

Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games

Umang Bhaskar* Yu Cheng[†] Young Kun Ko[‡] Chaitanya Swamy[§]

November 1, 2016

Abstract

We study the optimization problem faced by a perfectly informed principal in a Bayesian game, who reveals information to the players about the state of nature to obtain a desirable equilibrium. This *signaling problem* is the natural design question motivated by uncertainty in games and has attracted much recent attention. We present new hardness results for signaling problems in (a) Bayesian two-player zero-sum games, and (b) Bayesian network routing games.

For *Bayesian zero-sum games*, when the principal seeks to maximize the equilibrium utility of a player, we show that it is *NP*-hard to obtain an additive FPTAS. Our hardness proof exploits *duality* and the equivalence of separation and optimization in a novel way. Further, we rule out an additive PTAS assuming *planted clique* hardness, which states that no polynomial time algorithm can recover a planted clique from an Erdős-Rényi random graph. Complementing these, we obtain a PTAS for a structured class of zero-sum games (where obtaining an FPTAS is still *NP*-hard) when the payoff matrices obey a Lipschitz condition. Previous results ruled out an FPTAS assuming planted-clique hardness, and a PTAS only for *implicit* games with *quasi-polynomial-size strategy sets*.

For *Bayesian network routing games*, wherein the principal seeks to minimize the average latency of the Nash flow, we show that it is *NP*-hard to obtain a (multiplicative) $(\frac{4}{3} - \epsilon)$ -approximation, even for linear latency functions. This is the *optimal* inapproximability result for linear latencies, since we show that full revelation achieves a $\frac{4}{3}$ -approximation for linear latencies.

1 Introduction

In Bayesian games, players' payoffs 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

*Tata Institute of Fundamental Research, Mumbai, India. Part of this work was done while the author was a postdoctoral scholar at the University of Waterloo and was supported in part by Chaitanya Swamy's research grants. Email: umang@tifr.res.in.

[†]University of Southern California, CA 90089, USA. Supported in part by Shang-Hua Teng's Simons Investigator Award from the Simons Foundation. Email: yu.cheng.1@usc.edu.

[‡]Princeton University, NJ 08544, USA. Email: yko@cs.princeton.edu.

[§]Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON N2L 3G1, Canada. Supported in part by NSERC grant 32760-06, an NSERC Discovery Accelerator Supplement Award, and an Ontario Early Researcher Award. Email: cswamy@uwaterloo.ca.

players depend crucially on the information available from the signal they receive. Since applications are often rife with uncertainty, understanding the effect of information available to players is a fundamental problem in game theory; see, e.g., [Bla51, Ake70, Hir71, MW82, LRS10, BBM13]. Whereas for a single player, it is known that more information leads to better payoffs [Bla51], with multiple players, outcomes are more complex and often counterintuitive with “more” (information) not necessarily translating to “better” (payoffs). The latter was first observed by Hirshleifer [Hir71]; recently, Dughmi [Dug14] gave an example where neither full-revelation nor no-revelation is optimal.

While classical work has focused on the role of information in influencing strategies, the computational problem of *designing* optimal information structures for Bayesian games, commonly called the *signaling problem*, has received much recent attention [BMS12, DIR14, EFG⁺12, GD13]. Here, a perfectly-informed principal seeks to reveal selective information to the players to optimize some function of the resulting equilibrium, such as the revenue, or payoff of a particular player. Two-player zero-sum games and network routing games are natural starting points for investigating the signaling problem due to their fundamental importance and appealing structure. They admit a canonical, tractable choice of equilibrium; this also decouples the concerns of optimal-signaling computation and equilibrium computation.

Our results We study signaling in two widely studied classes of games: two-player zero-sum games, and network routing games. As in much of previous work, in our setting players share the same prior belief on the state of nature, and signaling schemes are symmetric: the principal reveals the same information to all players. Further, as previously, our results are for *additive* approximations in Bayesian zero-sum games, and for *multiplicative* approximations in Bayesian network routing games. Our main contribution is to derive hardness results for these classes of games that close the gap between what is achievable in polytime (or quasi-polytime) and what is intractable.

In Section 4, we consider *Bayesian (two-player) zero-sum games*, in which the principal seeks to maximize the value of the game — the equilibrium payoff of the row player.¹ First, we 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 4.1). Previous work by Dughmi [Dug14] ruled out an FPTAS assuming the planted clique hardness (see Conjecture 1). Thus, we replace an *average-case* hardness assumption with the much more conventional worst-case assumption of *NP*-hardness.

Next, we consider the hardness of obtaining a PTAS for the signaling problem. Since there is a quasi-polytime approximation scheme for signaling given by Cheng et al. [CCD⁺15], it is unlikely that a PTAS for signaling is *NP*-hard. We show that assuming planted-clique hardness, there does not exist a PTAS for the signaling problem (Theorem 4.4). Previously, the non-existence of a PTAS was shown (assuming planted-clique hardness) only for *implicit* zero-sum games with *quasi-polynomial-size strategy sets* [Dug14]. Complementing these hardness results, we devise a PTAS for a structured class of Bayesian zero-sum games (Theorem 4.14), when the payoff matrices obey a Lipschitz condition.

In Section 5, 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 5.1). This yields an *optimal* inapproximability

¹In zero-sum games, this also captures the problem of maximizing a weighted combination of players’ equilibrium payoffs.

result for linear latencies, since we show that full revelation obtains the *price of anarchy* of the routing game as its approximation ratio (Theorem 5.4), which is $\frac{4}{3}$ for linear latency functions [RT02]. These are the *first* results for the complexity of signaling in Bayesian network routing games.

We also obtain hardness results for two related signaling problems in Bayesian zero-sum games (Section 6). Firstly, we rule out a PTAS for computing the best prior (the *maximum prior problem*), under the *exponential time hypothesis* (ETH). Previously, [CCD⁺15] studied a mixture-selection problem and showed that in the absence of a property called *noise-stability*, obtaining a PTAS was hard, assuming planted-clique hardness. Our result shows that in their setting a QPTAS is in fact the best possible approximation obtainable, assuming the ETH. Finally, if the principal’s value depends on the players’ strategies, and not just their payoffs, we show that obtaining a PTAS is NP-hard (Theorem 6.1).

Our techniques Our results for Bayesian zero-sum games are obtained via two main ideas. Our NP-hardness result, the PTAS for a structured class of games, and the PTAS-hardness for the maximum prior problem, all follow 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 *dual signaling problem* (Section 3). Our key insight is that *the dual signaling problem is a rather useful tool for both deriving hardness results and devising approximation algorithms*. This usefulness stems from the equivalence of separation and optimization [GLS93], 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 dual signaling problem (Theorem 4.2), and a PTAS for the dual signaling problem yields a PTAS for signaling (Theorem 4.10).

This equivalence paves the way for our results. Whereas, typically, an (approximate) separation oracle is used to (approximately) solve the optimization problem, we exploit this equivalence in an unorthodox fashion by also leveraging the *hardness* of the dual signaling problem to prove hardness results for the signaling (i.e., primal optimization) problem. We show that it is NP-hard to obtain an FPTAS for the dual signaling problem, and thus obtain that it is NP-hard to obtain an FPTAS for the signaling problem. Notably, in contrast to the (weaker) planted-clique hardness result in [Dug14] for the signaling problem, we obtain our NP-hardness result with minimal effort, a fact that underscores the benefits of moving to the dual signaling problem.

On the positive side, we obtain a PTAS for the dual signaling problem for our structured class of Bayesian zero-sum games, which thus yields a PTAS for the signaling problem for this class. Interestingly, when cast in the mixture-selection framework of Cheng et al. [CCD⁺15], the signaling problem for our structured class *does not* satisfy the noise-stability property stated in [CCD⁺15]. In the absence of noise-stability, [CCD⁺15] showed planted-clique hardness for obtaining a PTAS. Our result bypasses this hardness result, and obtains a PTAS for a problem for which noise-stability does not hold. Finally, we show that a PTAS for the maximum-prior problem yields a PTAS for the dual signaling problem, and we rule out the latter via a simple, clean reduction from the *best-Nash* problem and the recent result of Braverman et al. [BKW15]. This result also strengthens the hardness result from [CCD⁺15] mentioned above, by showing that in the absence of noise-stability, a QPTAS is the best-possible approximation for the mixture selection problem, assuming the ETH.

Our second main idea, used to rule out a PTAS assuming planted-clique hardness, is a “direct” reduction that combines and strengthens techniques from [Dug14, FNS07]. We utilize the *planted clique cover* problem defined in [Dug14] — multiple cliques are now planted and one seeks to recover a constant fraction of them — and shown to be at least as hard as the planted clique problem. 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 (and hence, the planted clique) problem. The FPTAS-hardness reduction in [Dug14] creates a network security game (see Section 2) with payoffs of absolute value $\Omega(\log^2 n)$ (or alternatively, a quasi-polynomial-size strategy set for the column player) to enforce the above property. Payoffs of magnitude $\Omega(\log n)$ seem necessary with this kind of approach, which therefore only yields an $O(1/\log n)$ gap that is insufficient to rule out a PTAS. We abandon the use of network security games and instead leverage a device from [FNS07] to ensure the above “large-spreading” property. This idea is also used to show planted-clique hardness for the *best-Nash* problem [HK11]; however we are constrained to work with zero-sum games, and therefore need to apply this idea carefully. A subtle, but crucial, technical issue is that we need to significantly tighten the planted-clique recovery result in [Dug14]. To recover a specific planted clique S of size $k = \omega(\log^2 n)$ (in the presence of other such planted cliques), [Dug14] requires a set T with $|T| = \Theta(k)$, $|S \cap T| = \Omega(k)$, whereas we only require that $|T|, |S \cap T| = \Omega(\log n)$, and this is crucial since we can only ensure that spreading takes place over $O(\log n)$ -size sets.

Our hardness result for Bayesian routing games is a direct reduction from the problem of computing edge tolls that minimize the total (latency + toll)-cost of the resulting equilibrium flow, which is inapproximable within a factor of $\frac{4}{3}$.

Related work Whereas understanding the role of information in influencing strategies is a classical problem in game theory, the computational problem of designing optimal information structures has been studied more recently. Much of this work has focused on signaling in auctions, where the goal is to maximize revenue [EFG⁺12, BMS12, GD13] or social welfare [DIR14]. Dughmi [Dug14] initiated the computational study of signaling in Bayesian zero-sum games, and obtained various hardness results under the planted-clique hardness assumption. This work left open the question of whether hardness results can be obtained under standard worst-case assumptions, such as $P \neq NP$, a question that we answer in the affirmative. On the positive side, Cheng et al. [CCD⁺15] showed that for Bayesian normal form games with a constant number of players and for general objectives of the principal, an ϵ -approximate signaling scheme that maximizes the objective at an ϵ -approximate Nash equilibrium can be computed in quasi-polynomial time. This work left open the question of whether a PTAS is possible for signaling in Bayesian zero-sum games. We preclude this under the planted-clique hardness assumption, and complementing this, design a PTAS for a structured class of games. As noted earlier, the latter result does not follow from [CCD⁺15] since the resulting signaling problem fails to have small noise stability.

The planted-clique problem was introduced by Jerrum [Jer92] and Kučera [Kuč95], and despite extensive efforts (see, e.g., [AV14, FR10, DGGP11] and the references therein), no polytime algorithm is known for recovering cliques of size $k = o(\sqrt{n})$. There is a quasi-polytime algorithm known when $k \geq 2 \log_2 n$; on the other hand, various algorithmic strategies have been ruled out for this problem [Jer92, FK03, FGR⁺13]. The planted-clique problem has thus been used in various reductions (see, e.g., [HK11, JP00]), and is an example where an *average-case hardness assumption*

has been used to derive hardness results.

