6) is essentially tight.

Theorem 6.7 follows directly from the following two lemmas. Lemma 6.19 converts an $\frac{\epsilon^2}{4n}$-approximate Nash equilibrium into an ϵ-well-supported Nash equilibrium (see Definition 2.3), by reallocating each player's probabilities on strategies with low expected payoffs to the best-response strategy (first observed in [35]). Lemma 6.20 then uses a padding argument to show that, for ϵ-*well-supported* Nash equilibria, the question of whether there is a polynomial-time algorithm for $\epsilon = n^{-c}$ is equivalent for all constants $c > 0$.

Lemma 6.19. *For any n-player game whose payoffs are normalized to be between* $[0, 1]$*, if we have an oracle for computing players' payoffs, we can efficiently convert an $\frac{\epsilon^2}{4n}$-approximate equilibrium into an ϵ-well-supported equilibrium.*

Proof. Take an $\frac{\epsilon^2}{4n}$-approximate equilibrium of the game. We call a strategy "good" for a player if the strategy is an $\frac{\epsilon}{2}$-best response for the player, and we call it "bad" otherwise. A player can put at most probability $\frac{\epsilon}{2n}$ on the "bad" strategies without violating the $\frac{\epsilon^2}{4n}$-approximate equilibrium condition. We move all the probabilities on "bad" strategies for all players to (any one of) their best responses simultaneously. After moving the probabilities, every player assigns non-zero probabilities only to the "good" strategies. Since the total probability we moved is at most $\frac{\epsilon}{2}$ and the payoffs are in $[0, 1]$, the previously "good" strategies ($\frac{\epsilon}{2}$-best responses) are now ϵ-best responses. □

Lemma 6.20. *For n-player k-strategy anonymous games with $k = O(1)$, if an $\frac{1}{n^\gamma}$-well-supported equilibrium can be computed in time $O(n^d)$ for constants $\gamma, d > 0$, then there is an FPTAS for computing approximate-well-supported Nash equilibria in anonymous games.*

Proof. Let ϵ be the desired quality of the well-supported equilibrium. If $\frac{1}{n^\gamma} \leq \epsilon$ we are done, so we assume n is smaller. We set $n' = (1/\epsilon)^{1/\gamma}$, so that $\frac{1}{n'^\gamma} = \epsilon$. Given an n-player anonymous game G, we build an n'-player anonymous game G' as follows: we add $n' - n$ dummy players, and give the dummy players utility 1 on strategy 1, and 0 on any other strategies so in any ϵ-well-supported equilibria, the dummy player must all play strategy 1 with probability 1. (Note that this is only true for ϵ-well-supported Nash equilibrium; in an ϵ-approximate Nash equilibrium, the dummy players can put ϵ probability elsewhere.) We shift the utility function of the actual players to ignore the dummy players on strategy 1. Formally, the payoff structure of G' is given by:

- For each $i > n$,

$$u'^i_a(x) = \begin{cases} 1 & \text{if } a = 1 \\ 0 & \text{otherwise} \end{cases}$$

- For each $i \leq n$, we subtract the number of players on strategy 1 by $n' - n$ and then apply the original utility function. We define $\phi : \mathbb{Z}^k \to \mathbb{Z}^k$ as $\phi(x_1, \ldots, x_k) = (x_1 - (n' - n), x_2, \ldots, x_k)$,

$$u'^i_a(x) = \begin{cases} u^i_a(\phi(x)) & \text{if } x_1 \geq n' - n \\ 0 & \text{otherwise} \end{cases}$$

Since $\epsilon = \frac{1}{n'^\gamma}$, by assumption we can compute an ϵ-well-supported equilibrium of G' in time $O(n'^d)$, and we can simply remove the dummy players to obtain an ϵ-equilibrium of the original game G. The running time is $O(n'^d) = \text{poly}(n, 1/\epsilon)$ when $\gamma = \Theta(1)$. $\square$

Proof of Theorem 6.7. Assume that we can compute an $O(n^{-c})$-approximate equilibrium in polynomial time for some constant $c > 1$. Let $\gamma = c-1$, so we can compute an $O\left(\frac{1}{n^{1+\gamma}}\right)$-approximate equilibrium in polynomial time. By Lemma 6.19, we can convert it into an $O\left(\frac{1}{n^{\gamma/2}}\right)$-well-supported equilibrium. Lemma 6.20 then states that any polynomial-time algorithm that computes a well-supported Nash equilibrium of an inverse polynomial precision gives an FPTAS for computing well-supported Nash equilibria in anonymous games. $\square$

Chapter 7

Conclusion and Open Questions

Algorithmic game theory is rife with strategic interactions with uncertainty and information asymmetry. In this thesis, we examined the following question through a computational lens:

> *What is the best way to reveal information to other strategic players, and how hard is it to find the optimal information structure?*

We studied the design of information structures — a principal who is privy to private information must choose how to reveal information to induce a better outcome. We developed algorithms and proved matching hardness results for signaling in many important classes of games: normal form games, and succinct games including network routing games, second price auctions and majority voting.

