

On the Hardness of Signaling*

Shaddin Dughmi†

Department of Computer Science
 University of Southern California
 Los Angeles, California 90089-0781
 shaddin@usc.edu

Abstract—There has been a recent surge of interest in the role of information in strategic interactions. Much of this work seeks to understand how the realized equilibrium of a game is influenced by uncertainty in the environment and the information available to players in the game. Lurking beneath this literature is a fundamental, yet largely unexplored, algorithmic question: how should a “market maker” who is privy to additional information, and equipped with a specified objective, inform the players in the game? This is an informational analogue of the mechanism design question, and views the *information structure* of a game as a mathematical object to be designed, rather than an exogenous variable.

We initiate a complexity-theoretic examination of the design of optimal information structures in general Bayesian games, a task often referred to as *signaling*. We focus on one of the simplest instantiations of the signaling question: Bayesian zero-sum games, and a principal who must choose an information structure maximizing the equilibrium payoff of one of the players. In this setting, we show that optimal signaling is computationally intractable, and in some cases hard to approximate, assuming that it is hard to recover a planted clique from an Erdős-Rényi random graph. This is despite the fact that equilibria in these games are computable in polynomial time, and therefore suggests that the hardness of optimal signaling is a distinct phenomenon from the hardness of equilibrium computation.

Necessitated by the non-local nature of information structures, en-route to our results we prove an “amplification lemma” for the planted clique problem which may be of independent interest. Specifically, we show that even if we plant many cliques in an Erdős-Rényi random graph, so much so that most nodes in the graph are in some planted clique, recovering a constant fraction of the planted cliques is no easier than the traditional planted clique problem.

I. INTRODUCTION

Mechanism design is concerned with designing the rules of a game so as to effect desirable outcomes at equilibrium. This form of intervention, through the design of *incentives*, is fundamentally algorithmic in nature, and has led to a large body of work which examines the computational complexity of incentive-compatible mechanisms. This paper concerns a different, though arguably equally important, mode of intervention: through making available the right *information*.

A classic example illustrating the importance of information in games is Akerlof’s “market for lemons” [2]. Each

	Coop	Defect
Coop	$-1 + \theta$ $-1 + \theta$	0 $-5 + \theta$
Defect	$-5 + \theta$ 0	-4 -4

Figure 1. An incomplete information variant of the prisoners’ dilemma.

seller in this market is looking to sell a used car which is equally likely to be a “peach” (high quality) or a “lemon” (low quality). Prospective buyers value peaches at \$1000, and lemons at \$0, whereas sellers value keeping a peach at \$800 and a lemon at \$0. If both buyers and sellers are informed, then peaches are traded at some price between \$800 and \$1000, increasing the social welfare by \$200 per trade. In contrast, when only sellers are informed, each buyer would only be willing to pay his expected value of \$500 for a random car. Since this is less than a seller’s reservation value for a peach, the sellers of peaches are driven out of the market, and only lemons are ever traded.

The previous example illustrates that information deficiency can reduce the payoff of some or all players in a game. Somewhat counter-intuitively, revealing additional information can also degrade the payoffs of some or all players, and in fact optimal information structures may reveal some but not all the information available. For an illustrative example, consider the incomplete-information variant of the classical prisoners’ dilemma shown in Figure 1, in which the game’s payoffs are parametrized by a state of nature θ . When $\theta = 0$, this is the traditional prisoners’ dilemma in which cooperation is socially optimal, yet the unique Nash equilibrium is the one where both players defect, making both worse off. When, however, the reward from cooperation is uncertain, the equilibrium depends on players’ beliefs about θ .

Assuming that θ is uniformly distributed in $[-3, 3]$, and that players are risk-neutral and know nothing about θ besides its distribution, they play as if θ equals its expectation and the defection equilibrium persists. On the other hand, if both players learn the realization of θ before making their decisions, they would cooperate at equilibrium when $\theta \geq 1$, and defect otherwise, improving both their expected utilities.

*A full version of this paper is available at <http://arxiv.org/abs/1402.4194>.

†Supported in part by NSF CAREER Award CCF-1350900.

What is surprising, however, is that neither the opaque nor the transparent signaling scheme is optimal. Consider the following partially-informative scheme: when $\theta > -1$, both players receive the signal `High`, and otherwise receive the signal `Low`. On signal `High`, the posterior expected value of θ for both players is 1, inducing cooperation. On signal `Low`, the players defect. This information structure induces cooperation with greater probability than either of the opaque or transparent schemes, and thus improves the expected utility of both players.

Motivated by these intricacies, this paper begins a systematic complexity-theoretic examination of optimal information structure design in games, henceforth referred to as *signaling*. Our choice of backdrop is the simplest and most fundamental of all in game theory: 2-player zero-sum games. We consider incomplete-information zero-sum games, in which a state of nature θ determines the entries of the payoff matrix, and is drawn from a common-knowledge prior distribution λ represented explicitly. We then ask the following question: how should a principal privy to the realization of θ release information to players in order to effect a desired outcome at minimax equilibrium? Such a principal may be interested in maximizing some weighted combination of the players’ utilities, or more generally some (possibly linear) function of their chosen mixed strategies. Most such natural classes of objective functions include, as a special case, the task of maximizing one player’s utility at equilibrium. Therefore, we adopt that as our objective of choice for our hardness results.

Before describing our results and techniques, we elaborate on our modeling assumptions and chosen class of games. First we note that, like much of the prior work, we constrain ourselves to the design of *symmetric* signaling schemes: those which reveal the same information regarding θ to every player in the game. We believe such a restriction to a shared communication channel is natural and justified in many foreseeable applications. Moreover, the equilibria of a Bayesian game when agents have symmetric beliefs are simpler to characterize, better understood, and more often admits a canonical choice for equilibrium selection.¹ That being said, the complexity of *asymmetric* signaling schemes is an exciting avenue for future work.²

Second, we note that in addition to being simple and fundamental, zero-sum games admit a canonical and *tractable* choice of equilibrium — the minimax equilibrium. In addition to simplifying our analysis, this drives a larger conceptual point: the hardness of optimal signaling is a distinct phenomenon from the hardness of equilibrium computation.

¹e.g., as apparent in [1], auctions with information asymmetries suffer from the multiplicity of equilibria, and their analysis requires nontrivial refinements of standard equilibrium concepts.

²Asymmetric information introduces interdependencies and correlations in agents’ posterior estimates of the game’s payoffs. Asymmetric signaling appears related to optimization over the space of correlated equilibria. We refer the reader to the characterization of Bergemann and Morris [7].

Indeed, it is not hard to see that the optimal signaling task must, in general, “inherit” the computational complexity of the adopted equilibrium concept. Our choice of a class of games for which equilibrium can be easily computed essentially “disentangles” the complexity of optimal signaling from that of equilibrium computation.