Recently, Rubinstein [Rub15] has independently also obtained hardness results for signaling in zero-sum games. He shows that there is no additive PTAS assuming ETH, and obtaining a multiplicative PTAS is *NP*-hard. These results are orthogonal to ours, as there is no known reduction between ETH and planted-clique hardness. Further, *NP*-hardness of a multiplicative PTAS does not rule out an additive FPTAS.

In Bayesian network routing games, [VFH15] study the ability 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 with graphs consisting of parallel links. On these simple networks, navigation services (such as Waze or Google Maps) cannot do anything to improve the latency of the Nash flow. Our work, in contrast, studies the *computational complexity* of obtaining the best (public) signaling scheme in general graphs, and conclude that finding an $(\frac{4}{3} - \epsilon)$ -approximation is *NP*-hard.

2 Preliminaries and notation

We use \mathbb{R}_+ for the set of nonnegative reals. For integer n , $[n] := \{1, 2, \dots, n\}$. If $n \geq 1$, we use Δ_n to denote the $(n - 1)$ -dimensional simplex $\{x \in \mathbb{R}_+^n : \sum_i x_i = 1\}$. 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 vector containing 1 as its i -th entry and 0 elsewhere.

Bayesian zero-sum games and signaling schemes A *Bayesian zero-sum game* is specified by a tuple $(\Theta, \{\mathcal{A}^\theta\}_{\theta \in \Theta}, \lambda)$, where $\Theta = \{1, \dots, M\}$ denotes the states of nature, and λ is a prior distribution on the states of nature (thus $\lambda \in \Delta_M$). We assume the row and column player has r, c pure strategies respectively. For each state of nature $\theta \in \Theta$, $\mathcal{A}^\theta \in [-1, 1]^{r \times c}$ specifies the payoffs of the row player in a zero-sum game. Let $\mu \in \Delta_M$ be an arbitrary distribution over states of nature. Then $\mathcal{A}^\mu := \sum_{\theta \in \Theta} \mu_\theta \mathcal{A}^\theta$ is the matrix of expected payoffs for the row player under distribution μ .

A *signaling scheme* is a policy by which a principal reveals (partial) information about the state of nature. We focus on symmetric signaling schemes which reveals the same information to all the players. A signaling scheme specifies a set of signals Σ and a map $\varphi : \Theta \mapsto \Delta_{|\Sigma|}$ from the states of nature Θ to distributions over the signals in Σ . Thus, $\varphi(\theta)_\sigma$ is the probability that the principal selects signal σ when the state of nature is θ . When the state θ is revealed, the principal computes a signal $\sigma \sim \varphi(\theta)$. Both players receive σ and correspondingly update their belief on the state-distribution to μ^σ , where for each state θ ,

$$\mu_\theta^\sigma = \frac{\Pr(\sigma|\theta) \Pr(\theta)}{\Pr(\sigma)} = \frac{\varphi(\theta)_\sigma \lambda_\theta}{\sum_{\theta' \in \Theta} \varphi(\theta')_\sigma}.$$

The players then, based on their posterior belief, play the zero-sum game given by \mathcal{A}^{μ^σ} .

Each signal σ thus yields a posterior distribution $\mu^\sigma \in \Delta_M$, and these posterior distributions form a convex decomposition of the prior $\lambda = \sum_\sigma \Pr(\sigma) \mu^\sigma$. As observed in [Dug14], specifying a signaling scheme (Σ, φ) is in fact equivalent to specifying a distribution α over posterior distributions $\mu \in \Delta_M$ that yield a convex decomposition of the prior λ . Thus, a signaling scheme can also be described as $\alpha := (\alpha_\mu)_{\mu \in \Delta_M}$, where $\sum_{\mu \in \Delta_M} \alpha_\mu \mu = \lambda$. The signals Σ in such a signaling scheme are described implicitly, and correspond to the posteriors μ for which $\alpha_\mu > 0$. This will be our

perspective on signaling schemes throughout. In Section 4.2, we will explicitly need to describe the signals, and then use μ^σ for the posterior corresponding to signal σ and α_σ for $\Pr(\sigma)$.

Let $\text{val} : \Delta_M \mapsto \mathbb{R}$ be the principal's objective function. For the bulk of our results, we consider the objective function $\text{val}(\mu)$ for $\mu \in \Delta_M$ that evaluates to the row-player's payoff at equilibrium in the zero-sum game specified by \mathcal{A}^μ . Note that $\text{val}(\mu)$ is unique, $\text{val}(\mu) := \max_{x \in \Delta_r} \min_{j \in [c]} (x^T \mathcal{A}^\mu)_j$, although there could be multiple Nash equilibria.

The quality of a signaling scheme α for a Bayesian zero-sum game is then given by $\sum_{\mu \in \Delta_M} \alpha_\mu \text{val}(\mu)$. The *signaling problem* in a Bayesian zero-sum game is to find a signaling scheme α that maximizes $\sum_{\mu \in \Delta_M} \alpha_\mu \text{val}(\mu)$. Let $\text{opt}(\mathcal{I})$ denote the value of the optimal signaling scheme for a Bayesian zero-sum game \mathcal{I} . We note that $\text{opt}(\mathcal{I})$ 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 . By Caratheodory's theorem, $M + 1$ posteriors (equivalently, signals) suffice to specify any convex decomposition of the prior. Together, this implies that an optimal signaling scheme can be specified by at most $M + 1$ posteriors.

We say that an algorithm for the signaling problem is an (additive) ε -*approximation algorithm* if for every instance \mathcal{I} the algorithm runs in polytime and returns a signaling scheme of value at least $\text{opt}(\mathcal{I}) - \varepsilon$. A *polytime approximation scheme* (PTAS) is an algorithm that runs in polytime and returns a solution of value at least $\text{opt}(\mathcal{I}) - \varepsilon$ for every instance \mathcal{I} and constant $\varepsilon > 0$; an FPTAS is a PTAS whose running time for an instance \mathcal{I} and parameter ε is $\text{poly}(\text{size of } \mathcal{I}, \frac{1}{\varepsilon})$.

Security games Some of our results utilize a class of zero-sum games that we call *extended security games*, wherein the payoff matrix for state θ is given by

$$\mathcal{A}^\theta := \bar{A} + b^\theta \mathbb{1}_c^T + \mathbb{1}_r (d^\theta)^T, \quad \text{where} \quad b^\theta \in \mathbb{R}^r, d^\theta \in \mathbb{R}^c. \quad (1)$$

Let B and D be matrices having columns $\{b^1, \dots, b^M\}$, and $\{d^1, \dots, d^M\}$ respectively. We obtain the following expressions for \mathcal{A}^μ and $\text{val}(\mu)$ for $\mu \in \Delta_M$.

$$\mathcal{A}^\mu = \bar{A} + (B\mu) \mathbb{1}_c^T + \mathbb{1}_r (\mu^T D^T), \quad \text{val}(\mu) = \max_{x \in \Delta_r} \left\{ x^T B\mu + \min_{j \in [c]} (x^T \bar{A} + \mu^T D^T)_j \right\}. \quad (2)$$

A special case of an extended security game (and the reason for this terminology) is the *network security game* defined by [Dug14]. 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 $\bar{A} = D^T = -\rho I_{n \times n}$. 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 $d = \theta = a$.

Planted clique and planted clique cover Some of our hardness results are based on the hardness of the *planted-clique* and *planted clique cover* problems. The latter problem was introduced by Dughmi [Dug14].

Definition 2.1 (Planted clique cover problem **PCover**(n, p, k, r) [Dug14]). 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 $i = 1, \dots, r$, 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, \dots, S_r , given $G \sim \mathcal{G}(n, p, k, r)$.

In the *planted clique problem* $\mathbf{PClique}(n, p, k)$, there is a single planted clique ($r = 1$) and the goal is to recover this clique. The following hardness assumption for the planted-clique problem has been used in deriving various hardness results.

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

The ellipsoid method. We utilize the ellipsoid method to translate hardness and approximation results for the dual of the signaling problem to signaling.

Theorem 2.2 (Chapters 4, 6 in [GLS93]; Section 9.2 in [NY83]). *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 $\text{poly}(\text{size of } y, L)$ 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 $\text{poly}(n, L)$.*
- (ii) *Let $h : \mathbb{R}^n \mapsto \mathbb{R}$ 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 $T = \text{poly}(n, L, \log(\frac{K}{\epsilon}))$ and using at most T queries to the value oracle for h .*

3 The dual signaling problem

The signaling problem can be formulated as the following mathematical program,

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

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).

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

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

Definition 3.1 (Dual signaling with precision parameter ϵ). Given a Bayesian zero-sum game $(\Theta, \{\mathcal{A}^\theta\}_{\theta \in \Theta}, \lambda)$, $w \in \mathbb{R}^M$, and $\epsilon > 0$, distinguish between:

- (i) $\text{val}(\mu) \geq w^T \mu + \epsilon$ for some $\mu \in \Delta_M$; if so return $\mu \in \Delta_M$ s.t. $\text{val}(\mu) \geq w^T \mu - \epsilon$;
- (ii) $\text{val}(\mu) < w^T \mu - \epsilon$ for all $\mu \in \Delta_M$.

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

Notice that the dual signaling problem is unconstrained: λ plays no role.

4 Bayesian zero-sum games

We now prove the following results for signaling in Bayesian zero-sum games. We show that the signaling problem does not admit an FPTAS unless $P=NP$ (Theorem 4.1) and does not admit a PTAS assuming the hardness of the planted-clique problem (Theorem 4.4). Complementing these hardness results, we present a PTAS for a structured class of extended security games (Theorem 4.14).

4.1 NP-hardness of obtaining an FPTAS

Theorem 4.1 (Corollary of Theorems 4.2 and 4.3). *There is no FPTAS for the signaling problem, even for network security games, unless $P=NP$.*

Theorem 4.2. *There is a polynomial $q(M)$ such that an $\frac{\varepsilon}{q(M)}$ -approximation algorithm \mathcal{B} for the signaling problem \mathcal{I} yields a polytime algorithm for the threshold signaling problem $(\mathcal{I}, \eta \mathbb{1}_M, \varepsilon)$. Thus, an FPTAS for the signaling problem yields an FPTAS for the threshold signaling problem.*

Proof. Let $(\Theta, \{\mathcal{A}^\theta\}, \lambda), \eta \mathbb{1}_M, \varepsilon$ be the input to the threshold signaling problem, with precision parameter ε . Note that for any $\mu \in \Delta_M$, we have $-1 \leq \text{opt}(\mu) \leq 1$ since $|\mathcal{A}_{i,j}^\theta| \leq 1$ for all θ, i, j . Let $p(M)$ be the polynomial given by part (ii) of Theorem 2.2. Set $q(M) = p(M) + 1$. We utilize part (ii) of Theorem 2.2 with $X = \Delta_M$, $\delta = \frac{\varepsilon}{q(M)}$, $h(\cdot) = \text{opt}(\cdot)$ (which is concave, as noted earlier), $K = 2$, and using \mathcal{B} as the imperfect value oracle, to find $z \in \Delta_M$ in polytime such that $\text{opt}(z) \geq \max_{\mu \in \Delta_M} \text{opt}(\mu) - p(M)\delta$. We run \mathcal{B} on the prior z to obtain a signaling scheme α of value $v \geq \text{opt}(z) - \varepsilon$. If $v \geq \eta$, then we return that we are in case (i) and one of the points $\mu \in \Delta_M$ with $\alpha_\mu > 0$ must satisfy $\text{val}(\mu) \geq \eta$. If $v < \eta$, then we have $\max_{\mu \in \Delta_M} \text{val}(\mu) \leq \max_{\mu \in \Delta_M} \text{opt}(\mu) < \eta + (p(M) + 1)\delta$, so we return that we are in case (ii). ■

Theorem 4.3. *There is no FPTAS for the threshold signaling problem, even for network security games, unless $P=NP$.*

Proof. The proof follows readily via a reduction from the *balanced complete bipartite subgraph* (BCBS) problem [GJ79], which illustrates the convenience of working with the dual signaling problem. In BCBS, given a bipartite graph $G = (V \cup W, E)$ and an integer $r \geq 0$, we want to determine if G contains $K_{r,r}$ (i.e., an $r \times r$ biclique). Given a BCBS instance, set $\varepsilon = \frac{1}{2n^8}$, where $n = |V| + |W|$, and $\eta = 1 - (2n + 1)\varepsilon$. We create a Bayesian network security game by letting G be the graph in the network security game, and setting $\rho = 2rn\varepsilon$. Recall that this means that states of nature correspond to nodes of G , so $\Theta = V \cup W$, and the payoff matrix for a distribution $\mu \in \Delta_\Theta$ is given by (2) where B is the adjacency matrix of G and $\bar{A} = D^T = -\rho I_{n \times n}$. This creates an instance of the threshold signaling problem with precision parameter ε ; the prior λ is irrelevant. We show that solving this instance would decide the BCBS-instance.