We saw the role of information revelation changes from chapter to chapter. In informational variants of Braess' paradox and prisoner's dilemma, a principal tries to hide information to help the players fight their selfishness. In normal form games, a principal who wants to help his friend must identify which portion of the information helps one of the players but not the other; which may require her to identify dense subgraphs in a given graph. In second price auctions, a principal who seeks to maximize her revenue must reveal some but not all information to induce the right amount of competition in the market.

The computational complexity of optimal signaling also changes, and becomes easier from chapter to chapter. For network routing games, in the worst case, the principal has to solve NP-hard problems to do better than revealing full information. In normal form games, the principal can compute a near-optimal signaling scheme in quasipolynomial time. As we move to anonymous games like second price auctions and voting, the principal can signal approximately optimally in polynomial time. By settling the computational complexity of these signaling problems, we improved our understanding of information asymmetry in games, as well as the power and limitations of strategic information revelation.

The investigation of optimal information revelation has also led to powerful algorithmic frameworks. Driven by the desire for fundamental insights, we identified the mixture selection problem — an algorithmic problem that arises naturally in the design of optimal information structures. We presented two complexity measures that seem to dictate the complexity of mixture selection and optimal signaling, and solved a number of signaling problems near-optimally under the mixture selection framework.

The design of information structures is emerging as a new area in algorithmic game theory, an area that is still largely unexplored. This thesis addresses the optimal signaling in several basic families of Bayesian games, and there are many exciting problems to be discovered and solved. We list a few open questions below.

Open Questions

Problem 7.1 (Private signaling)**.** *How does the computational complexity change if the principal is allowed to reveal different information to different players?*

In this thesis, we study public signaling schemes, where the principal must reveal the same information to all players. Does private communication make the principal more powerful, and how does the complexity of optimal signaling change? Dughmi and Xu [51] showed that, for multi-player games with n players, the gap between the value of the optimal public and private signaling schemes is at least $\Theta(n)$. They also settle the complexity of public and private signaling when there are no externalities[1]. It remains an interesting open question how to signal efficiently in games with externalities, and whether the interaction between the players makes the signaling problem harder or not.

Problem 7.2 (Equivalence of optimization and separation)**.** *For a polytope P contained in the simplex, if we are given a PTAS for the separation (or membership) oracle of P — an oracle that runs in polynomial time for any constant $\epsilon > 0$ and has ϵ-additive error — can we obtain a PTAS for optimization over P?*

In other words, do we need a much more precise membership oracle to be able to optimize approximately? In Section 4.4, we ruled out an FPTAS for optimal signaling using FPTAS hardness of posterior selection. Recall that the posterior selection problem asks for the best posterior distribution, while the signaling problem asks for the best decomposition (of the prior distribution) into posteriors. It is often easier to show the posterior selection problem is hard, and then use the same intuition to derive a direct reduction for the hardness of signaling. For example, finding a planted clique in a random graph is hard, and for similar reasons finding a constant fraction of a planted clique cover (i.e., decomposing into dense subgraphs) is also hard; approximating the size of the maximum independent set is

[1] In games with no externalities, each player's payoff depends only on his own action (and also on the state of nature for Bayesian games), but not on the actions of other players.

hard; similarly approximating the chromatic number (i.e., decomposing into independent sets) is also hard. These ideas are used implicitly in [15, 29, 46, 51] to show PTAS hardness results for different signaling problems, and these results can be unified if Problem 7.2 can be resolved in the positive.

Problem 7.3 (Nash equilibria in anonymous games)**.** *Is there an FPTAS for computing Nash equilibria in anonymous games?*

Almost all the algorithmic results for equilibrium computation in anonymous games can be viewed as first guessing the outcome of the game, and then trying to reconstruct this outcome using only the best response of each player. New ideas seem to be needed for qualitatively faster algorithms. On the other hand, for ruling out an FPTAS, it is unlikely that the approach in [25] can work directly. This is because $1/\operatorname{poly}(n)$ precision is only enough to de-anonymize $O(\log n)$ players, but $O(\log n)$-player $O(1)$-strategy games can be solved in time $n^{O(\log \log n)}$ (rather than quasipolynomial time) due to the existence theory of the reals.