Results and Techniques

Our results are based on the conjectured hardness of the planted clique problem. Specifically, we assume that it is hard to recover a planted k -clique from an Erdős-Rényi random graph $\mathcal{G}(n, p)$ for $p = \frac{1}{2}$ and some $k = k(n)$ satisfying $k = \omega(\log^2 n)$ and $k = o(\sqrt{n})$.

We prove two main results assuming this conjecture in Section IV. For explicitly-represented 2-player Bayesian zero-sum games, we prove that no algorithm computes a player-optimal signaling scheme in polynomial time. We strengthen this to a hardness of approximation result, for an additive absolute constant, for an implicitly-represented zero-sum game which nevertheless permits efficient equilibrium computation.

Our hardness results hold for a Bayesian zero-sum game in which the states of nature, as well as each player’s pure strategies, are the nodes V of a graph G . Specifically, we construct such a game in which signaling schemes correspond to “fractional partitions” of the states of nature V . Roughly speaking, the first player’s utility function favors signaling schemes which (fractionally) partition the graph into large yet highly connected clusters of nodes. If G is a random graph in which many k -cliques have been planted randomly throughout, so much so that the planted cliques roughly *cover* the nodes of G , then the best signaling scheme roughly corresponds to a partition of G into the planted cliques. Moreover, we exhibit an algorithm which, given a near-optimal signaling scheme, recovers a constant fraction of the planted clique cover.

In order to base our results on the planted clique conjecture, in Section III we prove an “amplification lemma” extending the hardness of recovering a planted clique to recovering a constant fraction of a planted clique cover. This lemma, which may be of independent interest, appears surprising because its analogue for the “distinguishing version” of the planted cover and planted clique problems is false.

Additional Related Work

The study of the effects of information on strategic interactions, and mechanisms for signaling, has its roots in the early works of Akerlof [2] and Spence [32]. Hirshleifer [22] was the first to observe that more information sometimes leads to worse market outcomes, in contrast to earlier work by Blackwell [9] which implied that more information is always better for a single agent in a non-competitive environment. Since then, many works have examined the effects of additional information on players’ utilities. Lehrer

et al. [28] showed that information improves players' utilities in *common interest* games — games where players have identical payoffs in each outcome, and Bassan et al. [6] exhibited a polyhedral characterization of games in which more information improves the individual utility of *every* player. The work of Peski [31] is related to ours, and considers *asymmetric* information structures in zero-sum games. They show that maximizing one player's utility requires revealing as much information as possible to that player, and withholding as much as possible from his opponent.

Many recent works have established striking characterizations of the space of equilibria attainable by signaling, in natural game domains. Bergemann et al. [8] considers signaling in a buyer/seller price-discrimination game, and provides a polyhedral characterization of the space of potential equilibria and their associated payoffs. Syrgkanis et al. [33] characterize the equilibrium of a 2-player common value auction supplemented with a signaling scheme. Bergemann and Morris [7] characterize the space of equilibria attainable by (asymmetric) information structures in general games, and relate it to the space of correlated equilibria.

Despite appreciation of the importance of information in strategic interactions, it is only recently that researchers have started viewing the information structure of a game as a mathematical object to be designed, rather than merely an exogenous variable. Kamenica and Gentzkow [26] examine settings in which a sender must design a signaling scheme to convince a less informed receiver to take a desired action. Recent work in the CS community, including by Emek et al. [14], Bro Miltersen and Sheffet [10] and Guo and Deligkas [20], examines revenue-optimal signaling in an auction setting, and presents polynomial-time algorithms and hardness results for computing it. Dughmi et al. [13] examine welfare-optimal signaling in an auction setting under exogenous constraints, and presents polynomial-time algorithms and hardness results.

II. PRELIMINARIES

A. Games

At its most general, a *Bayesian game* is a family of *games of complete formation* parametrized by a state of nature θ , where θ is assumed to be drawn from a known prior distribution. In this paper, we focus on *finite, Bayesian, 2-player zero-sum games*, each of which is described by the following parameters: Nonnegative integers r and c , denoting the number of pure strategies of the row player and column player respectively; a finite family $\Theta = \{1, \dots, M\}$ of *states of nature*, which we index by θ ; a family of *payoff matrices* $\mathcal{A}^\theta \in \mathbb{R}^{r \times c}$, indexed by states of nature $\theta \in \Theta$; and a *prior distribution* distribution $\lambda \in \Delta_M$ on the states of nature. Naturally, $\mathcal{A}^\theta(i, j)$ denotes the payoff of the row player when the row player plays i , the column player plays j , and the state of nature is θ . The column player's payoff in the same situation is $-\mathcal{A}^\theta(i, j)$.

We use two different representations of zero-sum games. In the *explicit representation*, the matrices $\{\mathcal{A}^\theta\}_\theta$, as well as the prior distribution λ , are given explicitly. We also relax this somewhat for our second result; specifically, we consider a game in which the individual matrices \mathcal{A}^θ are given *implicitly*, since one of the players has a number of strategies which is exponential in the natural description of the game. Nevertheless, a low-rank bilinear structure permits efficient computation of equilibria, and the value of the game, in the implicitly represented game we consider.³

B. Signaling Schemes

We examine policies whereby a principal reveals partial information regarding the state of nature θ to the players. Crucially, we require that the principal reveal the same information to both players in the game. A *symmetric signaling scheme* is given by a set Σ of *signals*, and a (possibly randomized) map φ from states of nature Θ to signals Σ . Abusing notation, we use $\varphi(\theta, \sigma)$ to denote the probability of announcing signal $\sigma \in \Sigma$ when the state of nature is $\theta \in \Theta$. We restrict attention to signaling schemes with a finite set of signals Σ , and this is without loss of generality when Θ is finite. We elaborate on this after describing the convex decomposition interpretation of a signaling scheme.

We note that signaling schemes are in one-to-one correspondence with *convex decompositions* of the prior $\lambda \in \Delta_M$ — namely, distributions supported on the simplex Δ_M , and having expectation λ . Formally, a signaling scheme $\varphi : \Theta \rightarrow \Sigma$ corresponds to the convex decomposition $\lambda = \sum_{\sigma \in \Sigma} \alpha_\sigma \cdot x_\sigma$, where $\alpha_\sigma = \Pr[\varphi(\theta) = \sigma] = \sum_{\theta \in \Theta} \lambda(\theta) \varphi(\theta, \sigma)$, and $x_\sigma(\theta) = \Pr[\theta | \varphi(\theta) = \sigma] = \frac{\lambda(\theta) \varphi(\theta, \sigma)}{\alpha_\sigma}$. Note that $x_\sigma \in \Delta_M$ is the *posterior distribution* of θ conditioned on signal σ , and α_σ is the probability of signal σ . The proof of the converse direction, namely that every convex decomposition of λ corresponds to a signaling scheme, is elementary yet thought provoking, and hence left to the reader.