If G has the required subgraph V', W' , set $\mu_v = 1/r$ for all $v \in V'$ and $x_v = 1/r$ for all $v \in W'$. Then, by (2), we have $\text{val}(\mu) \geq x^T B \mu - \rho \|\mu + x\|_\infty \geq 1 - \rho/r = \eta + \varepsilon$. where we have $x^T B \mu = 1$ since V', W' form a complete bipartite subgraph.

Suppose there exists $\mu \in \Delta_M$ so that $\text{val}(\mu) \geq \eta - \varepsilon$. We show then that G contains $K_{r,r}$. Let x be the equilibrium strategy of the attacker, so $\text{val}(\mu) = x^T B \mu - \rho \|\mu + x\|_\infty$. Let $V' := \{v \in V \cup W : \mu_v \geq 1/n^3\}$ and $W' := \{v \in V \cup W : x_v \geq 1/n^3\}$. Then $\sum_{v \in V'} \mu_v = 1 - \sum_{v \notin V'} \mu_v > 1 - 1/n^2$. Similarly $\sum_{v \in W'} x_v > 1 - 1/n^2$. Every vertex in V' must be adjacent to every vertex in W' , otherwise $x^T B \mu \leq 1 - 1/n^6 < \eta$. Thus, V' and W' must be in different partitions. Assume $V' \subseteq V$

and $W' \subseteq W$. For each vertex v , $\mu_v + x_v \leq \frac{(1+1/n)}{r}$, otherwise $\text{val}(\mu) < 1 - (2n + 2)\varepsilon$. Hence, $|V'| \geq \frac{\sum_{v \in V'} \mu_v}{(1+1/n)/r} > r \frac{1-1/n^2}{(1+1/n)} = r(1 - 1/n)$, and therefore $|V'| \geq r$. Similarly $|W'| \geq r$, and this yields the $r \times r$ biclique. \blacksquare

We conjecture that Theorem 4.2 can, in fact, be strengthened to show that an ε -approximation for signaling yields an $O(\varepsilon)$ -approximation for threshold signaling, so that a PTAS for signaling yields a PTAS for threshold signaling. This would *rule out a sub-quasipolytime approximation scheme* (i.e., an $n^{\tilde{\Omega}(\log^{1-o(1)} n)}$ -time approximation scheme) for signaling under the (deterministic) *exponential time hypothesis* (ETH), since we prove in Section 6 that there is no sub-quasi-PTAS for threshold signaling assuming ETH.

This would be an *optimal* hardness result since a quasi-PTAS follows from [CCD⁺15]. Recently, Rubinstein [Rub15] obtains this hardness result via a *direct* reduction that builds upon ideas in [AIM14]. However, tightening part (ii) of Theorem 2.2 would give a much simpler proof. We leave this as an intriguing open question. Below, we rule out a PTAS for signaling under an orthogonal hardness assumption.

4.2 Planted-clique hardness of obtaining a PTAS

Theorem 4.4. *There is a constant ε_0 such that, assuming the planted-clique hardness conjecture (Conjecture 1), there is no ε_0 -approximation for the signaling problem in Bayesian zero-sum games.*

Our hardness result strengthens the one in [Dug14], which rules out an FPTAS assuming the planted-clique conjecture. The reduction therein creates a network security game from a graph $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ (see Section 2). The idea is that if a signaling scheme achieves value close to 1, then it must place a large weight on posteriors and attacker mixed-strategies that randomize over a large set of nodes. Further, the posterior and attacker must essentially identify dense components of G , as otherwise the attacker's value would be close to the background density $\frac{1}{2}$. As noted earlier, a limitation of this type of construction is that the parameter ρ used in the network security game needs to be roughly $\Omega(\log n)$ to ensure that the posterior and the attacker's mixed strategies are supported on an $\Omega(\log n)$ -size set of nodes. This only yields an $\Theta(\frac{1}{\text{polylog}(n)})$ gap, which is insufficient to rule out a PTAS. We overcome this obstacle by moving away from a network security game, and instead exploiting an idea of [FNS07] to eliminate all equilibria of $O(\log n)$ -size support from the game. Theorem 4.4 follows immediately by combining Lemmas 4.5 and 4.6.

Lemma 4.5. *Let $\varepsilon > 0$, $k = k(n)$ satisfy $k = \omega(\log n)$ and $k = o(\sqrt{n})$, and $r = \Theta(n/k)$. Suppose there is a polytime algorithm that takes as input $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ with planted cliques $\{S_i\}$, and outputs a family $\mathcal{T} \subseteq 2^V$ of clusters satisfying the following with constant probability, for any constant $c_3 \geq 10^3$:*

$$\text{for an } \varepsilon\text{-fraction of } \{S_i\}, \exists T \in \mathcal{T} \text{ with } |T \cap S_i| \geq \max\{\varepsilon|T|, c_3 \log n\}. \quad (*)$$

Then there is a polynomial-time algorithm for $\mathbf{PClique}(n, \frac{1}{2}, k)$ having constant success probability.

Lemma 4.6. *Let $k = k(n)$ satisfy $k = \omega(\log n)$ and $k = o(\sqrt{n})$, and $r = \frac{5n}{k}$. There is a polynomial-time randomized reduction that takes a graph $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ as input and outputs a Bayesian zero-sum game such that the following hold with high probability.*

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

(Soundness) Given a signaling scheme of value at least 0.97, one can obtain in polytime a collection \mathcal{T} of clusters satisfying condition (*) in Lemma 4.5.

Above, and throughout this section, when we say with high probability, we mean success probability $1 - \frac{1}{\text{poly}(n)}$. The Bayesian zero-sum game we construct always admits a signaling scheme of large value; however *finding* a near-optimal signaling scheme in polytime would refute the planted-clique conjecture. Lemma 4.5 (proved in Appendix A) is similar to a planted-clique recovery result proved in [Dug14]. While we utilize similar ideas, our result works under *much weaker* requirements. Our lemma allows clusters in \mathcal{T} to have size $\Theta(\log n)$ — which is crucial for Lemma 4.6 — whereas in [Dug14], the clusters need to have size $\omega(\log^2 n)$. In the rest of this section, we prove Lemma 4.6. We use the following parameters.

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

To keep the presentation simple, we give a construction where $\mathcal{A}_{i,j}^\theta \in [-Z, Z]$ (as opposed to $[-1, 1]$). Let A_G denote the $(n \times n)$ adjacency matrix of $G = (V, E)$. We split G into G^- and G^+ with corresponding adjacency matrices A_G^- and A_G^+ where G^- are the background edges and G^+ are the clique edges added in steps (1) and (2) of Definition 2.1 respectively. The states of nature and the row-player's strategies correspond to the nodes of G . The prior λ is $\mathbb{1}_n/n$, thus each state of nature (each vertex) is equally likely to occur. For every $\theta \in \Theta = V$, the payoff matrix $\mathcal{A}^\theta \in [-Z, Z]^{n \times (2N+1)}$ is given by $[a^\theta \ B \ \mathbb{1}_n(d^\theta)^T]$, which are defined as follows:

- (1) a^θ is the θ -th column of the adjacency matrix A_G , so $a_i^\theta = 1$ 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) $d^\theta \in [-Z, Z]^N$, where each entry d_j^θ is set independently to $2 - Z$ with probability $\frac{3}{4Z}$, and 2 otherwise.

We use **Row** and **Col** to denote the row and column players respectively. Let D be the $n \times N$ matrix having rows $(d^\theta)^T$ for $\theta \in \Theta$.

To gain some intuition, observe that for a posterior μ and **Row**'s mixed strategy x , the row vector $x^T \mathcal{A}^\mu$ yielding **Col**'s payoffs is $[x^T A_G \mu \ x^T B \ \mu^T D]$. 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

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

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 or D is $\frac{5}{4} > \max_i a_i^\theta$. Exploiting this, Claim 4.7(i) implies that if x and μ both randomize uniformly over a large set of vertices, **Col** plays column 1. The completeness proof now follows from the oft-used idea of (roughly speaking) choosing posteriors and mixed strategies for **Row** that randomize uniformly over the planted cliques. Conversely, if x or μ has support of size at most $c_2 \log n$, then Claim 4.7(ii) implies that **Col** can play some column of B or D and make $\text{val}(\mu)$ 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 satisfying condition (*) of Lemma 4.5. This yields the soundness proof.

The following properties about the above construction will be useful.

Claim 4.7. Let $R \subseteq V$. (i) If $|R| = \omega(\log n)$, with high probability, for every $j \in [N]$, $\frac{1}{|R|} \sum_{i \in R} B_{i,j} > 1$ and $\frac{1}{|R|} \sum_{i \in R} D_{i,j} > 1$. (ii) If $|R| \leq c_2 \log n$, with high probability, $\exists j, k \in [N]$ such that $B_{i,j} = 2 - Z = D_{i,k}$ for all $i \in R$.

Proof. We first prove (i). The proof is a standard application of Chernoff bounds, and is also essentially shown in [HK11]. We prove the statement for B ; the argument for D is identical. Fix a column $j \in [N]$. We have $\mathbf{E}\left[\frac{\sum_{i \in R} B_{i,j}}{|R|}\right] = \frac{5}{4}$, 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 $\Pr\left[\frac{\sum_{i \in R} B_{i,j}}{|R|} < \frac{9}{8}\right] \leq \frac{1}{2N \text{poly}(n)}$. The union bound over all N columns yields the claim.

We now prove (ii). The proof again follows from Chernoff bounds, and is the key insight in [FNS07] (also utilized in [HK11]). Fix some $R \subseteq V$ with $|R| = c_2 \log n$. We prove the statement for B ; the proof for D is identical. 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

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

Taking the union bound over all $R \subseteq V$ with $|R| = c_2 \log n$, we obtain

$$\begin{aligned} & \Pr\left[\exists R \subseteq V \text{ with } |R| = c_2 \log n \text{ s.t. no } j \in [N] \text{ satisfies } B_{i,j} = 2 - Z \text{ for all } i \in R\right] \\ & \leq \binom{n}{c_2 \log n} \left[1 - \left(\frac{3}{4Z}\right)^{|R|}\right]^N \leq \exp\left(c_2 \log^2 n - N \left(\frac{3}{4Z}\right)^{c_2 \log n}\right) \\ & \leq 1 - 1/\text{poly}(n). \quad \blacksquare \end{aligned}$$

Lemma 4.8 (Proposition B.2 in [Dug14] quantified). Let $\varepsilon > 0$, and $c \geq 24 \cdot 2.1 \cdot \max\{1, \frac{1+\varepsilon}{\varepsilon^2}\}$. For all $n \geq 2$, we have

$$\Pr[\exists S, T \subseteq V \text{ with } |S|, |T| \geq c \log n, \text{ bi-density}_{G^-}(S, T) > \frac{1+\varepsilon}{2}] \leq \frac{2}{n^3}$$

Lemma 4.9 (Corollary of Lemma 4.8). For $c_2 = 10^5$ and $\varepsilon = 0.03$ With high probability, for all $S, T \subseteq V$ with $|S|, |T| \geq c_2 \log n$, $\text{bi-density}_{G^-}(S, T) \leq \frac{1+\varepsilon}{2}$.