Problem 7.4 (Routing games with non-linear latencies)**.** *Is there a better signaling scheme than full revelation for Bayesian routing games with non-linear latencies?*

We showed that no polynomial time algorithm can do better than 4/3 in the worst case for signaling in network routing games. The best signaling algorithm we know of, which simply reveals full information, is a multiplicative approximation with the ratio equal to the price of anarchy. It remains open what is the best possible ratio we can obtain in polynomial time for non-linear latency functions.

Problem 7.5 (Planted clique conjecture)**.** *Is there a formal connection between planted clique and widely used worst-case hardness assumptions, e.g., the Exponential Time Hypothesis (ETH)?*

It was shown that computing ϵ-best Nash equilibrium in two-player normal form games requires quasipolynomial time for a small enough constant $\epsilon > 0$, assuming either the planted clique conjecture [63] or the ETH [20]. Two of the hardness results in this thesis, optimal signaling in normal form games, and mixture selection in the absence of noise stability, can both be obtained by assuming either the planted clique conjecture [15] or the ETH [86]. Is there a formal connection between the planted clique conjecture and the ETH?

Reference List

[1] Abbott, T., Kane, D., and Valiant, P. (2005). On the complexity of two-player win-lose games. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 113–122.

[2] Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500.

[3] Alon, N., Andoni, A., Kaufman, T., Matulef, K., Rubinfeld, R., and Xie, N. (2007). Testing k-wise and almost k-wise independence. In *Proceedings of the 38th ACM Symposium on Theory of Computing (STOC)*, pages 496–505.

[4] Alon, N., Krivelevich, M., and Sudakov, B. (1998). Finding a large hidden clique in a random graph. In *Proceedings of the 9th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 594–598.

[5] Alonso, R. and Câmara, O. (2016). Persuading voters. *The American Economic Review*, 106(11):3590–3605.

[6] Althöfer, I. (1994). On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355.

[7] Babichenko, Y. and Barman, S. (2016). Computational aspects of private Bayesian persuasion. *arXiv preprint 1603.01444*.

[8] Babichenko, Y., Barman, S., and Peretz, R. (2014). Simple approximate equilibria in large games. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 753–770.

[9] Babichenko, Y., Papadimitriou, C., and Rubinstein, A. (2015). Can almost everybody be almost happy? PCP for PPAD and the inapproximability of Nash. In *Proceedings of the 7th Conference on Innovations in Theoretical Computer Science*, pages 1–9.

[10] Babichenko, Y. and Rubinstein, A. (2017). Communication complexity of approximate Nash equilibria. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, pages 878–889.

[11] Bárány, I., Vempala, S., and Vetta, A. (2005). Nash equilibria in random games. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 123–131.

[12] Barman, S. (2015). Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory's theorem. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing*, pages 361–369.

[13] Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *The American Economic Review*, 105(3):921–957.

[14] Bergemann, D. and Morris, S. (2017). Information design: A unified perspective. *Cowles Foundation Discussion Paper No. 2075*.

[15] Bhaskar, U., Cheng, Y., Ko, Y. K., and Swamy, C. (2016). Hardness results for signaling in Bayesian zero-sum and network routing games. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 479–496.

[16] Blackwell, D. (1951). Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102.

[17] Blonski, M. (1999). Anonymous games with binary actions. *Games and Economic Behavior*, 28(2):171–180.

[18] Blonski, M. (2005). The women of Cairo: Equilibria in large anonymous games. *Journal of Mathematical Economics*, 41(3):253–264.

[19] Braess, D. (1968). Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12(1):258–268.

[20] Braverman, M., Ko, Y. K., and Weinstein, O. (2015). Approximating the best Nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 970–982.

[21] Bro Miltersen, P. and Sheffet, O. (2012). Send mixed signals: earn more, work less. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 234–247.

[22] Carathéodory, C. (1907). Über den variabilitätsbereich der koeffizienten von potenzreihen, die gegebene werte nicht annehmen. *Mathematische Annalen*, 64(1):95–115.

[23] Chen, X., Cheng, Y., and Tang, B. (2017). Well-supported versus approximate Nash equilibria: Query complexity of large games. In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*.