We judge the quality of a signaling scheme by the outcome it induces signal by signal. Specifically, the principal is equipped with an objective function of the form $\sum_{\sigma} \alpha_\sigma \cdot f(x_\sigma)$, where $f : \Delta_M \rightarrow \mathbb{R}$ is some function mapping a posterior distribution to the quality of the equilibrium chosen by the players. For example, f may be the social welfare at the induced equilibrium, a weighted combination of players' utilities at equilibrium, or something else entirely. In this setup, one can show that there always exists a signaling scheme with a finite set of signals which maximizes our objective, so long as the states of nature are finitely many. The optimal choice of signaling scheme is related to the *concave envelope* f^+ of the function f .⁴ Specifically, such

³Implicitly-represented 2-player zero-sum games often admit efficient algorithms for equilibrium computation. Fairly general conditions under which this is possible are well exposted by Immorlica et al. [23].

⁴ f^+ is the point-wise lowest concave function h for which $h(x) \geq f(x)$ for all x in the domain. Equivalently, the hypograph of f^+ is the convex hull of the hypograph of f .

a signaling scheme achieves $\sum_{\sigma} \alpha_{\sigma} \cdot f(x_{\sigma}) = f^{+}(\lambda)$. Application of Caratheodory’s theorem to the hypograph of f , therefore, shows that $M + 1$ signals suffice.

Our impossibility results rule out algorithms for computing near optimal signaling schemes, almost agnostic to how such schemes are represented as output. For concreteness, the reader can think of a signaling scheme φ as represented by the matrix of pairwise probabilities $\varphi(\theta, \sigma)$. Since we only consider games where the states of nature, and therefore also the number of signals w.l.o.g., are polynomially many in the description size of the game, this is a compact representation. The representation of φ as a convex decomposition would do equally well, as both representations can be efficiently computed from each other.

C. Strategies, Equilibria, and Objectives

Given a Bayesian game, a symmetric signaling scheme (α, x) with signals Σ induces $|\Sigma|$ sub-games, one for each signal. The subgame corresponding to signal $\sigma \in \Sigma$ is played with probability α_{σ} , and players’ (common) beliefs regarding the state of nature in this subgame are given by the posterior distribution $x_{\sigma} \in \Delta_M$. The quality of a symmetric signaling scheme in such a game is contingent on a choice of an *equilibrium concept* and an *objective function*.

Equilibrium Concept: An equilibrium concept distinguishes a mixed strategy profile for every posterior belief $x \in \Delta_M$. This permits defining an objective function on signaling schemes, as described in Section II-B. In general Bayesian games, evaluating the quality of a signaling scheme may be complicated by issues of equilibrium selection — for example, general sum games often admit many Bayes-Nash equilibria. However, our restriction to two-player zero-sum games side-steps such complications: all standard equilibrium concepts are payoff-equivalent in zero-sum games, and correspond to the minimax equilibrium in each subgame.

Our restriction to zero-sum games avoids an additional complication, which arises for more general games. Equilibrium computation in full-information zero-sum games is *tractable*, permitting efficient computation of equilibrium in the Bayesian game, for every posterior distribution over the states of nature. This avoids “inheriting” the computational complexity of equilibrium computation into our optimal signaling problems, and therefore suggests that the hardness of optimal signaling is a distinct phenomenon from the complexity of computing equilibria.

Formally, given a Bayesian zero-sum game $(\{\mathcal{A}^{\theta}\}_{\theta=1}^M, \lambda)$ as described in Section II-A, and a signaling scheme corresponding to the convex decomposition (α, x) of λ , this induces a distribution over sub-games. Specifically, for each signal $\sigma \in \Sigma$, the Bayesian zero-sum game $(\{\mathcal{A}^{\theta}\}_{\theta=1}^M, x_{\sigma})$ is played with probability α_{σ} . We use $\mathcal{A}^{\sigma} = \mathbf{E}_{\theta \sim x_{\sigma}}[\mathcal{A}^{\theta}]$ to denote the matrix of posterior expected payoffs conditioned on signal σ . A *Bayesian Nash equilibrium* corresponds to an equilibrium of each sub-game $(\{\mathcal{A}^{\theta}\}_{\theta=1}^M, x_{\sigma})$ in which

players play as they would in the complete information zero-sum game \mathcal{A}^{σ} . In the sub-game corresponding to signal σ , the row player’s payoff is simply his payoff in the complete information zero-sum game \mathcal{A}^{σ} , namely $u_r(\mathcal{A}^{\sigma}) = \max_{y \in \Delta_r} \min_{j=1}^c (y^{\top} \mathcal{A}^{\sigma})_j$. The expected payoff of the row player over the entire game is then given by $u_r(\alpha, x) = \sum_{\sigma \in \Sigma} \alpha_{\sigma} u_r(\mathcal{A}^{\sigma})$.

Objective Function: We adopt u_r , as given above, as our objective function. By symmetry, this is technically equivalent to adopting objective function u_c , the utility of the column player. To justify this choice for our hardness results, observe that most natural classes of objective functions — say weighted combination of players’ utilities, or linear functions of players’ mixed strategies — include u_r and u_c as special cases.

Additionally, we are motivated by the fact that the space of all payoff profiles achievable by signaling is spanned by the two schemes maximizing the utility of one of the players, in the following sense. Let φ_r be a signaling scheme maximizing the row player’s utility, and φ_c be a signaling scheme maximizing the column player’s utility. Moreover, let $u_r^{\max} = -u_c^{\min}$ denote the row player’s maximum utility (equivalently, the negation of the column player’s minimum utility), and similarly let $u_c^{\max} = -u_r^{\min}$ denote the column player’s maximum utility (equivalently, the negation of the row player’s minimum utility). If there exists a signaling scheme φ inducing a utility profile (u_r, u_c) , then it can be formed as a convex combination of the two extreme schemes φ_r and φ_c — specifically, by running φ_r with probability $\frac{u_r - u_r^{\min}}{u_r^{\max} - u_r^{\min}}$ and φ_c otherwise. Consequently, “mapping out” the space of possible payoff profiles, as well constructing a signaling scheme attaining a desired payoff profile, both reduce to computing the extreme schemes φ_r and φ_c .

D. Graphs, Clusters, and Density

An *undirected graph* G is a pair (V, E) , where V is a finite set of *nodes* or *vertices*, and $E \subseteq \binom{V}{2}$ is a set of undirected *edges*. We usually use n to denote the number of vertices of a graph, and m to denote the number of edges. Given a graph $G = (V, E)$, a *cluster* is some $S \subseteq V$. Given a cluster S , we define the intra-cluster edges $E(S)$ as those edges with both endpoints in S . The *induced subgraph* of S is the graph $H = (S, E(S))$. Moreover, given two clusters $S, T \subseteq V$, we define the inter-cluster edges $E(S, T)$ as the edges with at least one endpoint in each of S and T .