4.2.1 Completeness proof in Lemma 4.6

We use a deterministic signaling scheme that groups together states of nature in the same planted clique. Let S_1, \dots, S_r be the planted cliques in G in some arbitrary order. 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 $A := V \setminus \bigcup_j S_j$ as the remaining vertices. Finally, $S'_0 = A \cup \{v \in S'_i : |S'_i| < \frac{k}{10^4}\}$. Our signaling scheme is (Σ, α, μ) where the set of signals is $\Sigma = \{0\} \cup \{i \in [r] : |S'_i| \geq \frac{k}{10^4}\}$. For each signal σ , $\alpha_\sigma = \frac{|S'_\sigma|}{n}$ and μ_σ is the uniform distribution over S'_σ . Note that the signaling scheme is independent of B and D .

For posterior μ^σ , where $\sigma \neq 0$, consider the strategy x^σ where **Row** plays the uniform distribution on S'_σ . Claim 4.7(i) implies that **Col**'s best response to x^σ is to play column 1. Therefore, $\text{val}(\mu^\sigma) \geq \text{bi-density}(S'_\sigma, S'_\sigma) = 1 - \frac{1}{|S'_\sigma|} \geq 1 - \frac{10^4}{k}$. With $r = \frac{5n}{k}$, we have $|A| \leq e^{0.1} \cdot \mathbf{E}[|A|] \leq e^{-4.9} n$ with high probability due to standard Chernoff bounds (since the events $\{v \in A\}_{v \in V}$ are negatively correlated). Therefore, for suitably large n , with high probability, $|S'_0| \leq |A| + \frac{5n}{k} \cdot \frac{k}{10^4} \leq e^{-4.7} n$. So, with high probability, the signaling scheme has value at least $\sum_{\sigma \in \Sigma \cap [r]} \alpha_\sigma \left(1 - \frac{10^4}{k}\right) \geq (1 - e^{-4.7}) \left(1 - \frac{10^4}{k}\right) \geq 0.99$.

4.2.2 Soundness proof in Lemma 4.6

For a signal $\sigma \in \Sigma$ with corresponding posterior μ_σ , let x_σ denote **Row**'s equilibrium strategy for \mathcal{A}^{μ_σ} . We first filter out the set of ‘‘useful’’ signals, i.e., those with relatively high value. Let $\Sigma_1 = \{\sigma \in \Sigma : \text{val}(\mu_\sigma) \geq 1 - \sqrt{\epsilon}\}$. 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. Fix $\epsilon = 0.03$. For every signal $\sigma \in \Sigma_1$, define $T_\sigma = \left\{i : e_i^T A_G \mu_\sigma \geq 1 - \frac{Z\sqrt{\epsilon}}{Z-2}\right\}$, and let \tilde{x}_σ be the uniform distribution on T_σ . We output $\mathcal{T} = \{T_\sigma : \sigma \in \Sigma_1\}$.

We show that \mathcal{T} satisfies condition (*) in Lemma 4.5. The value of the signaling scheme is $\sum_{\sigma \in \Sigma} \alpha_\sigma \text{val}(\mu_\sigma) \geq 1 - \epsilon$. Noting that $\text{val}(\mu) \leq 1$ for all μ , by Markov's inequality, we have $\alpha(\Sigma_1) \geq 1 - \sqrt{\epsilon}$. (Given a vector $v \in \mathbb{R}^k$, and $S \subseteq [k]$, we use $v(S)$ to denote $\sum_{i \in S} v_i$.) Assume that the high probability event in Claim 4.7(ii) happens.

Fix $\sigma \in \Sigma_1$. For any $R \subseteq V$ with $|R| \leq c_2 \log n$, we must have $x_\sigma(R) \leq \frac{2}{Z}$ and $\mu_\sigma(R) \leq \frac{2}{Z}$. Otherwise, suppose $x_\sigma(R) > \frac{2}{Z}$ (the argument for μ_σ is similar). Then, considering the column j of B having $B_{i,j} = 2 - Z$ for all $i \in R$, we have $\sum_{i \in [n]} (x_\sigma)_i B_{i,j} \leq (2 - Z)x_\sigma(R) + 2(1 - x_\sigma(R)) < 0$, which implies that $\text{val}(\mu_\sigma) < 0$. Now since $1 - \sqrt{\epsilon} \leq \text{val}(\mu_\sigma) \leq 1$, by the definition of T_σ and Markov's inequality, we have $x_\sigma(T_\sigma) \geq \frac{2}{Z}$, and hence $|T_\sigma| \geq c_2 \log n$. We now switch from x_σ to \tilde{x}_σ in order to relate the value of the signaling scheme to bi-density and deduce that \mathcal{T} satisfies condition (*). As before, G^- are the background edges and G^+ are the clique edges added in steps (1) and (2) of Definition 2.1 respectively, 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 $A_G \leq A_G^- + A_G^+ \leq A_G^- + \sum_{i=1}^r A_G^i$.

Let R denote the $c_2 \log n$ largest entries in $\tilde{x}_\sigma^T A_G^-$, and let $\tilde{\mu}_\sigma$ be the uniform distribution on R . Since $\tilde{\mu}_\sigma$ and \tilde{x}_σ are uniform distributions over R and T_σ respectively (which have size at least $c_2 \log n$), we have $\tilde{x}_\sigma^T A_G^- \tilde{\mu}_\sigma = \text{bi-density}(T_\sigma, R) \leq \frac{1+\epsilon}{2}$ due to Lemma 4.9. Moreover, $\mu_\sigma(R) \leq \frac{1}{10}$, and since the maximum entry of $\tilde{x}_\sigma^T A_G^-$ outside of R is at most the average entry in R , we have $\tilde{x}_\sigma^T A_G^- \mu_\sigma \leq \frac{1}{10} + \frac{9}{10} \cdot \tilde{x}_\sigma^T A_G^- \tilde{\mu}_\sigma < 0.6$.

Finally, we also have $\sum_{\sigma \in \Sigma_1} \alpha_\sigma (\tilde{x}_\sigma^T A_G \mu_\sigma) \geq (1 - \sqrt{\epsilon})(1 - \frac{Z\sqrt{\epsilon}}{Z-2}) > 0.85$. Therefore,

$$\begin{aligned} \frac{1}{4} &< \sum_{\sigma \in \Sigma_1} \alpha_\sigma \tilde{x}_\sigma^T (A_G - A_G^-) \mu_\sigma \leq \sum_{\sigma \in \Sigma_1} \alpha_\sigma \sum_{i=1}^r \tilde{x}_\sigma^T A_G^i \mu_\sigma \\ &= \sum_{i=1}^r \sum_{\sigma \in \Sigma_1} \alpha_\sigma \mu_\sigma(S_i) \frac{|T_\sigma \cap S_i|}{|T_\sigma|} \leq \sum_{i=1}^r \left(\sum_{\sigma \in \Sigma_1} \alpha_\sigma \mu_\sigma(S_i) \right) \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T|} \right) \\ &\stackrel{(**)}{\leq} \sum_{i=1}^r \frac{|S_i|}{n} \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T|} \right) = \frac{5}{r} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T|} \right). \end{aligned}$$

Inequality (**) follows since for every $v \in \Theta$, we have $\sum_{\sigma \in \Sigma_1} \alpha_\sigma (\mu_\sigma)_v$ is at most $\sum_{\sigma \in \Sigma} \alpha_\sigma (\mu_\sigma)_v = \lambda_v = \frac{1}{n}$. Therefore $\frac{1}{r} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T|} \right) \geq \frac{1}{20}$. This implies that at least a $\frac{1}{39}$ -fraction of S_1, \dots, S_r satisfy $\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T|} \geq \frac{1}{40}$. Since $|T| \geq c_2 \log n$ for all $T \in \mathcal{T}$, \mathcal{T} satisfies condition (*) in Lemma 4.5.

4.3 A PTAS for structured extended security games

We now devise a PTAS for a structured class of extended security games (Theorem 4.14). First, we reduce the signaling problem to the dual signaling problem using the ellipsoid method (Theo-

rem 4.10). This reduction applies to *all* Bayesian zero-sum games. Next, we devise a PTAS for the dual signaling problem for our class of extended network security games (Theorem 4.14).

Theorem 4.10 (Dual signaling to signaling). *A polytime algorithm for the dual signaling problem with precision ε gives a 5ε -approximation algorithm for the signaling problem. In particular, a PTAS for the dual signaling problem yields a PTAS for the signaling problem.*

To prove Theorem 4.10, we utilize the ellipsoid method, specifically, part (i) of Theorem 2.2, adapting the standard transformation from separation to optimization to take into account the additive error in the dual separation problem. To circumvent the technical difficulties caused by the infinite-dimensionality of (P), we approximate (P) by a finite-dimensional LP, where we restrict the variables in (P), and, analogously the constraints in (D) to a suitable δ -net of Δ_M . Let $\mathcal{I} = (\Theta, \{\mathcal{A}^\theta\}, \lambda)$ be a Bayesian zero-sum game. Recall that $|\mathcal{A}_{i,j}^\theta| \leq 1$ for all θ, i, j . For $\delta \in (0, 1]$ with $1/\delta \in \mathbb{Z}$, and $\mu \in \Delta_M$, define

$$S_\delta := \{\mu' \in \Delta_M : \mu'_\theta/\delta \in \mathbb{Z} \quad \forall \theta \in \Theta\}, \quad S_\delta(\mu) := \{\mu' \in S_\delta : \|\mu - \mu'\|_\infty \leq \delta\}.$$

Claim 4.11. *Fix $\mu \in \Delta_M$. For any $\mu' \in S_\delta(\mu)$, we have $|\text{val}(\mu) - \text{val}(\mu')| \leq M\delta$. Hence, we can efficiently find $\hat{\mu} \in S_\delta(\mu)$ such that $w^T \hat{\mu} - \text{val}(\hat{\mu}) \leq w^T \mu - \text{val}(\mu) + M\delta$.*

Proof. The entries in \mathcal{A}^μ and $\mathcal{A}^{\mu'}$ differ by at most $M\delta$. Hence for every mixed-strategy profile $(x, y) \in \Delta_r \times \Delta_c$, we have $|x^T(\mathcal{A}^\mu - \mathcal{A}^{\mu'})y| \leq M\delta$, and therefore $|\text{val}(\mu) - \text{val}(\mu')| \leq M\delta$.

We can efficiently find $\hat{\mu} \in S_\delta(\mu)$ that minimizes $w^T \mu'$ over $\mu' \in S_\delta(\mu)$ since this can be cast as an LP. Then, we have $w^T \hat{\mu} - \text{val}(\hat{\mu}) \leq w^T \mu - (\text{val}(\mu) - M\delta)$. \blacksquare

We work with the following finite-dimensional counterparts of (P) and (D) and argue that this approximation only yields a small error.

$$\begin{aligned} \max \quad & \sum_{\mu \in S_\delta} \alpha_\mu \text{val}(\mu) \\ \text{s.t.} \quad & \sum_{\mu \in S_\delta} \alpha_\mu \mu = \lambda; \quad \alpha \geq 0. \end{aligned} \quad (\text{P}_\delta) \quad \min \quad w^T \lambda \quad (\text{D}_\delta) \\ \text{s.t.} \quad & w^T \mu \geq \text{val}(\mu) \quad \forall \mu \in S_\delta.$$

Since $\lambda \in \text{conv}(S_\delta(\lambda))$, (P_δ) is feasible for any $\lambda \in \Delta_M$. Clearly, any solution to (P_δ) gives a solution to (P) of equal value. The converse is also approximately true.