[24] Chen, X., Deng, X., and Teng, S.-H. (2009). Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):1–57.

[25] Chen, X., Durfee, D., and Orfanou, A. (2015). On the complexity of Nash equilibria in anonymous games. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 381–390.

[26] Cheng, Y., Cheung, H. Y., Dughmi, S., Emamjomeh-Zadeh, E., Han, L., and Teng, S.-H. (2015). Mixture selection, mechanism design, and signaling. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1426–1445.

[27] Cheng, Y., Cheung, H. Y., Dughmi, S., and Teng, S.-H. (2014). Signaling in quasipolynomial time. *arXiv preprint 1410.3033*.

[28] Cheng, Y., Diakonikolas, I., and Stewart, A. (2017). Playing anonymous games using simple strategies. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 616–631.

[29] Cheng, Y. and Ko, Y. K. (2017). Finding best mixture of Nash is hard. *Manuscript*.

[30] Coja-Oghlan, A. (2010). Graph partitioning via adaptive spectral techniques. *Combinatorics, Probability & Computing*, 19(2):227.

[31] Cole, R., Dodis, Y., and Roughgarden, T. (2006). How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72(3):444–467.

[32] Conitzer, V. and Sandholm, T. (2006). Computing the optimal strategy to commit to. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*, pages 82–90.

[33] Daskalakis, C., De, A., Kamath, G., and Tzamos, C. (2016). A size-free CLT for Poisson multinomials and its applications. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, pages 1074–1086.

[34] Daskalakis, C., Diakonikolas, I., and Servedio, R. A. (2012). Learning Poisson binomial distributions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 709–728.

[35] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259.

[36] Daskalakis, C. and Papadimitriou, C. H. (2007). Computing equilibria in anonymous games. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 83–93.

[37] Daskalakis, C. and Papadimitriou, C. H. (2008). Discretized multinomial distributions and Nash equilibria in anonymous games. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 25–34.

[38] Daskalakis, C. and Papadimitriou, C. H. (2009). On oblivious PTAS's for Nash equilibrium. In *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*, pages 75–84.

[39] Daskalakis, C. and Papadimitriou, C. H. (2015). Approximate Nash equilibria in anonymous games. *Journal of Economic Theory*, 156:207–245.

[40] Dekel, Y., Gurel-Gurevich, O., and Peres, Y. (2011). Finding hidden cliques in linear time with high probability. In *Proceedings of the 8th Workshop on Analytic Algorithmics and Combinatorics*, pages 67–75.

[41] Diakonikolas, I., Kane, D. M., and Stewart, A. (2016a). The Fourier transform of Poisson multinomial distributions and its algorithmic applications. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, pages 1060–1073.

[42] Diakonikolas, I., Kane, D. M., and Stewart, A. (2016b). Optimal learning via the Fourier transform for sums of independent integer random variables. In *Proceedings of the 29th Conference on Learning Theory*, pages 831–849.

[43] Diakonikolas, I., Kane, D. M., and Stewart, A. (2016c). Properly learning Poisson binomial distributions in almost polynomial time. In *Proceedings of the 29th Conference on Learning Theory*, pages 850–878.

[44] Doerr, B. (2011). Analyzing randomized search heuristics: tools from probability theory. In *Theory of randomized search heuristics*. World Scientific.

[45] Dubhashi, D. P. and Panconesi, A. (2009). *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press.

[46] Dughmi, S. (2014). On the hardness of signaling. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 354–363.

[47] Dughmi, S. (2017). Algorithmic information structure design: a survey. *ACM SIGecom Exchanges*, 15(2):2–24.

[48] Dughmi, S., Immorlica, N., and Roth, A. (2014). Constrained signaling in auction design. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1341–1357.

[49] Dughmi, S., Kempe, D., and Qiang, R. (2016). Persuasion with limited communication. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 663–680.

[50] Dughmi, S. and Xu, H. (2016). Algorithmic Bayesian persuasion. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, pages 412–425.

[51] Dughmi, S. and Xu, H. (2017). Algorithmic persuasion with no externalities. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 351–368.

[52] Emek, Y., Feldman, M., Gamzu, I., Paes Leme, R., and Tennenholtz, M. (2012). Signaling schemes for revenue maximization. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 514–531.

[53] Etessami, K. and Yannakakis, M. (2010). On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597.

[54] Feder, T., Nazerzadeh, H., and Saberi, A. (2007). Approximating Nash equilibria using small-support strategies. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pages 352–354.