The *density* of G is the fraction of all potential edges in E — namely $\text{density}(G) = \frac{2|E|}{|V|(|V|-1)}$. More generally, the density of a cluster S in G is the density of the subgraph of G induced by S , specifically $\text{density}_G(S) = \frac{2|E(S)|}{|S|(|S|-1)}$. To precisely state and prove our results, we require a slightly different notion of density, defined between pairs of clusters. We define the *bi-density* between clusters S and T as the fraction of all pairs $(u, v) \in S \times T$ connected by an edge; formally,

bi-density(S, T) = $\frac{1}{|S||T|} \cdot |\{(u, v) \in S \times T \mid \{u, v\} \in E\}|$. Equivalently, the bi-density between S and T can be thought of in terms of the adjacency matrix of A of G ; specifically, $\text{bi-density}(S, T) = \frac{1}{|S||T|} \sum_{i \in S} \sum_{j \in T} A_{ij}$. Observe that density and bidensity are closely related, in that $\text{bi-density}(S, S) = \text{density}(S)(1 - \frac{1}{|S|})$.

E. Random Graphs

We make use of *Erdős-Rényi random graphs*. Given $n \in \mathbb{N}$ and $p \in [0, 1]$, the random graph $\mathcal{G}(n, p)$ has vertices $V = \{1, \dots, n\}$, and its edges include every $e \in \binom{V}{2}$ independently with probability p . Naturally, for $G \sim \mathcal{G}(n, p)$, every cluster S of G has expected density p . Moreover, every pair of clusters S and T has expected bi-density $p \left(1 - \frac{|S \cap T|}{|S||T|}\right)$.⁵ We make use of the following standard probabilistic bound, proved in the full version.

Proposition II.1. *Let $p \in (0, 1)$ and $\alpha > 1$ be absolute constants (independent of n), and let $G \sim \mathcal{G}(n, p)$. There is an absolute constant $\beta = \beta(p, \alpha)$ such that, with high probability, $\text{bi-density}_G(X, Y) \leq \alpha p$ simultaneously for all clusters X and Y with $|X|, |Y| > \beta \log n$.*

As in the above proposition, we make references throughout this paper to guarantees which hold *with high probability* over a family of graphs parametrized by the number of nodes n . By this we mean that the claimed property holds with probability at least $1 - 1/\text{poly}(n)$, for some polynomial in the number of nodes.

III. PLANTING CLIQUES AND CLIQUE COVERS

A. The Planted Clique Problem

For our hardness results, we assume that it is hard to recover a planted clique from an Erdős-Rényi random graph. Specifically, we consider the following problem for parameters $n, k \in \mathbb{N}$ and $p \in [0, 1]$.

Definition III.1 (The Planted Clique Problem $\text{PCLIQUE}(n, p, k)$). *Let $G \sim \mathcal{G}(n, p, k)$ be a random graph with vertices $[n] = \{1, \dots, n\}$, constructed as follows: (1) Every edge is included in G independently with probability p ; (2) A set $S \subseteq [n]$ with $|S| = k$ is chosen uniformly at random; (3) All edges with both endpoints in S are added to G . Given a sample from $G \sim \mathcal{G}(n, p, k)$, recover the planted clique S .*

We refer to the parameter p as the *background density*, S as the *planted clique*, and k as the *size* of the clique. Typically, we will think of k as a function of n , and p as an absolute constant independent of n . Our results will hinge on the conjectured hardness of recovering the clique with constant probability. We use the following conjecture.

Assumption III.2. *For some function $k = k(n)$ satisfying $k = \omega(\log^2 n)$ and $k = o(\sqrt{n})$, there is no probabilistic*

⁵The discrepancy from p is due to the absence of self-loops.

polynomial-time algorithm for $\text{PCLIQUE}(n, \frac{1}{2}, k)$ with constant success probability.

By *constant success probability*, we mean that the probability of the algorithm recovering the clique is bounded below by a constant independent of n , over the random draw of the graph $G \sim \mathcal{G}(n, p, k)$ as well as the internal random coins of the algorithm.

The problem of recovering a planted clique, as well as the (no harder) problem of distinguishing a draw from $\mathcal{G}(n, p)$ from a draw from $\mathcal{G}(n, p, k)$, has been the subject of much work since it was introduced by Jerrum [24] and Kučera [27]. Almost all of this work has focused on $p = \frac{1}{2}$. On the positive side, there is a quasipolynomial time algorithm for recovering the clique when $k \geq 2 \log n$. The best known polynomial-time algorithms, on the other hand, can recover planted cliques of size $\Omega(\sqrt{n})$, through a variety of different algorithmic techniques (e.g. [3, 29, 11, 15, 5, 17, 12]).

Despite this extensive body of work, there are no known polynomial-time algorithms for recovering, or even detecting, planted cliques of size $k = o(\sqrt{n})$. There is evidence the problem is hard: Jerrum [24] ruled out Markov chain approach for $k = o(\sqrt{n})$; Feige and Krauthgamer [16] ruled out algorithms based on the Lovász-Schrijver family of semi-definite programming relaxations for $k = o(\sqrt{n})$; and recently Feldman et al. [18] defined a family of “statistical algorithms,” and ruled out such algorithms for recovering planted cliques of size $k = o(\sqrt{n})$. Consequently, the planted clique conjecture has been used as a hardness assumption in a variety of contexts (e.g. [4, 25, 21, 30]).

B. From Planting Cliques to Planting Covers

For our results, we construct games in which the random state of nature is a node in some graph G . Recalling that a signaling scheme is a kind of fractional partition of the states of nature, with each “fractional part” corresponding to a signal, we design our games so that the quality of a signal is proportional to the density of that part of the graph. Since a signaling scheme’s quality is measured in aggregate over the entire distribution over states of nature, any hardness result must rule out recovering a family of dense clusters (one per signal) scattered throughout the graph, rather than merely a unique dense cluster as in the planted clique problem. In this section, we define the *planted clique-cover problem* in which many cliques are planted throughout the graph, and prove a kind “amplification lemma” extending the hardness of recovering a planted clique to recovering a constant fraction of the planted cover.

The planted clique-cover problem is parametrized by $n, k, r \in \mathbb{N}$ and $p \in [0, 1]$, and defined as follows.

Definition III.3 (The Planted Clique Cover Problem $\text{PCCOVER}(n, p, k, r)$). *Let $G \sim \mathcal{G}(n, p, k, r)$ be a random graph on vertices $[n] = \{1, \dots, n\}$, constructed as follows: (1) Include every edge in G independently with probability*

p ; (2) For $i = 1$ to r , choose $S_i \subseteq [n]$ with $|S_i| = k$ uniformly at random, and add all edges with both endpoints in S_i to G . Given a sample from $G \sim \mathcal{G}(n, p, k, r)$, recover a constant fraction of the planted cliques S_1, \dots, S_r .⁶

As in the planted clique problem, we refer to the parameter p as the *background density*, S_1, \dots, S_r as the *planted cliques*, k as the *size* of each clique, and r as the *number* of cliques. Moreover, we will think of k and r as functions of n , and p as an absolute constant independent of n . Note that the planted clique problem is the special case of the planted clique-cover problem when $r = 1$. For our results, we will use the planted clique-cover problem for $r = \Theta(\frac{n}{k})$, guaranteeing that a constant fraction of the nodes are in at least one of the planted cliques with high probability.