Lemma 4.12. *Any feasible solution α to (P) of value v gives a solution to (P_δ) of value at least $v - M\delta$. Hence, $\text{opt}(\text{P}_\delta) \geq \text{opt}(\text{P}) - M\delta$.*

Proof. This is an easy consequence of Claim 4.11. For any $\mu \in S$, let $\tau^{(\mu)} \in \Delta_{S_\delta(\mu)}$ be some convex decomposition of μ . Then

$$\lambda_\theta = \sum_{\mu \in S} \alpha_\mu \mu_\theta = \sum_{\mu \in S} \alpha_\mu \sum_{\mu' \in S_\delta(\mu)} \tau_{\mu'}^{(\mu)} \mu'_\theta = \sum_{\mu' \in S_\delta} \mu'_\theta \sum_{\mu \in S} \alpha_\mu \tau_{\mu'}^{(\mu)}.$$

Thus, setting $\alpha'_{\mu'} := \sum_{\mu \in S} \alpha_\mu \tau_{\mu'}^{(\mu)}$ for all $\mu' \in S_\delta$, we obtain that α' is a feasible solution to (P_δ) . To compare the objective values of α and α' , note that

$$\begin{aligned} \sum_{\mu \in S} \alpha_\mu \text{val}(\mu) &= \sum_{\mu} \alpha_\mu \text{val}(\mu) \sum_{\mu' \in S_\delta(\mu)} \tau_{\mu'}^{(\mu)} \leq \sum_{\mu} \alpha_\mu \sum_{\mu' \in S_\delta(\mu)} \tau_{\mu'}^{(\mu)} (\text{val}(\mu') + M\delta) \\ &= \sum_{\mu' \in S_\delta} \alpha'_{\mu'} \text{val}(\mu') + M\delta. \end{aligned} \quad \blacksquare$$

Now the basic idea is to solve (D_δ) with the ellipsoid method using the algorithm \mathcal{B} to obtain a separation oracle for (D_δ) with an additive error. In the course of solving (D_δ) , we also obtain a polynomial-size LP consisting of the violated inequalities of (D_δ) returned by the separation oracle during the execution of the ellipsoid method whose optimal value is the same as $\text{opt}(D_\delta)$. Taking the dual of this compact LP yields an LP of the same form as (P_δ) but with α_μ variables for only polynomially many points in S_δ ; solving this yields the desired approximate signaling scheme. The additive error in the separation oracle for (D_δ) complicates the arguments slightly.

We now discuss the details. Set $\delta = \varepsilon/M$. Let \mathcal{B} be the algorithm for solving the dual signaling problem with precision ε . For a given $\nu, \epsilon \in \mathbb{R}$, consider the set $Q(\nu, \epsilon) := \{w \in \mathbb{R}^M : w^T \lambda \leq \nu, w^T \mu \geq \text{val}(\mu) - \epsilon \ \forall \mu \in S_\delta\}$. Note that the constraints of $Q(\nu, \epsilon)$ have encoding length $\text{poly}(M, \text{size of } (\lambda, \nu, \delta, \epsilon))$. For a given ν and $w \in \mathbb{R}^M$, we can determine if $w \in Q(\nu, \varepsilon)$, or find a hyperplane separating w from $Q(\nu, -2\varepsilon)$, as follows. We first check if $w^T \lambda \leq \nu$ and if not, then return this as the separating hyperplane. We run \mathcal{B} on the input $(\mathcal{I}, w, \varepsilon)$. If \mathcal{B} determines that we are case (i), then it also returns $\mu \in \Delta_M$ with $\text{val}(\mu) \geq w^T \mu - \varepsilon$. By Claim 4.11, we can then find $\hat{\mu} \in S_\delta(\mu)$ such that $w^T \hat{\mu} - \text{val}(\hat{\mu}) \leq 2\varepsilon$, so we can use $w^T \hat{\mu} - \text{val}(\hat{\mu}) \geq 2\varepsilon$ to separate w from $Q(\nu, -2\varepsilon)$. If \mathcal{B} determines that we are in case (ii), then we are certainly not in case (i), so we have $\text{val}(\mu) \leq w^T \mu + \varepsilon$ for all $\mu \in \Delta_M$, which implies that $w \in Q(\nu, \varepsilon)$.

So for a fixed ν , in polynomial time, the ellipsoid method either certifies that $Q(\nu, -2\varepsilon) = \emptyset$ or returns a point in $Q(\nu, \varepsilon)$. We find the smallest ν (via binary search) such that the latter case happens; call this value ν^* . Then,

$$\nu^* \geq \left(\min w^T \lambda \quad \text{s.t.} \quad w^T \mu \geq \text{val}(\mu) - \varepsilon \ \forall \mu \in S_\delta \right) = \text{opt}(P_\delta) - \varepsilon.$$

The equality above follows since the dual of the minimization LP is above is (P_δ) with the objective function changed to $\sum_{\mu \in S_\delta} \alpha_\mu (\text{val}(\mu) - \varepsilon) = \sum_{\mu \in S_\delta} \alpha_\mu \text{val}(\mu) - \varepsilon$. For any $\epsilon > 0$, running the ellipsoid method for $\nu = \nu^* - \epsilon$ yields a polynomial-size certificate for the emptiness of $Q(\nu^* - \epsilon, -2\varepsilon)$ consisting of the inequality $w^T \lambda \leq \nu^* - \epsilon$ and the polynomially many violated inequalities $w^T \mu - \text{val}(\mu) \geq 2\varepsilon$ returned during the execution of the ellipsoid method. Let $T \subseteq S_\delta$ be the polynomial-size set of points for which we obtain these violated inequalities. By duality,

$$\begin{aligned} \nu^* - \epsilon &< \left(\min w^T \lambda \quad \text{s.t.} \quad w^T \mu \geq \text{val}(\mu) + 2\varepsilon \ \forall \mu \in T \right) \\ &= 2\varepsilon + \left(\max \sum_{\mu \in T} \alpha_\mu \text{val}(\mu) \quad \text{s.t.} \quad \sum_{\mu \in T} \mu \alpha_\mu = \lambda, \quad \alpha \geq 0 \right). \end{aligned}$$

Thus, solving the polynomial-size LP inside the parentheses yields a signaling scheme of value at least $\text{opt}(P_\delta) - 3\varepsilon - \epsilon$, so taking $\epsilon = \varepsilon$ and using Lemma 4.12, we obtain a signaling scheme of value at least $\text{opt}(P) - 5\varepsilon$. \blacksquare

Definition 4.13 (γ -Lipschitz). A matrix $A \in \mathbb{R}^{r \times c}$ is γ -Lipschitz if $\|x^T A - x'^T A\|_\infty \leq \gamma \|x - x'\|_\infty$ for all $x, x' \in \Delta_r$. An extended security game specified by matrices \overline{A}, B, D (see (1), (2)) is γ -Lipschitz if D^T is γ -Lipschitz. We place no constraints on the matrices \overline{A} and B .

Observe that an extended security game specified by matrices \overline{A}, B, D (see (1), (2)) is γ -Lipschitz if D^T is γ -Lipschitz. We place no constraints on the matrices \overline{A} and B . We design a simple PTAS for the dual signaling problem on γ -Lipschitz extended security games, for constant γ . By Theorem 4.10, this yields a PTAS for the signaling problem for γ -Lipschitz extended security games.

Theorem 4.14. *There is a PTAS for the dual signaling problem on γ -Lipschitz extended security games. This yields a PTAS for the signaling problem on γ -Lipschitz extended-security games.*

Proof. Given Theorem 4.10, we only need to prove the first statement. Let $(\mathcal{I}, w, \varepsilon)$ be the input to the dual signaling problem where \mathcal{I} is a γ -Lipschitz extended security game. Set $\varepsilon' = \varepsilon/\gamma$. Our algorithm simply finds $\hat{\mu} = \operatorname{argmax}_{\mu \in S_{\varepsilon'}} (\operatorname{val}(\mu) - w^T \mu)$ by exhaustive search. If $\operatorname{val}(\hat{\mu}) - w^T \hat{\mu} \geq 0$, we state that we are in case (i) and return $\hat{\mu}$; else we state that we are in case (ii).

First, note that the algorithm runs in time $\operatorname{poly}(\operatorname{size of } \mathcal{I}, M^{\frac{\gamma}{\varepsilon}})$, since $|S_{\varepsilon'}| \leq \binom{M}{1/\varepsilon'} (\frac{1}{\varepsilon'})^{1/\varepsilon'}$ (there are $\binom{M}{1/\varepsilon'}$ choices for the support, and at most $\frac{1}{\varepsilon'}$ choices for each of the at most $\frac{1}{\varepsilon'}$ coordinates in the support).

Let μ^* maximize $\operatorname{val}(\mu) - w^T \mu$, and x^* be the equilibrium strategy for the row player in the resulting zero-sum game. We claim that $\operatorname{val}(\mu^*) - w^T \mu^* \leq \operatorname{val}(\hat{\mu}) - w^T \hat{\mu} + \varepsilon$, which shows that we correctly solve the dual signaling problem: if case (i) applies, then $\operatorname{val}(\hat{\mu}) - w^T \hat{\mu} \geq 0$; if case (ii) applies, then clearly, $\operatorname{val}(\hat{\mu}) - w^T \hat{\mu} < -\varepsilon$.

We now prove the claim. Since $\mu^* \in \operatorname{conv}(S_{\varepsilon'}(\mu^*))$, there exists some $\mu' \in S_{\varepsilon'}(\mu^*)$ such that $x^{*T} B \mu^* - w^T \mu^* \leq x^{*T} B \mu' - w^T \mu'$. Further, since D^T is γ -Lipschitz, for all $j \in [c]$, $(x^{*T} \bar{A} + \mu^{*T} D^T)_j \leq (x^{*T} \bar{A} + \mu'^T D^T)_j + \gamma \varepsilon'$. Combining these inequalities yields that $\operatorname{val}(\mu') - w^T \mu' \geq \operatorname{val}(\mu^*) - w^T \mu^* - \varepsilon$. \blacksquare

5 Bayesian network routing games

We now consider the signaling problem in Bayesian network routing games and prove an *optimal* inapproximability result for linear latency functions: It is *NP*-hard to obtain a multiplicative approximation better than $4/3$ (Theorem 5.1), and this approximation is achieved for linear latency functions by a simple signaling scheme that simply reveals the state of nature (Theorem 5.4).

A *network routing game* is a tuple $\Gamma = (G = (V, E), \{l_e\}_{e \in E}, \{(s_i, t_i, d_i)\}_{i \in [k]})$, where G is a directed graph with latency function $l_e : \mathbb{R}_+ \mapsto \mathbb{R}_+$ 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 f is given by $l_e(f_e)$, where f_e is the total flow on e . The latency of a path P is $l_P(f) := \sum_{e \in P} l_e(f_e)$. The total latency of a flow f is $C(l; f) := \sum_{e \in E} f_e l_e(f_e)$; an optimal flow is a feasible flow with minimum latency. A feasible flow f in a routing game is a *Nash flow* (also called a *Wardrop flow*), if each player chooses a minimum latency path; that is, for all i , all s_i - t_i paths P, Q with $f_e > 0$ for all $e \in P$, $l_P(f) \leq l_Q(f)$. All Nash flows have the same total latency (see, e.g., [RT02]).

