Algorithmic Persuasion with No Externalities

SHADDIN DUGHMI, University of Southern California HAIFENG XU, University of Southern California

We study the algorithmics of information structure design -a.k.a. persuasion or signaling - in a fundamental special case introduced by Arieli and Babichenko: multiple agents, binary actions, and no inter-agent externalities. Unlike prior work on this model, we allow many states of nature. We assume that the principal's objective is a monotone set function, and study the problem both in the public signal and private signal models, drawing a sharp contrast between the two in terms of both efficacy and computational complexity.

When private signals are allowed, our results are largely positive and quite general. First, we use linear programming duality and the equivalence of separation and optimization to show polynomial-time equivalence between (exactly) optimal signaling and the problem of maximizing the objective function plus an additive function. This yields an efficient implementation of the optimal scheme when the objective is supermodular or anonymous. Second, we exhibit a (1 - 1/e)-approximation of the optimal private signaling scheme, modulo an additive loss of ϵ , when the objective function is submodular. These two results simplify, unify, and generalize results of [3, 4], extending them from a binary state of nature to many states (modulo the additive loss in the latter result). Third, we consider the binary-state case with a submodular objective, and simplify and slightly strengthen the result of [4] to obtain a (1 - 1/e)-approximation via a scheme which (i) signals independently to each receiver and (ii) is "oblivious" in that it does *not* depend on the objective function so long as it is monotone submodular.

When only a public signal is allowed, our results are negative. First, we show that it is NP-hard to approximate the optimal public scheme, within any constant factor, even when the objective is additive. Second, we show that the optimal private scheme can outperform the optimal public scheme, in terms of maximizing the sender's objective, by a polynomial factor.

Additional Key Words and Phrases: Bayesian Persuasion; Signaling; Information Structure.

1 INTRODUCTION

Information structure design studies how beliefs influence behavior, both of individuals and of groups, and how to shape those beliefs in order to achieve desired outcomes. The key object of study here is the *information structure*, which describes "who knows what" about the parameters governing the payoff function of a game of incomplete information. These parameters, collectively termed the *state of nature*, encode uncertainty in the environment, and their prior distribution together with the information structure determine agents' equilibrium behavior.

When studied descriptively, information structures lend structural insight into the space of potential equilibria of a game of incomplete information (termed Bayes correlated equilibria in [6]), and can provide answers to comparative statics questions regarding the provision or withholding

Contact information: {shaddin, haifengx}@usc.edu .

This work is supported by NSF grant CCF-1350900.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'17, June 26–30, 2017, Cambridge, Massachusetts, USA. © 2017 ACM ISBN 978-1-4503-4527-9/17/06...\$15.00. DOI: http://dx.doi.org/10.1145/3033274.3085152

of information. The associated prescriptive questions are fundamentally algorithmic, and are often termed *persuasion* or *signaling*. Persuasion is the task faced by a principal – call her the *sender* – who has privileged access to the state of nature, and can selectively provide information to the agent(s) in the game – call him/them the *receiver(s)* – in order to further her own objectives. Such a sender effectively *designs* the information structure, and faces the algorithmic problem of selecting which of her information to reveal, to whom, and whether and how to add noise to this information. As in much of the prior work in this area, when we think of the information structure as an algorithm implemented by the sender, we call it a *signaling scheme*.

Both descriptive and prescriptive questions in this space have enjoyed an explosion of research interest of late, in particular at the intersection of economics and computer science. Our discussion can not do justice to the large literature in this area, including the wealth of work on related models. We therefore refer the curious reader to the recent survey by [16] and the references therein. The model we focus on in this paper, and for which we provide a thorough exploration through the algorithmic lens, is the fundamental special case of the multi-agent persuasion problem when restricted to binary actions and no externalities.

Our Model and its Significance

Our technical and conceptual starting point, and arguably the most influential model in this area, is the *Bayesian persuasion* model of [22]. Here the sender faces a single receiver, and her task is to persuade the receiver to take an action which is more favored by the sender. As in persuasion more generally, the sender's "leverage" is her informational advantage — namely, access to the realization of the state of nature, which determines the payoffs of the actions to both the receiver and the sender. The receiver, in contrast, a-priori knows nothing about the state of nature besides its prior distribution (the common prior). Kamenica and Gentzkow [22] assume that the sender can *commit* to her policy for revealing the information prior to realization of the state of nature, and study how and when this signaling scheme can increase the sender's utility.

The model we use is one recently proposed by Arieli and Babichenko [3], both generalizing and restricting aspects of the model of [22]. Here the sender interacts with multiple receivers, each of whom is restricted to a binary choice of actions. Without loss, we denote these actions by 1 and 0. As mentioned in [3, 4], settings like this arise when a manager seeks to persuade investors to invest in a project, or when a principal persuades opinion leaders in a social network with the goal of maximizing social influence. Each receiver's utility depends only on his own action and the state of nature, but crucially not on the actions of other receivers – the *no externality* assumption. The sender's utility, on the other hand, depends on the state of nature as well as the *profile* of receiver actions – since actions are binary, the sender's utility can be viewed as a set function on receivers. As in [22], the state of nature is drawn from a common prior, and the sender can commit to a policy of revealing information regarding the realization of the state of nature. Since there are multiple receivers, this policy - the information structure - is more intricate, since it can reveal different, and correlated, information to different receivers. As made clear in [3], such flexibility is crucial to the sender unless receivers are homogeneous and the sender's utility function highly structured (for example, additively separable across receivers). In particular, if restricted to a public communication channel, the sender is limited in her ability to discriminate between receivers and correlate their actions, whereas a private communication channel provides more flexibility.