At first glance, it might appear that planting many cliques as we do here makes the problem easier than planted clique. Somewhat surprisingly — and we elaborate on why later — this is not the case. We exhibit a reduction from the planted clique problem to the planted clique-cover problem with an arbitrary parameter r , showing that indeed the planted clique-cover problem is no easier. We prove the following.

Lemma III.4. *Let k and r be arbitrary functions of n , and p be an arbitrary constant. If there is a probabilistic polynomial-time algorithm for $\text{PCCOVER}(n, p, k, r)$ with constant success probability, then there is such an algorithm for $\text{PCLIQUE}(n, p, k)$.*

Proof: The reduction proceeds as follows: Given a graph $G \sim \mathcal{G}(n, p, k) = \mathcal{G}(n, p, k, 1)$, we construct a graph $G' \sim \mathcal{G}(n, p, k, r)$ by planting $r - 1$ additional k -cliques at random, as in Definition III.3. In a sense, we “continue where planted clique left off” by adding $r - 1$ more cliques placed randomly in the graph. Let S_1 denote the original planted clique, and S_2, \dots, S_r be the “additional” cliques planted through the reduction.

The key observation is that the cliques S_1, \dots, S_r are indistinguishable to any algorithm operating on a sample from $\mathcal{G}(n, p, k, r)$. This is by a symmetry argument: permuting the order in which cliques are planted does not change the distribution $\mathcal{G}(n, p, k, r)$. As a result, any algorithm which recovers a constant fraction of the planted cliques from G' with constant probability must recover each of S_1, \dots, S_r with constant probability. In particular, applying an algorithm for the planted cover problem to the outcome of our reduction yields a list of cliques which includes S_1 — the original planted clique — with constant probability. This leads to an algorithm for the planted clique problem with constant success probability. ■

Lemma III.4 is surprising since it does not appear to hold for the *distinguishing* variants of PCCOVER and

⁶By this we mean that the algorithm should output a list of k -cliques in G , at least αr of which are in $\{S_1, \dots, S_r\}$, for some constant α independent of n .

PCLIQUE. Specifically, when $k = n^\gamma$ for a constant $\gamma < \frac{1}{2}$, whereas it is conjectured that no algorithm can distinguish a sample from $\mathcal{G}(n, \frac{1}{2})$ from a sample from $\mathcal{G}(n, \frac{1}{2}, k)$ with constant probability, a simple statistical test succeeds in distinguishing a sample from $\mathcal{G}(n, \frac{1}{2})$ from a sample from $\mathcal{G}(n, \frac{1}{2}, k, n/k)$ with constant probability. The test in question simply counts the edges in the graph. The random graph $\mathcal{G}(n, \frac{1}{2})$ has $\frac{n^2}{4} \pm O(n)$ edges with probability approaching 1, whereas $\mathcal{G}(n, \frac{1}{2}, k, n/k)$ has $\frac{n^2}{4} + \Omega(\frac{n}{k} \cdot k^2) = \frac{n^2}{4} + \Omega(n^{1+\gamma})$ edges in expectation — well above the confidence interval for the number of edges in $\mathcal{G}(n, \frac{1}{2})$.

C. Approximate Recovery

To simplify our hardness results, we show that producing a set of vertices with sufficient overlap with a planted clique is polynomial-time equivalent to recovering that entire clique, with high probability.

Lemma III.5. *Let $\epsilon > 0$ and $p \leq \frac{1}{2}$ be constants, $k = k(n)$ satisfy $k = \omega(\log^2 n)$ and $k = o(\sqrt{n})$, and $r = O(\frac{n}{k})$. There is a probabilistic polynomial-time algorithm which takes as input $G \sim \mathcal{G}(n, p, k, r)$ and a cluster $T \subseteq [n]$, and outputs every planted k -clique S in G for which $|T \cap S| > \epsilon|T \cup S|$, with high probability.*

In other words, to recover a planted k -clique S it suffices to produce a set T of size $O(k)$ for which $|T \cap S| = \Omega(k)$. The analogous statement for the (single) planted clique problem is folklore. Our proof is similar, though requires some additional accounting because our cliques may overlap. We relegate the proof of Lemma III.5 to the full version.

IV. ZERO-SUM GAMES

In this section, we show the intractability of optimal signaling in zero sum games, when the objective is maximizing one player’s payoff at equilibrium. We prove two results which follow from Assumption III.2: Hardness of optimal signaling for *explicit* zero sum games — those games with a polynomial number of strategies and matrices given explicitly, and hardness of approximation for *implicit* zero-sum games. Our former result, for explicit games, requires the real numbers in the payoff matrices to scale with the input size, and therefore does not qualify as a hardness of approximation result per se. Our latter result for implicitly-described games does rule out a constant additive approximation relative to range of possible payoffs, and crucially holds for a game in which equilibrium computation is *tractable*, in the sense described in Section II-C.

Theorem IV.1. *Assumption III.2 implies that there is no polynomial-time algorithm which computes a signaling scheme maximizing a player’s payoff in an explicitly represented 2-player Bayesian zero-sum game.*

Theorem IV.2. *Assumption III.2 implies there is an implicitly-described, yet tractable, Bayesian zero-sum game*

with payoffs in $[-1, 1]$, and an absolute constant ϵ , such that there is no polynomial-time algorithm for computing a signaling scheme which ϵ -approximately maximizes (in the additive sense) one player’s payoff.

For both our results, we use variants of the following security game. In a security game, an *attacker* chooses a target to attack, and a *defender* chooses targets to defend. The payoff of both players depends on the choice of target and the extent to which it is protected. We consider a zero-sum security game on a network, in which an attacker is looking to take down edges of the network, and a defender is looking to defend those edges. The state of nature determines a vulnerable node θ , and an attacker can only take down edges adjacent to θ by attacking its other endpoint. The defender, on the other hand, can select a small set of nodes and protect the edges incident on them from attack.