In a *Bayesian network routing game*, the edge latency functions $\{l_e^\theta\}_{e \in E}$ 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 $l_e^\mu(x_e) := \sum_{\theta \in \Theta} \mu_\theta l_e^\theta(x_e)$. Define $\operatorname{val}(\mu) := C(l^\mu; f^\mu)$, where f^μ is the Nash flow for latency functions $\{l_e^\mu\}$. The *signaling problem in a Bayesian routing game* is to determine $(\alpha_\mu)_{\mu \in \Delta_M} \geq 0$ of finite support specifying a convex decomposition of λ (i.e., $\sum_{\mu \in \Delta_M} \alpha_\mu \mu = \lambda$) that minimizes the expected latency of the Nash flow, $\sum_{\mu \in \Delta_M} \alpha_\mu \operatorname{val}(\mu)$.

Theorem 5.1. *For any $\epsilon > 0$, obtaining a $(4/3 - \epsilon)$ -approximation for the signaling problem in Bayesian routing games is NP-hard, even in single-commodity games with linear latency functions.*

Let (G, s, t, d) be a single-commodity routing game. We reduce from the problem of determining edge tolls $\tau \in \mathbb{R}_+^E$ that minimize $C(l + \tau; f^{NE}(\tau))$, where $l + \tau$ denotes the collection of latency functions $\{l_e(x) + \tau_e\}_e$ and $f^{NE}(\tau)$ is the Nash flow for $l + \tau$. Note that $C(l + \tau; f) = \sum_e f_e(l_e(f_e) + \tau_e)$ takes into account the contribution from tolls; we refer to this as the total cost of f . By *optimal tolls*, we mean tolls τ that minimize $C(l + \tau; f^{NE}(\tau))$.

Theorem 5.2 ([CDR06]). *There are optimal tolls where the toll on every edge is 0 or ∞ . If $P \neq NP$, there is no $(\frac{4}{3} - \epsilon)$ -approximation algorithm for the problem of computing optimal tolls in networks with linear latency functions, for any $\epsilon > 0$.*

Let $\Gamma = (G = (V, E), l, s, t, d)$ be an instance of a routing game with linear latencies. Let $m = |E| \geq 2$. By scaling latency functions suitably, we may assume that $d = 1$. Then, for any latency functions l' , the latency of the Nash flow for l' equals the common delay of all flow-carrying s - t paths. Let $L = C(l; f^{NE})$ be the latency of the Nash flow for l . Let τ^* be optimal $\{0, \infty\}$ -tolls, $L^* = C(l + \tau^*, f^{NE}(\tau^*))$ be the optimal cost, and $K^* := \{e \in E : \tau_e^* = \infty\}$. We can view τ^* as simulating the removal of edges in K^* .

We create the following Bayesian routing game. Let $(G_1 = (V_1, E_1), s_1, t_1)$ and $(G_2 = (V_2, E_2), s_2, t_2)$ 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 $e' \in E$, set the latency function in the new graph $h_e(x) = l_{e'}(x)$, and set $h_e(x) = 0$ for $e = (s, s_1), (s, s_2), (t_1, t), (t_2, t)$. The states of nature correspond to edges in H . We set $\lambda_\theta = 1/m^2$ for all $\theta \in E_1 \cup E_2$; the remaining $1 - \frac{2}{m}$ mass is spread equally on $(s, s_1), (s, s_2)$. We set $h_e^\theta(x) = h_e(x) + 8m^3L$ if $\theta = e$ and $h_e(x)$ otherwise. Our Bayesian routing game is $((G, \{h_e^\theta\}_{\theta, e}, s, t, d), \lambda)$.

The idea here is that state θ encodes the removal of edge θ : specifically, if $\mu_\theta = \Omega(\frac{1}{m})$ for a posterior μ , then h^μ simulates removing edge θ due to the large constant term $8m^3L$. Let K_i be the edge-set corresponding to K^* in G_i , for $i = 1, 2$. The prior λ is set up so that: (a) it admits a convex decomposition into posteriors μ^1, μ^2 , where h^{μ^1} simulates that $G_i \setminus K_i$ is connected to s and G_{3-i} is disconnected from s ; and (b) any convex-decomposition of λ must be such that a large weight is placed on posteriors μ , where h^μ simulates that only one of G_i is connected to s , so that $\{\mu_e 8m^3L\}_{e \in E_i}$ yields tolls τ for edges in E such that $C(l + \tau, f^{NE}(\tau)) \leq \text{val}(\mu)$. Lemma 5.3 makes the statements in (a) and (b) precise, and Theorem 5.1 follows immediately from Lemma 5.3 and Theorem 5.2.

Lemma 5.3. *There is a signaling scheme for the above Bayesian routing game with latency L^* . Further, given a signaling scheme α for the above Bayesian routing game with expected latency L' , one can obtain tolls τ such that the routing game $(G, l + \tau, s, t, d)$ has Nash latency at most $\frac{L'}{1-4/m}$.*

Proof. We first show the existence of a signaling scheme with latency L^* . Define posterior $\mu^1 \in \Delta_{EH}$ as: $\mu_\theta^1 = 2/m^2$ for all $\theta \in K_1 \cup E_2 \setminus K_2$, $\mu_{(s, s_2)}^1 = (1 - 2/m)$. Define μ^2 symmetrically as: $\mu_\theta^2 = 2/m^2$ for all $\theta \in K_2 \cup E_1 \setminus K_1$, $\mu_{(s, s_1)}^2 = (1 - 2/m)$. Then $\lambda = (\mu^1 + \mu^2)/2$, and this is our signaling scheme. We will show that $\text{val}(\mu^1) = \text{val}(\mu^2) \leq L^*$, proving the lemma.

Consider distribution μ^1 ; the argument for μ^2 is symmetrical. The idea is that an edge e with $\mu_e^1 > 0$ has $h_e^{\mu^1}(x) \geq 8mL$, which effectively deletes e from H ; other edges have $h_e^{\mu^1}(x) = h_e(x)$. So

μ^1 simulates retaining edges in $G_1 \setminus K_1$. Let $f = f^{NE}(\tau^*)$ be the Nash flow in the routing game $(G, l + \tau^*, s, t, d)$. So $C(l + \tau^*; f) = L^*$. Recall that $d = 1$, so every s - t path in G has latency at least L^* . Then the flow that sends d on edges (s, s_1) and (t_1, t) and f on edges of G_1 , is feasible. On every edge $e \in E(H)$ with positive flow, $\mu_e^1 = 0$, so the latency of this flow under h^{μ^1} is L^* . Further, this is a Nash flow for h^{μ^1} : any s - t path P either contains an edge with $\mu_e^1 > 0$, and if not, contains an s_1 - t_1 path; in the latter case, there is a corresponding s - t path Q in G , and the latency of P under h^{μ^1} equals $(l + \tau^*)_Q(f)$, which is at least L^* since f is the Nash flow for $l + \tau^*$.

Next, we show how to obtain the required tolls from the signaling scheme α (with expected latency L'). Assume $L' \leq L$, otherwise $\tau = 0$ suffices. At least $(1 - 4/m)$ of the probability mass of α 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 $\text{val}(\mu') \leq \frac{L'}{1 - 4/m}$. Assume $\mu'_{(s, s_1)} \geq \frac{1}{2m}$; the other case is symmetric. Let $f = f^{\mu'}$ be the Nash flow for latency functions $h^{\mu'}$. (Again, since $d = 1$, every s - t path P in H with $f_e > 0$ for all $e \in P$ satisfies $h_P^{\mu'}(f) = \text{val}(\mu')$.)

Since $h_{(s, s_1)}^{\mu'} \geq 4m^2L > \text{val}(\mu')$, we must have $f_{(s, s_1)} = 0$, so f is supported on G_2 . Abusing notation, for $e \in E_2$, we also use e to denote the corresponding edge in E . For every $e \in E_2$, we have $h_e^{\mu'}(x) = l_e(x) + \mu'_e 8m^3L$. Thus, defining $\tau_e = \mu'_e 8m^3L$ for all $e \in E$, we obtain that f restricted to E_2 is a Nash flow for $(G, l + \tau, s, t, d)$, and its latency is at most $\text{val}(\mu')$. This is easy to see, since every s - t path P in G corresponds to an s_2 - t_2 path Q in H , and $(l + \tau)_P(f) = h_Q^{\mu'}(f)$. ■

Theorem 5.4. *The full-revelation signaling scheme, i.e., revealing the state of nature, has the price of anarchy for the underlying latency functions as its approximation ratio. In particular, for linear latencies, it achieves a $\frac{4}{3}$ -approximation.*

Proof. 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 $\frac{4}{3}$ [RT02].

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.

Slightly abusing notation, we use f^θ to denote the Nash flow with respect to the latency functions $\{l_e^\theta\}$. We use \tilde{f}^θ to denote the optimal flow for latency functions $\{l_e^\theta\}$. Let ρ be the price of anarchy for the collection $\{l_e^\theta\}_{e \in E, \theta \in \Theta}$ of latency functions, so we have $C(l^\theta; \tilde{f}^\theta) \geq C(l^\theta; f^\theta)/\rho$ for all $\theta \in \Theta$. The full-revelation signaling scheme has cost $\sum_{\theta \in \Theta} \lambda_\theta C(l^\theta; f^\theta)$.

Consider any signaling scheme α . Its cost is

$$\begin{aligned} \sum_{\mu \in \Delta_M} \alpha_\mu \text{val}(\mu) &= \sum_{\mu} \alpha_\mu C(l^\mu; f^\mu) = \sum_{\mu} \alpha_\mu \sum_{\theta \in \Theta} \mu_\theta C(l^\theta; f^\mu) \geq \sum_{\mu, \theta} \alpha_\mu \mu_\theta C(l^\theta; \tilde{f}^\theta) \\ &\geq \sum_{\mu, \theta} \alpha_\mu \mu_\theta C(l^\theta; f^\theta)/\rho = \sum_{\theta} \lambda_\theta C(l^\theta; f^\theta)/\rho. \end{aligned} \quad \blacksquare$$

6 Extensions: hardness results for related problems

6.1 Maximum prior problem

We study the closely-related problem of finding $\mu \in \Delta_M$ that maximizes $\text{opt}(\mu)$. The proof of Theorem 4.2 in fact shows that *a PTAS for the maximum-prior problem yields a PTAS for threshold signaling*. Theorem 6.1 uses this implication to rule out a PTAS for the maximum prior problem under the exponential time hypothesis (ETH) by giving a simple, clean reduction from the best-Nash problem in *general-sum* two-player games, for which a PTAS is ruled out by [BKW15]. Theorem 6.1 establishes the *optimal* hardness result for the maximum prior problem, since a quasi-PTAS for the maximum prior problem was recently presented in [CCD⁺15].

Theorem 6.1 also implies that the general maximum prior problem studied in [CCD⁺15] does not have a PTAS under the ETH, when the objective function is $O(1)$ -Lipschitz but not $O(1)$ -noise-stable. (For this case, [CCD⁺15] ruled out a PTAS assuming hardness of planted clique.) This is because the objective function (minimax value) for signaling in zero-sum games is $O(1)$ -Lipschitz.

Theorem 6.1. *Assuming ETH, there is a constant ε_0 such that, any algorithm that returns an (additive) ε_0 -approximation for the maximum prior problem, even for extended security games, must run in quasipolynomial, i.e., $n^{\tilde{\Omega}(\log^{1-o(1)} n)}$, time. In particular, assuming ETH, there is no PTAS for the maximum prior problem, even for extended security games.*

Recall that, as noted earlier, the proof of Theorem 4.2 shows that, for any ε , a polytime ε -approximation for the maximum prior problem yields a polytime algorithm for the threshold signaling problem with precision parameter 2ε . Thus, it suffices to show that, assuming ETH, there is some constant ε_0 such that solving the threshold signaling problem with precision parameter ε_0 , even for extended security games, requires quasipolynomial running time. To show this, we reduce from the problem of finding an ε -Nash equilibrium in a general two-player game with ε -approximate social welfare, and utilize the following hardness result for this problem.