[55] Feige, U. and Krauthgamer, R. (2003). The probable value of the Lovász–Schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2):345–370.

[56] Feige, U. and Ron, D. (2010). Finding hidden cliques in linear time. In *Proceedings of the 21st Conference on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms*, pages 189–204.

[57] Feldman, V., Grigorescu, E., Reyzin, L., Vempala, S., and Xiao, Y. (2013). Statistical algorithms and a lower bound for detecting planted cliques. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC)*, pages 655–664.

[58] Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.

[59] Ghosh, A., Nazerzadeh, H., and Sundararajan, M. (2007). Computing optimal bundles for sponsored search. In *Proceedings of the 3rd Workshop on Internet and Network Economics (WINE)*, pages 576–583.

[60] Goldberg, P. W. and Turchetta, S. (2015). Query complexity of approximate equilibria in anonymous games. In *Proceedings of the 11th Workshop on Internet and Network Economics (WINE)*, pages 357–369.

[61] Grötschel, M., Lovász, L., and Schrijver, L. (1993). Geometric algorithms and combinatorial optimization. *Algorithms and Combinatorics*, 2:1–362.

[62] Guo, M. and Deligkas, A. (2013). Revenue maximization via hiding item attributes. In *Proceedings of the 23rd international joint conference on Artificial Intelligence*, pages 157–163.

[63] Hazan, E. and Krauthgamer, R. (2011). How hard is it to approximate the best Nash equilibrium? *SIAM Journal on Computing*, 40(1):79–91.

[64] Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4):561–574.

[65] Holt, C. and Roth, A. (2004). The Nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences*, 101(12):3999–4002.

[66] Jerrum, M. (1992). Large cliques elude the Metropolis process. *Random Structures & Algorithms*, 3(4):347–359.

[67] Juels, A. and Peinado, M. (2000). Hiding cliques for cryptographic security. *Designs, Codes and Cryptography*, 20(3):269–280.

[68] Kahn, J., Kalai, G., and Linial, N. (1988). The influence of variables on boolean functions. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 68–80.

[69] Kamenica, E. and Gentzkow, M. (2009). Bayesian persuasion. Technical report, National Bureau of Economic Research.

[70] Kannan, R. and Theobald, T. (2007). Games of fixed rank: a hierarchy of bimatrix games. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1124–1132.

[71] Khot, S. and Saket, R. (2012). Hardness of finding independent sets in almost q-colorable graphs. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 380–389.

[72] Kučera, L. (1995). Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2):193–212.

[73] Lehrer, E., Rosenberg, D., and Shmaya, E. (2010). Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682.

[74] Lipton, R. J., Markakis, E., and Mehta, A. (2003). Playing large games using simple strategies. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 36–41.

[75] Mehta, R. (2014). Constant rank bimatrix games are PPAD-hard. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 545–554.

[76] Milchtaich, I. (1996). Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124.

[77] Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5).

[78] Minder, L. and Vilenchik, D. (2009). Small clique detection and approximate Nash equilibria. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 673–685.

[79] Mossel, E., O'Donnell, R., and Oleszkiewicz, K. (2010). Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171(1):295–341.

[80] Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49.

[81] Nemirovski, A. S. and Yudin, D. B. (1983). *Problem complexity and method efficiency in optimization.* Wiley.

[82] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic Game Theory.* Cambridge University Press.

[83] Okuno-Fujiwara, M., Postlewaite, A., and Suzumura, K. (1990). Strategic information revelation. *The Review of Economic Studies*, 57(1):25–47.

[84] Papadimitriou, C. H. (2005). Computing correlated equilibria in multi-player games. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pages 49–56.

[85] Roughgarden, T. and Tardos, E. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2):236–259.

[86] Rubinstein, A. (2015a). ETH-hardness for signaling in symmetric zero-sum games. *arXiv preprint 1510.04991*.

[87] Rubinstein, A. (2015b). Inapproximability of Nash equilibrium. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 409–418.

[88] Rubinstein, A. (2016). Settling the complexity of computing approximate two-player Nash equilibria. In *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 258–265.

[89] Tsaknakis, H. and Spirakis, P. (2007). An optimization approach for approximate Nash equilibria. In *Proceedings of the 3rd Workshop on Internet and Network Economics (WINE)*, pages 42–56.

[90] Vasserman, S., Feldman, M., and Hassidim, A. (2015). Implementing the wisdom of Waze. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 660–666.

[91] Von Stackelberg, H. (1934). *Marktform und gleichgewicht*. Springer.