Definition IV.3 (Network Security Game). *Instances of this Bayesian zero-sum security game are described by an undirected graph $G = (V, E)$ with n nodes and m edges, representing a communication network, an integer $d \geq 1$ equal to the number of nodes the defender can simultaneously protect, and a real number $\rho \geq 0$ equal to the utility gain to the defender for protecting a vulnerable or attacked node. States of nature correspond to vertices V . When the realized state of nature is $\theta = v$, this indicates that the edges incident on node v are vulnerable to attack. We assume that $\theta \sim \lambda$, where the prior λ is the uniform distribution over V . The attacker’s strategies also correspond to the nodes V . The defender’s strategies correspond to subsets of V of size at most d , representing the choice of nodes to defend. When the state of nature is $\theta \in V$, the attacker attacks $a \in V$, and the defender defends $D \subseteq V$ with $|D| \leq d$, the payoff of the attacker is defined as follows:*

$$\mathcal{A}^\theta(a, D) = |\{(\theta, a)\} \cap E| - \rho|D \cap \{\theta, a\}|$$

When d is a constant, the size of the matrices \mathcal{A}^θ is polynomial in the representation size of the game, and we can think of the game as being represented *explicitly* as a set of n matrices $\mathcal{A}^\theta \in [-2\rho, 1]^{n \times \binom{n}{d}}$. However, when d is superconstant, the game matrices have a superpolynomial $\binom{n}{d}$ number of columns. Nevertheless, even then equilibrium computation is still *tractable* in the sense described in Section II-C — i.e., our representation permits computation of the equilibrium, and corresponding utilities, for every posterior distribution $x \in \Delta_n$ over states of nature. This is a consequence of the “low dimensional” nature of the defender’s mixed strategy space, and linearity of players’ utilities in this low-dimensional representation. Specifically, a mixed strategy of the defender can be summarized as a vector in the matroid polytope $\mathcal{P}_d = \{z \in [0, 1]^n : \sum_{i=1}^n z_i \leq d\}$. A vector $z \in \mathcal{P}_d$ encodes the probability by which the defender defends each target. It is not hard to see that this representation is loss-less, in that (a) the vector z

summarizing a defender’s mixed strategy, together with the attacker’s mixed strategy $y \in \Delta_n$, suffices for computing the payoff of each player in a subgame with beliefs x — specifically, using A to denote the adjacency matrix of graph G , the attacker’s payoff is $x^\top A y - \rho(z^\top x + z^\top y)$; (b) Given a vector $z \in \mathcal{P}_d$, a corresponding mixed strategy of the defender with small support can be efficiently recovered as the convex decomposition of z into the corner points of the matroid polytope \mathcal{P}_d ; ⁷ and (c) Given a vector $z \in \mathcal{P}_d$ summarizing the defender’s mixed strategy in a subgame with beliefs x , a best response for the attacker can be computed efficiently.

Theorems IV.1 and IV.2 follow from the two lemmas below, as well as Lemma III.4. Specifically, Theorem IV.1 instantiates both lemmas below with $d = 1$ and $\rho = \frac{k}{150}$, and Theorem IV.2 sets $d = \frac{k}{150}$ and $\rho = 1$.

Lemma IV.4. *Let k and r be such that $r = \frac{3n}{k}$. For a network security game on $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ with $d, \rho \geq 1$ satisfying $\rho d = \frac{k}{150}$, with high probability there is a signaling scheme attaining expected attacker utility at least 0.8.*

Lemma IV.5. *Let $\epsilon > 0$ be an absolute constant, k satisfy $k = \omega(\log^2 n)$ and $k = o(\sqrt{n})$, and $r = \Theta(\frac{n}{k})$. For a network security game on $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ with $d, \rho \geq 1$ satisfying $\rho d = \Theta(k)$, there is a probabilistic polynomial-time algorithm which, given any signaling scheme obtaining attacker utility at least $0.5 + \epsilon$, outputs an $\Omega(\epsilon)$ fraction of the planted cliques in G , with high probability.*

Lemma IV.4 follows from analyzing a deterministic signaling scheme which groups together states of nature in the same planted clique (breaking ties arbitrarily). The proof is simple, and is relegated to the full version. We prove Lemma IV.5 next.

Proof of Lemma IV.5: We exhibit an algorithm which, given as input a graph $G \sim \mathcal{G}(n, \frac{1}{2}, k, r)$ and a signaling scheme φ attaining the claimed attacker utility, outputs a family of clusters \mathcal{T} from which a constant fraction of the planted cliques can be recovered in polynomial time. Specifically, we show that for a constant fraction of the planted cliques S_1, \dots, S_r , there is a cluster $T \in \mathcal{T}$ which satisfies the requirements of Lemma III.5. To simplify our proof, we assume that $\rho d/2$ is an integer, though unsurprisingly this assumption can easily be removed.

Consider Algorithm 1. In step (2), the algorithm “zeroes-out” those vertices which are “overrepresented” in either the attacker’s strategy or the posterior distribution over states of nature. This is justified because, as will become clear in the proof, we designed the game’s payoffs so that overrepresented nodes can easily be protected by the defender, and hence account for only a small fraction of the attacker’s total utility. After zeroing-out large entries, the posteriors \hat{x}_σ and

⁷This is the constructive version of Caratheodory’s theorem, as shown by Grötschel et al. [19].

Algorithm 1 Algorithm for Computing Approximation of Planted Cliques from a Signaling Scheme

Input: Graph $G = (V, E)$ with adjacency matrix $A \in \mathbb{R}^{n \times n}$

Input: A signaling scheme $\varphi : V \rightarrow \Sigma$, represented in convex decomposition form (α, x) where $\alpha \in \Delta_\Sigma$, and $x_\sigma \in \Delta_V$ for each $\sigma \in \Sigma$.

Output: A list \mathcal{T} of subsets of V , each of size $\rho d/2$.

- 1: For each $\sigma \in \Sigma$, compute the attacker's minimax strategy $y_\sigma \in \Delta_V$ in the subgame corresponding to signal σ .
- 2: For each $\sigma \in \Sigma$, discard from each of x_σ and y_σ all entries greater than $\frac{2}{\rho d}$. Formally, let $\hat{x}_\sigma(v)$ equal $x_\sigma(v)$ if $x_\sigma(v) \leq \frac{2}{\rho d}$, and equal 0 otherwise. $\hat{y}_\sigma(v)$ is defined analogously.
- 3: For each $\sigma \in \Sigma$, let $z_\sigma \in \Delta_V$ be an extreme-point solution of the following linear program. (Note that \hat{x}_σ is fixed)

$$\begin{aligned} & \text{maximize} && (\hat{x}_\sigma)^\top A z_\sigma \\ & \text{subject to} && \|z_\sigma\|_\infty \leq \frac{2}{\rho d} \\ & && \sum_{v \in V} z_\sigma(v) \leq 1 \\ & && z_\sigma(v) \geq 0, \quad \text{for } v \in V. \end{aligned} \quad (1)$$

- 4: Let T_σ be the support of $z_\sigma \in \Delta_V$, for each $\sigma \in \Sigma$.
 - 5: Output $\mathcal{T} = \{T_\sigma : \sigma \in \Sigma\}$.
-

attacker strategies \hat{y}_σ have entries no larger than $1/\Theta(k)$, and therefore behave roughly as uniform distributions over $\Theta(k)$ vertices. Steps (3) and (4) are ‘‘cleanup steps’’, which convert the attacker's mixed strategy to a uniform distribution over $\frac{\rho d}{2} = \Theta(k)$ vertices T_σ . We will show that the resulting family \mathcal{T} of clusters must overlap substantially with the planted cliques, in the sense of Lemma III.5, for the utility of the attacker to exceed the background density of $\frac{1}{2}$ by a constant. We formalize all this through a sequence of propositions.