Theorem 6.2 ([BKW15]). *Assuming ETH, there is a constant $\varepsilon^* > 0$ such that any algorithm for finding an ε^* -approximate Nash equilibrium with social welfare at least $OPT - \varepsilon^*$ in a general bimatrix game requires $n^{\tilde{\Omega}(\log^{1-o(1)} n)}$ time, where OPT is the optimal welfare of a Nash equilibrium.*

Proof of Theorem 6.1. Let $(\mathcal{R}, \mathcal{C})$ be a bimatrix game, where $\mathcal{R}, \mathcal{C} \in [-1, 1]^{m \times n}$ are the payoffs for the row- and column- players respectively. To avoid confusion with the extended security game, we refer to the row- and column- players in the bimatrix game as the \mathcal{R} - and \mathcal{C} - players. A pair of mixed strategies (x, y) for the \mathcal{R} - and \mathcal{C} - players respectively is an ε -approximate equilibrium if:

$$x^T(\mathcal{R} + \mathcal{C})y - \max_{i \in [m]}(\mathcal{R}y)_i - \max_{j \in [n]}(x^T\mathcal{C})_j \geq -\varepsilon. \quad (4)$$

The social welfare of (x, y) is defined as $x^T(\mathcal{R} + \mathcal{C})y$. Let OPT be the maximum social welfare of a (mixed) Nash equilibrium of $(\mathcal{R}, \mathcal{C})$. Note that $-2 \leq OPT \leq 2$.

We construct an extended security game where the states of nature correspond to the pure strategies of the \mathcal{C} -player (in the bimatrix game), and the row-player's pure strategies (in the extended security game) correspond to the \mathcal{R} -player's strategies (in the bimatrix game). We will set things up so that the expected payoff in the extended security game to the row player under a posterior distribution μ and when he plays a mixed strategy x is a linear combination of the LHS of (4) (viewing (x, μ) as a mixed-strategy profile for the bimatrix game $(\mathcal{R}, \mathcal{C})$) and the social welfare

of (x, μ) in the bimatrix game $(\mathcal{R}, \mathcal{C})$. Let $\epsilon > 0$ be a parameter. The payoffs in the extended security game will have absolute value at most $1 + O(1/\epsilon)$. We will show that solving the threshold signaling problem for the resulting extended security game with threshold $\eta = \eta' - \epsilon$, and precision parameter ϵ yields a 6ϵ -approximate Nash equilibrium of $(\mathcal{R}, \mathcal{C})$ with social welfare at least $\eta' - 2\epsilon$, whenever there is a Nash equilibrium of $(\mathcal{R}, \mathcal{C})$ with social welfare at least η' or we state that we are in case (i) of the threshold signaling problem. So via binary search, we can obtain a 6ϵ -approximate Nash equilibrium of $(\mathcal{R}, \mathcal{C})$ with social welfare at least $OPT - 3\epsilon$. Thus, setting $\epsilon_0 = \Theta(\epsilon^{*2})$,² where ϵ^* is as given by Theorem 6.2, we obtain that, assuming ETH, the threshold signaling problem with precision parameter ϵ_0 requires quasipolynomial time, completing the proof.

We proceed to describe the extended security game and prove the desired claim. We set $\Theta = [n]$, so $M = n$. The row-player's pure strategy set is $[m]$, and the column-player's pure-strategy set is $[m] \times [n]$, so the row- and column- players have $r = m$ and $c = mn$ pure strategies respectively. The $r \times c$ matrix \bar{A} , $r \times M$ matrix B , and $c \times M$ matrix D in the extended security game are

$$\begin{aligned}\bar{A}_{i,(i',j)} &= -\frac{1}{\epsilon}\mathcal{C}_{i,j} && \forall i \in [m], (i', j) \in [m] \times [n] \\ B_{i,j} &= \left(1 + \frac{1}{\epsilon}\right)(\mathcal{R}_{i,j} + \mathcal{C}_{i,j}) && \forall i \in [m], j \in [n] \\ D_{(i,j'),j} &= -\frac{1}{\epsilon}\mathcal{R}_{i,j} && \forall (i, j') \in [m] \times [n], i \in [m].\end{aligned}$$

Claim 6.3. *For all $\mu \in \Delta_M$, $x \in \Delta_r$, we have*

$$\min_{k \in [c]} (x^T \bar{A} + \mu^T D^T)_j = -\frac{1}{\epsilon} \left(\max_{i \in [m]} (\mathcal{R}\mu)_i + \max_{j \in [n]} (x^T \mathcal{C})_j \right).$$

Proof. Consider any column-player strategy $k = (i', j') \in [m] \times [n]$. We have

$$\begin{aligned}(x^T \bar{A})_k + (\mu^T D^T)_k &= \sum_{i \in [m]} x_i \bar{A}_{i,(i',j')} + \sum_{j \in [n]} \mu_j D_{(i',j'),j} \\ &= -\frac{1}{\epsilon} \left(\sum_{i \in [m]} x_i \mathcal{C}_{i,j'} + \sum_{j \in [n]} \mu_j \mathcal{R}_{i',j} \right) = -\frac{1}{\epsilon} \left((x^T \mathcal{C})_{j'} + (\mathcal{R}\mu)_{i'} \right). \quad \blacksquare\end{aligned}$$

It follows from Claim 6.3 that for any $\mu \in \Delta_M$, we have

$$\begin{aligned}\text{val}(\mu) &= \max_{x \in \Delta_r} \left(x^T B \mu + \min_{j \in [c]} (x^T \bar{A} + \mu^T D^T)_j \right) \\ &= \max_{x \in \Delta_m} \left[\left(1 + \frac{1}{\epsilon}\right) x^T (\mathcal{R} + \mathcal{C}) \mu - \frac{1}{\epsilon} \left(\max_{i \in [m]} (\mathcal{R}\mu)_i + \max_{j \in [n]} (x^T \mathcal{C})_j \right) \right] \\ &= \max_{x \in \Delta_m} \left[x^T (\mathcal{R} + \mathcal{C}) \mu - \frac{1}{\epsilon} \left(\max_{i \in [m]} (\mathcal{R}\mu)_i + \max_{j \in [n]} (x^T \mathcal{C})_j - x^T (\mathcal{R} + \mathcal{C}) \mu \right) \right] \quad (5)\end{aligned}$$

Now suppose we solve the threshold signaling problem with threshold $\eta = \eta' - \epsilon$ (where $\eta' \leq 2$) and precision parameter ϵ . Suppose (x^*, μ^*) is a Nash equilibrium of $(\mathcal{R}, \mathcal{C})$ with social welfare at least η' . It follows from (5) that $\text{val}(\mu^*) \geq \eta'$. So we are not in case (ii) of the threshold signaling

²The $\Theta(\epsilon^{*2})$ is because we need additive error $\Theta(\epsilon^*)$ when payoffs are bounded in absolute value by $1 + O(1/\epsilon^*)$; when we scale payoffs so that they lie in $[-1, 1]$, this translates to an $\Theta(\epsilon^{*2})$ -approximation.

problem, and must obtain $\mu \in \Delta_n$ such that $\text{val}(\mu) \geq \eta - \epsilon = \eta' - 2\epsilon$. From (5), this implies that there is $x \in \Delta_m$ such that $x^T(\mathcal{R} + \mathcal{C})\mu \geq \eta' - 2\epsilon$ and

$$\max_{i \in [m]} (\mathcal{R}\mu)_i + \max_{j \in [n]} (x^T \mathcal{C})_j - x^T(\mathcal{R} + \mathcal{C})\mu \leq \epsilon \left(x^T(\mathcal{R} + \mathcal{C})\mu - \eta' + 2\epsilon \right) \leq 6\epsilon$$

where the last inequality follows since $\mathcal{R}, \mathcal{C} \in [-1, 1]^{m \times n}$. The same calculation holds whenever we state that we are in case (i) and return $\mu \in \Delta_n$. \blacksquare

6.2 Hardness with other equilibrium notions

It is known that in zero-sum games, correlated equilibria and Nash equilibria are payoff-equivalent, that is, they yield the same payoffs (this was also noted in [Dug14]). Thus, our hardness results extend to the case of correlated equilibria, as well as other notions of stability that are payoff-equivalent to Nash equilibria in zero-sum games. To see this, note that for $\mu \in \Delta_M$, due to the payoff equivalence, $\text{val}(\mu)$ is also the payoff of the row player in any correlated equilibrium in the zero-sum game specified by \mathcal{A}^μ . Hence, the statement of the signaling problem and its optimal value remain unchanged. Further, any signaling scheme for correlated equilibria gives a signaling scheme for Nash equilibria of equal value. This immediately extends *all* our hardness results (Theorem 4.1, Theorem 4.3, Theorem 4.4, Theorem 6.1) to correlated equilibria (and other payoff-equivalent equilibria).

6.3 Signaling with general objective functions

We now consider a more general signaling problem in Bayesian zero-sum games, where the principal's value may depend on the players' strategies, and show that it is *NP*-hard to obtain a PTAS. Formally, we have a Bayesian zero-sum game $(\Theta, \{\mathcal{A}^\theta\}_{\theta \in \Theta}, \lambda)$ and a $\Theta \times r \times c$ *principal objective tensor* $\mathcal{F} = (\mathcal{F}^\theta(i, j))$; that is, $\mathcal{F}^\theta \in [-1, 1]^{r \times c}$ for all $\theta \in \Theta$. We now define $\text{val}(\mu) = \max_{(x_\mu, y_\mu) \in \text{NE}(\mathcal{A}^\mu)} x_\mu^T (\sum_{\theta} \mu_\theta \mathcal{F}^\theta) y_\mu$, where $\text{NE}(\mathcal{A}^\mu)$ is the set of all (exact) Nash equilibria of \mathcal{A}^μ . As before, we seek a signaling scheme (Σ, α, μ) that maximizes $\sum_{\sigma \in \Sigma} \alpha_\sigma \text{val}(\mu_\sigma)$.

Theorem 6.4. *Given a Bayesian zero-sum game $(\Theta, \{\mathcal{A}^\theta\}_{\theta \in \Theta}, \lambda)$, and a principal objective tensor \mathcal{F} , it is *NP*-hard to distinguish whether the optimal signaling scheme has value 0 or at least $\frac{1}{2}$.*

The *NP*-hardness proof follows from a reduction from the *balanced vertex cover* (BVC) problem proposed in [CS06]. 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 and the prior is $\lambda = \mathbb{1}_n/n$.

The row player's pure strategy is to pick a node $v_1 \in V$, and the column player's pure strategy is to either pick a vertex v , an edge e , or a special strategy s . We design the *column player's* payoff as follows. The payoff for strategy

$$v \text{ is } \begin{cases} \frac{n}{n-2} & \text{if } v \notin \{\theta, v_1\}, \\ 0 & \text{otherwise.} \end{cases} \quad e \text{ is } \begin{cases} \frac{n}{n-2} & \text{if } e \text{ is not incident with } \theta, \\ 0 & \text{otherwise.} \end{cases} \quad s \text{ is } 1.$$

The principal's objective tensor is set up so that he is interested only in getting the column player to play the strategy s , that is, $\mathcal{F}^\theta(v, s) = 1$ for all $\theta, v \in V$; all other entries of \mathcal{F} are 0.

Lemma 6.5. *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 $\mu_v^1 = \frac{2}{n}$ for all $v \in C$ (and 0 otherwise), and $\mu_v^2 = \frac{2}{n}$ for all $v \notin C$. For posterior μ^1 , there is a Nash equilibrium where the row player chooses the mixed strategy x that picks $v_1 \in V \setminus C$ uniformly at random and the column player chooses strategy s ; thus, the principal gets a value of 1. This is because every node and edge is “protected” with probability at least $\frac{2}{n}$; the payoff of the column player for a pure strategy v or e is therefore at most $\frac{n-2}{n-2}(1 - \frac{2}{n}) \leq 1$. Since μ^1 is chosen 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 $\text{val}(\mu) > 0$, then G has a BVC solution. Let (x, y) be a Nash equilibrium that attains value $\text{val}(\mu)$, that is, $\text{val}(\mu) = x^T(\sum_{\theta} \mu_{\theta} \mathcal{F}^{\theta})y$. Since $\text{val}(\mu) > 0$, we must have $y_s > 0$. For this to happen, every node in V must be protected with probability at least $\frac{2}{n}$. That is, we must have $\frac{n-2}{n-2}(1 - x_v)(1 - \mu_v) \leq 1$ for all $v \in V$. Then, $n - 2 + \sum_v x_v \mu_v = \sum_v (1 - x_v)(1 - \mu_v) \leq n - 2$, which implies that we must have $x_v \mu_v = 0$ and $(1 - x_v)(1 - \mu_v) = 1 - \frac{2}{n}$ for all $v \in V$. So it must be that for all $v \in V$, exactly one of μ_v and x_v is equal to $\frac{2}{n}$. Let $C = \{v : \mu_v > 0\}$. It follows that $|C| = \frac{n}{2}$. The payoff of a column player for an edge $e = (u, v)$ is $\frac{n-2}{n-2}(1 - \mu_u - \mu_v)$, which must be at most 1, so we have $\mu_u + \mu_v \geq \frac{2}{n}$. It follows that C is a vertex cover of G . ■

We remark that it is important to allow the principal’s payoff to depend on specific strategies, and also to enforce exact Nash equilibrium. Intuitively, these two ingredients together make the objective function $\text{val}(\mu)$ very “sensitive” in μ . Moreover, these two conditions are essentially necessary for an NP-hardness result, as Cheng et al. [CCD⁺15] gave a bi-criteria quasi-PTAS for this general signaling problem, i.e., a quasi-polytime algorithm that loses an additive ϵ in the objective as well as in the Nash equilibrium constraints.

Acknowledgments.

The authors are grateful to Shaddin Dughmi for various suggestions, including on the planted clique reduction and for suggesting the network routing games problem. We also thank David Kempe, and Li Han for helpful discussions.

References

- [AIM14] Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz. AM with multiple Merlins. In *IEEE 29th Conference on Computational Complexity, CCC 2014*, pages 44–55, 2014.
- [Ake70] George A Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500, 1970.
- [AV14] Brendan PW Ames and Stephen A Vavasis. Convex optimization for the planted k-disjoint-clique problem. *Mathematical Programming*, 143(1-2):299–337, 2014.
- [BBM13] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *Economic Theory Center Working Paper*, (052-2013), 2013.

- [BKW15] Mark Braverman, Young Kun Ko, and Omri Weinstein. Approximating the best nash equilibrium in $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.
- [Bla51] David Blackwell. Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102, 1951.
- [BMS12] Peter Bro Miltersen and Or Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 234–247, 2012.
- [CCD⁺15] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In *56th Annual Symposium on Foundations of Computer Science (FOCS)*, 2015.
- [CDR06] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.
- [CS06] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic Commerce (EC)*, 2006.
- [DGGP11] Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. Finding hidden cliques in linear time with high probability. In *ANALCO*, pages 67–75. SIAM, 2011.
- [DIR14] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *the 25th ACM Symposium on Discrete Algorithms (SODA)*, 2014.
- [Doe11] Benjamin Doerr. Analyzing randomized search heuristics: tools from probability theory. In *Theory of randomized search heuristics*. World Scientific, 2011.
- [DP09] Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- [Dug14] Shaddin Dughmi. On the hardness of signaling. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 354–363, 2014.
- [EFG⁺12] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 514–531, 2012.
- [FGR⁺13] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh Vempala, and Ying Xiao. 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. ACM, 2013.
- [FK03] Uriel Feige and Robert Krauthgamer. The probable value of the Lovász–Schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2), 2003.
- [FNS07] Tomas Feder, Hamid Nazerzadeh, and Amin Saberi. Approximating nash equilibria using small-support strategies. In *Proceedings of the 8th ACM conference on Electronic Commerce (EC)*, pages 352–354. ACM, 2007.

- [FR10] Uriel Feige and Dorit Ron. Finding hidden cliques in linear time. *DMTCS Proceedings*, (01):189–204, 2010.
- [GD13] Mingyu Guo and Argyrios Deligkas. Revenue maximization via hiding item attributes. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 157–163. AAAI Press, 2013.
- [GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.
- [GLS93] Martin Grötschel, László Lovász, and Lex Schrijver. Geometric algorithms and combinatorial optimization. *Algorithms and Combinatorics*, 2:1–362, 1993.
- [Hir71] Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4):561–574, 1971.
- [HK11] Elad Hazan and Robert Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM Journal on Computing*, 40(1):79–91, 2011.
- [Jer92] Mark Jerrum. Large cliques elude the metropolis process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [JP00] Ari Juels and Marcus Peinado. Hiding cliques for cryptographic security. *Designs, Codes and Cryptography*, 20(3):269–280, 2000.
- [Kuč95] Luděk Kučera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2):193–212, 1995.
- [LRS10] Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2), 2010.
- [MW82] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5), 1982.
- [NY83] Arkadiĭ Semenovich Nemirovski and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. John Wiley and Sons, 1983.
- [RT02] T. Roughgarden and Eva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259, 2002.
- [Rub15] Aviad Rubinfeld. Eth-hardness for signaling in symmetric zero-sum games. *CoRR*, abs/1510.04991, 2015.
- [VFH15] Shoshana Vasserman, Michal Feldman, and Avinatan Hassidim. Implementing the wisdom of waze. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 660–666, 2015.

A Proof of Lemma 4.5

Recall that $\epsilon > 0$, $c_3 \geq 10^3$, and $k = k(n)$ satisfies $k = \omega(\log n)$ and $k = o(\sqrt{n})$, and $r = \Theta(n/k)$. Let $p = \frac{1}{2}$.

First, we proceed as in [Dug14] to reduce the planted-clique problem to the planted-clique-cover problem. Given an instance G of $\mathbf{PClique}(n, p, k)$, we can generate an instance G' of $\mathbf{PCover}(n, p, k, r)$ by planting $r - 1$ additional random k -cliques into G (as in step (2) of Definition 2.1). As noted in [Dug14], because the cliques S_1, \dots, S_r are indistinguishable, recovering a constant fraction of the planted cliques from G' would recover each of S_1, \dots, S_r with constant probability. In particular, it can recover the original planted clique with constant probability.

So our task is the following. Given a graph $G \sim \mathcal{G}(n, p, k, r)$, fix one of the planted k -cliques $S \subseteq V$. We need to show that given a cluster $T \subseteq V$ satisfying $|S \cap T| \geq \epsilon|T|$ and $|S \cap T| \geq c_3 \log n$, we can recover S with high probability. We assume that $r = \frac{5n}{k}$ in the sequel. Our algorithm is similar to (and in fact, simpler than) the one used in [Dug14] to prove a similar planted-clique recovery result (Lemma 3.5 therein). However, we need to recover the planted clique under a *much weaker* (both qualitatively and quantitatively) assumption. In our case, the above requirements on $|S \cap T|$ allow $|T| = \Theta(\log n)$ (which is crucial for the soundness proof in Lemma 4.6 to go through); in [Dug14], the requirement is that $|S \cap T| = \Omega(|S \cup T|)$ with $|S| = \omega(\log^2 n)$, so that we must have $|T| = \omega(\log^2 n)$. This difference in the magnitude of $|T|$ (and hence $|S \cap T|$) poses certain challenges and necessitates certain key changes to the analysis in [Dug14].

We use the following algorithm to recover S :

1. Pick an arbitrary set R of $c_3 \log n$ vertices from $S \cap T$.
2. Let S' be all the common neighbors of R .
3. Let \hat{S} be the vertices in S' with at least $k - 1$ neighbors in S' .

Since S is unknown, we use the following process to simulate Step (1). We first sample roughly $\frac{c_3 \log n}{\epsilon}$ vertices uniformly from T , and try Step (2) and (3) on every subset of $c_3 \log n$ of the sampled vertices. The number of subsets we need to check is polynomial. Moreover, because $|S \cap T| \geq \epsilon|T|$, with high probability, the sampled subset of T will contain $c_3 \log n$ vertices from $S \cap T$, and will encounter this set of $c_3 \log n$ vertices from $S \cap T$ in our enumeration.

We partition the edges of G into E^- and E^+ , where E^- are the background edges added in Step (1) of Definition 2.1, and E^+ are the extra clique-related edges added in Step (2) of Definition 2.1. Let E^i denote the edges of S_i . It is easy to verify that all the nodes in S will survive Step (2) and (3), so $S \subseteq \hat{S}$. We show that in fact, with high probability, no other vertices survive Step (2) and (3) through the following claims.

Claim A.1. *With high probability, we have $|E^-(v, S)| \leq 0.6|S|$ for all $v \notin S$.*

Proof. Since $|S| = \omega(\log n)$, this follows from a straightforward application of the Chernoff bound and the union bound. ■

Claim A.2. *With high probability, there are at most $c_3 \log n$ vertices $v \notin R$ with $|E^-(v, R)| \geq 0.8|R|$.*

Proof. Let $A = \{v \notin R : |E^-(v, R)| \geq 0.8|R|\}$. Then, $\text{bi-density}_{G^-}(R, A) \geq 0.8$. The constant $c_3 = 10^3$ and $\epsilon = 0.3$ satisfy the conditions of Lemma 4.8. So since $|R| \geq c_3 \log n$, we have $|A| < c_3 \log n$ with high probability. ■

Claim A.3. *With high probability, we have $|E^+(v, S)| \leq 12 \log n$ for all $v \notin S$.*

The following lemma will be useful in proving the above claim.

Lemma A.4 (see Ex. 1.13 in [DP09], Lemma 1.19 in [Doe11]). *Let X_1, \dots, X_n be arbitrary binary random variables. Suppose for every i , and every $x_1, \dots, x_{i-1} \in \{0, 1\}$, we have $\Pr[X_i = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_{i-1} = x_{i-1}] \leq p_i$. Let Y_1, \dots, 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 $\varepsilon \in (0, 1)$ and $\mu \geq \sum_{i=1}^n p_i$, we have $\Pr[\sum_{i=1}^n X_i > (1 + \varepsilon)\mu] \leq e^{-\varepsilon^2 \mu / 3}$.

Proof of Claim A.3. Fix $v \notin S$, and let X denote the random variable $|E^+(v, S)|$. Let S_1, \dots, 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' and E' to denote probabilities and expectations in the space where we condition 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 A.4 and infer that $\Pr'[X > 12 \log n] \leq e^{-\frac{6 \log n}{3}}$.

Finally, we have

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

By Claims A.2 and A.3, and since $|R| \geq c_3 \log n$, with high probability, for all but at most $c_3 \log n$ nodes $v \notin S$, we have

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

Hence, with high probability, at most $c_3 \log n$ nodes outside of S survive Step (2), i.e., $|S' \setminus S| \leq c_3 \log n$.

Claim A.5. *With high probability, we have $|E(v, S)| \leq 0.7|S|$ for all $v \notin S$.*

Proof. Since $12 \log n = o(|S|)$ (for sufficiently large n), by Claims A.1 and A.3, with probability, for all $v \notin S$, we have $|E(v, S)| = |E^-(v, S)| + |E^+(v, S)| \leq 0.6|S| + o(|S|) \leq 0.7|S|$. \blacksquare

By Claim A.5 and because $|S' \setminus S| \leq c_3 \log n$, with high probability, every node $v \in S' \setminus S$ has $|E(v, S')| \leq |E(v, S)| + c_3 \log n \leq 0.8|S|$. Therefore, no vertex $v \in S' \setminus S$ survives Step (3) and $\hat{S} = S$.