Proposition IV.6. *The attacker's utility $u(\varphi)$ from signaling scheme $\varphi = (\alpha, x)$ satisfies $u(\varphi) \leq \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A \hat{y}_\sigma$, where A denotes the adjacency matrix of G , and \hat{x}_σ and \hat{y}_σ are as defined in step (2) of Algorithm 1.*

Proof: We prove this bound signal-by-signal. Fix a signal $\sigma \in \Sigma$ with posterior x_σ and attacker strategy y_σ , and as shorthand use $x = x_\sigma$, $y = y_\sigma$, $\hat{x} = \hat{x}_\sigma$, and $\hat{y} = \hat{y}_\sigma$. We distinguish nodes which are *overrepresented* in x or y — specifically, the nodes $O = \left\{v : \max(x(v), y(v)) \geq \frac{2}{\rho d}\right\}$. There are at most ρd such nodes, since each of x and y , by virtue of being a probability distribution, can have at most $\rho d/2$ entries exceeding $2/\rho d$.

We consider a defender who chooses $D \subseteq O$ with $|D| = \min(d, |O|)$ uniformly at random. Observe that such a defender protects each node in O with probability at least

$\frac{\min(d, |O|)}{|O|} \geq \frac{d}{\rho d} = \frac{1}{\rho}$. The utility of the attacker is equal to $x^\top A y - \rho(\Pr[\theta \in D] + \Pr[a \in D])$. Note that $\Pr[\theta \in D] = \Pr[\theta \in O] \Pr[\theta \in D | \theta \in O] \geq \frac{1}{\rho} \Pr[\theta \in O] = \frac{x(O)}{\rho}$. Similarly $\Pr[a \in D] \geq \frac{y(O)}{\rho}$. Therefore we have that

$$u^\sigma(\varphi) \leq x^\top A y - (x(O) + y(O)) \quad (2)$$

Next, we let $\bar{x} = x - \hat{x}$ and $\bar{y} = y - \hat{y}$, and bound $x^\top A y$.

$$\begin{aligned} x^\top A y &= \hat{x}^\top A \hat{y} + \hat{x}^\top A \bar{y} + \bar{x}^\top A \hat{y} + \bar{x}^\top A \bar{y} \\ &\leq \hat{x}^\top A \hat{y} + \hat{x}^\top \bar{y} + \bar{x}^\top \hat{y} + \bar{x}^\top \bar{y} \\ &\leq \hat{x}^\top A \hat{y} + \bar{x}^\top y + \bar{y}^\top x \\ &\leq \hat{x}^\top A \hat{y} + \|\bar{x}\|_1 \|y\|_\infty + \|\bar{y}\|_1 \|x\|_\infty \quad (\text{Holder's ineq.}) \\ &\leq \hat{x}^\top A \hat{y} + \|\bar{x}\|_1 + \|\bar{y}\|_1 \\ &= \hat{x}^\top A \hat{y} + x(O) + y(O) \end{aligned} \quad (3)$$

Combining (2) with (3) completes the proof. \blacksquare

Proposition IV.7. $u(\varphi) \leq \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A z_\sigma$

Proof: This follows from Proposition IV.6 and the fact that \hat{y}_σ is a feasible solution to linear program (1). \blacksquare

Proposition IV.8. z_σ is a uniform distribution over $\rho d/2$ vertices, namely T_σ .

Proof: Recall that we assumed that $\rho d/2$ is an integer. Therefore, a simple argument shows that an extreme-point solution to LP (1) must set each variable either to 0 or to $\frac{2}{\rho d}$. In particular, the optimal solution must set precisely $\rho d/2$ entries of z_σ to $\frac{2}{\rho d}$, as needed. \blacksquare

We now upperbound the contribution of edges outside the planted cliques to the utility of the attacker. Write the adjacency matrix A as $A = A^- + A^+$, where A^- are the background edges added in Step (1) of Definition III.3, and A^+ are the clique edges added in Step (2) of Definition III.3.

Proposition IV.9. $(\hat{x}_\sigma)^\top A^- z_\sigma \leq 0.5 + \frac{\epsilon}{2}$ simultaneously for all signals $\sigma \in \Sigma$, with high probability.

Proof: Pick an arbitrary signal σ . Since $\|\hat{x}_\sigma\|_\infty \leq \frac{2}{\rho d}$, we have that $(\hat{x}_\sigma)^\top A^- z_\sigma$ is bounded from above by the value of the linear program which holds z_σ fixed, and maximizes $x^\top A^- z_\sigma$ over $x \in \Delta_V$ with $\|x\|_\infty \leq \frac{2}{\rho d}$. By an argument similar to in Proposition IV.8, at optimality x is a uniform distribution over some $\rho d/2$ vertices R . Therefore, since z_σ is also a uniform distribution over the $\rho d/2$ vertices T_σ (Proposition IV.8), the value of our linear program is equal to bi-density (R, T_σ) . Since $\rho d/2 = \Theta(k) = \omega(\log n)$, Proposition II.1 implies the claimed bound. \blacksquare

We can now wrap up the proof of Lemma IV.5. We will show that, on average over our planted cliques S_1, \dots, S_r , there is some $T \in \mathcal{T}$ output by Algorithm 1 overlapping with a constant fraction of the clique vertices. As notation, we let A^i denote the adjacency matrix of the clique S_i , for $i = 1, \dots, r$. Note that $A^+ \leq \sum_{i=1}^r A^i$. First, we show

that foreground edges contribute $\Omega(\epsilon)$ to the outerproduct $(\hat{x}_\sigma)^\top A z_\sigma$, on average over signals $\sigma \in \Sigma$.

$$\begin{aligned}
0.5 + \epsilon &\leq \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A z_\sigma && \text{(Prop. IV.7)} \\
&= \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top (A^- + A^+) z_\sigma \\
&\leq (0.5 + \frac{\epsilon}{2}) \sum_{\sigma \in \Sigma} \alpha_\sigma + \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A^+ z_\sigma && \text{(Prop. IV.9)} \\
&= 0.5 + \frac{\epsilon}{2} + \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A^+ z_\sigma
\end{aligned}$$

Therefore $\sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A^+ z_\sigma \geq \frac{\epsilon}{2}$. Next, we “break up” this sum of outerproducts into the constituent contributions of each planted clique.

$$\begin{aligned}
\frac{\epsilon}{2} &\leq \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A^+ z_\sigma \\
&\leq \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top \left(\sum_{i=1}^r A^i \right) z_\sigma \\
&= \sum_{i=1}^r \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma^\top A^i z_\sigma \\
&\leq \sum_{i=1}^r \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma(S_i) z_\sigma(S_i) \\
&= \sum_{i=1}^r \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma(S_i) \frac{|T_\sigma \cap S_i|}{|T_\sigma|}
\end{aligned}$$

Finally, we show that the average planted clique is well represented by some $T \in \mathcal{T}$.

$$\begin{aligned}
\frac{\epsilon}{2} &\leq \sum_{i=1}^r \sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma(S_i) \frac{|T_\sigma \cap S_i|}{|T_\sigma|} \\
&\leq \sum_{i=1}^r \left(\sum_{\sigma \in \Sigma} \alpha_\sigma \hat{x}_\sigma(S_i) \right) \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T_\sigma|} \right) \\
&\leq \sum_{i=1}^r \frac{|S_i|}{n} \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T_\sigma|} \right) \\
&= \frac{k}{n} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T_\sigma|} \right) \\
&= O(1) \cdot \frac{1}{r} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{|T_\sigma|} \right) \\
&= O(1) \cdot \frac{1}{r} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{k} \right)
\end{aligned}$$

The last inequality follows from $\sum_{\sigma} \alpha_\sigma \hat{x}_\sigma(v) \leq \sum_{\sigma} \alpha_\sigma x_\sigma(v) = \lambda_v = \frac{1}{n}$; the next to last equality from $r = \Theta(\frac{n}{k})$; and the last equality from $|T_\sigma| = \rho d/2 = \Theta(k)$.

Therefore $\frac{1}{r} \sum_{i=1}^r \left(\max_{T \in \mathcal{T}} \frac{|T \cap S_i|}{k} \right) \geq \Omega(\epsilon)$; i.e., the average planted clique intersects at least one of the sets in \mathcal{T} in a constant fraction $\Omega(\epsilon)$ of its vertices. This implies that an $\Omega(\epsilon)$ fraction of the planted cliques intersect at least one

of the sets in \mathcal{T} in $\Omega(\epsilon)$ fraction of its vertices, by a simple counting argument. Since each set in \mathcal{T} is of size $\Theta(k)$, this satisfies the requirements of Lemma III.5 for a constant fraction of the planted cliques, completing the proof. ■

V. CONCLUSION

Our results raise several open questions, and we close with some of them. First, note that our impossibility result for explicit zero sum games rules out approximation algorithms with additive error on the order of $\frac{1}{\text{polylog } n}$ relative to the range of player payoffs, where n is the number of a player strategies. Essentially, this rules out a *fully polynomial time approximation scheme* (FPTAS) for optimal signaling, but not a PTAS: an ϵ additive approximation algorithm for every constant ϵ independent of n . A more immediately attainable goal might be a quasipolynomial time approximation scheme (QPTAS), since a quasipolynomial time algorithm exists for the planted clique problem. In the event of such positive results, it makes sense to examine extensions to more general (multi-player and non-zero sum) games.

The reader might have noticed some similarities between optimal signaling and the problem of optimizing over Nash equilibria, which is known to admit a QPTAS, but not a PTAS by the reduction of Hazan and Krauthgamer [21] from the planted clique problem. However, those similarities do not appear to give any immediate answers for the signaling question. For example, whereas adaptations of the QPTAS for the best Nash equilibrium problem might be able to identify individual “good” signals, piecing those together into an information structure appears to be nontrivial.

Our results also point towards other potential signaling questions. Much of the related work mentioned in the introduction considers particular game domains in which signaling can be used to effect desirable outcomes. We believe that our ideas, and in particular the connection with dense subgraph detection, can shed light on the algorithmic component of some of these applications — we single out the auction domains studied in [14, 10, 13] as likely targets.

ACKNOWLEDGMENTS

We thank Aditya Bhaskara, Uriel Feige, Nicole Immorlica, Albert Jiang, and Shang-Hua Teng for helpful discussions. We also thank the anonymous FOCS reviewers for helpful feedback and suggestions.

REFERENCES

- [1] I. Abraham, S. Athey, M. Babaioff, and M. Grubb. Peaches, lemons, and cookies: designing auction markets with dispersed information. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 7–8, 2013.
- [2] G. A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500, 1970.

- [3] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. In *Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*, pages 594–598. Society for Industrial and Applied Mathematics, 1998.
- [4] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, and N. Xie. Testing k-wise and almost k-wise independence. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 496–505. ACM, 2007.
- [5] B. P. Ames and S. A. Vavasis. Convex optimization for the planted k-disjoint-clique problem. *Mathematical Programming*, 143(1-2):299–337, 2014.
- [6] B. Bassan, O. Gossner, M. Scarsini, and S. Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, 2003.
- [7] D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures. Technical report, Discussion paper, Cowles Foundation for Research in Economics, Yale University and Princeton University, 2012.
- [8] D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. *Economic Theory Center Working Paper*, (052-2013), 2013.
- [9] D. Blackwell. Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102, 1951.
- [10] P. Bro Miltersen and O. Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 234–247, 2012.
- [11] A. Coja-Oghlan. Graph partitioning via adaptive spectral techniques. *Combinatorics, Probability & Computing*, 19(2):227, 2010.
- [12] Y. Dekel, O. Gurel-Gurevich, and Y. Peres. Finding hidden cliques in linear time with high probability. In *ANALCO*, pages 67–75. SIAM, 2011.
- [13] S. Dughmi, N. Immorlica, and A. Roth. Constrained signaling in auction design. In *Proceedings of the 25th ACM Symposium on Discrete Algorithms (SODA)*, 2014.
- [14] Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 514–531, 2012.
- [15] U. Feige and R. Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures and Algorithms*, 16(2):195–208, 2000.
- [16] U. Feige and R. Krauthgamer. The probable value of the lovász–schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2):345–370, 2003.
- [17] U. Feige and D. Ron. Finding hidden cliques in linear time. *DMTCS Proceedings*, (01):189–204, 2010.
- [18] V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. Statistical algorithms and a lower bound for detecting planted cliques. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 655–664. ACM, 2013.
- [19] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.
- [20] M. Guo and A. Deligkas. Revenue maximization via hiding item attributes. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 157–163. AAAI Press, 2013.
- [21] E. Hazan and R. Krauthgamer. How hard is it to approximate the best nash equilibrium? *SIAM Journal on Computing*, 40(1):79–91, 2011.
- [22] J. Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4):561–574, 1971.
- [23] N. Immorlica, A. T. Kalai, B. Lucier, A. Moitra, A. Postlewaite, and M. Tennenholtz. Dueling algorithms. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 215–224, 2011.
- [24] M. Jerrum. Large cliques elude the metropolis process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [25] A. Juels and M. Peinado. Hiding cliques for cryptographic security. *Designs, Codes and Cryptography*, 20(3):269–280, 2000.
- [26] E. Kamenica and M. Gentzkow. Bayesian persuasion. Technical report, National Bureau of Economic Research, 2009.
- [27] L. Kučera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2):193–212, 1995.
- [28] E. Lehrer, D. Rosenberg, and E. Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682, 2010.
- [29] F. McSherry. Spectral partitioning of random graphs. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 529–537, 2001.
- [30] L. Minder and D. Vilenchik. Small clique detection and approximate nash equilibria. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 673–685. Springer, 2009.
- [31] M. Peski. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62(2):732–735, 2008.
- [32] M. Spence. Job market signaling. *The quarterly journal of Economics*, 87(3):355–374, 1973.
- [33] V. Syrgkanis, D. Kempe, and E. Tardos. Information asymmetries in common-value auctions with discrete signals. Preprint., 2013.