Our results, and the models of Kamenica and Gentzkow [22] and Arieli and Babichenko [3], are crucially underlied by the assumption that the sender has the power of commitment to a signaling scheme. The commitment assumption is not as unrealistic as it might first sound, and a number of arguments to that effect are provided in [16, 22, 24]. We mention one of those arguments here:

commitment arises organically at equilibrium if the sender and receiver(s) interact repeatedly over a long horizon, in which case commitment can be thought of as a proxy for "establishing credibility."

If one permits commitment by the sender, restricts attention to a common prior distribution, and postulates that receiver(s) are not privately informed through channels not controlled by the sender, then the model of [22] is *the* special case of persuasion for a single receiver, and the model of [3] is *the* special case of persuasion for multiple receivers with binary actions and no externalities. Viewed this way, both are fundamental special cases of the general information structure design problem, meaningfully restricted. It is thus only fitting that they be thoroughly explored via an algorithmic lens, en-route to a more general understanding of the algorithmics of information in multi-agent settings. A computational study of the single-receiver Bayesian persuasion model of [22] was undertaken by Dughmi and Xu [18]. As for multi-agent persuasion with binary actions and no externalities, a partial algorithmic understanding is provided by [3, 4] in the special case of a binary state of nature, and this paper picks up there.

Context: Private and Public Persuasion

At their most general, signaling schemes can reveal different information to different receivers, through *private* communication channels. Such schemes play a dual role: they *inform* receivers, possibly asymmetrically, and they *coordinate* their behavior by correlating the information provided. In some settings, however, such private communication channels are unrealistic, and the sender is constrained to a *public* communication channel. Work on both descriptive and prescriptive questions regarding information structure design in multi-agent settings can be classified along these lines, and the extent to which a public communication channel limits the sender's powers of persuasion is a fundamental question which has not been thoroughly explored.

Much of the earlier work on information structure design, in particular its computational aspects, focused on public signaling models. This includes work on signaling in auctions [8, 17, 19, 21], voting [2], routing [7], and abstract game models [7, 14, 15, 25]. The work of [14] is relevant to this paper, in that they identify conditions under which public persuasion problems are tractable to approximate, and prove impossibility results in some cases where those conditions are violated. Our hardness proof in Section 6 is in part based on some of their ideas.

Private persuasion has been less thoroughly explored, particularly through the computational lens. The space of (private channel) information structures is studied by Bergemann and Morris [6], who observe that these information structures and their associated equilibria form a generalization of correlated equilibria, and term the generalization the *Bayes Correlated Equilibrium (BCE)*. The space of BCEs is characterized in two-agent two-action games by [26]. Moreover, recent work explores private persuasion in the context of voting [5, 13, 27].

Arieli and Babichenko [3] pose the model we consider in this paper, with the goal of studying private persuasion. We believe that their model wisely simplifies the general multi-agent persuasion problem by removing the most thorny aspect limiting progress in the study of private schemes: externalities. Specifically, by assuming that receivers only influence each others' payoffs through the choices of the sender, they sidestep the equilibrium selection and computation concerns which would arise in more general settings. Assuming a binary state of nature, Arieli and Babichenko [3] provide explicit characterizations of the optimal private signaling scheme for three of the most natural classes of sender objective functions: viewing the sender's objective as a set function on receivers, those are supermodular, anonymous monotone submodular, and supermajority functions. In all three cases, their characterization can be easily converted to an efficient algorithm. Moreover Arieli and Babichenko [3] provide necessary and sufficient conditions under which public signaling schemes match the performance of private signaling schemes. In follow-up work, Babichenko and

Barman [4] also consider the binary state setting, and reduce the private persuasion problem to the problem of computing the concave closure of the sender's utility function. This connection yields a $(1 - \frac{1}{e} - \epsilon)$ -approximate private signaling scheme for monotone submodular utility functions, and an optimal private scheme for anonymous utility functions. Moreover, they show that the former approximation result is almost tight, assuming $P \neq NP$.

Our Results and Techniques

In Section 3, we examine private signaling in our model in the presence of many states of nature. When the sender's set-function objective is monotone non-decreasing, and the prior distribution is given as input, we show the polynomial-time equivalence of optimal (private) persuasion and the algorithmic problem of maximizing the sender's objective function plus an additive function (subject to no constraints).¹ The proof relies on a linear programming formulation of private persuasion, LP duality, and the equivalence of separation and optimization. This leads to polynomial-time persuasion algorithms when the sender's objective is anonymous or supermodular. Our results here are similar to the line of research on optimal mechanism design [9–11], in that both relate economic design problems with Bayesian incentive constraints to purely algorithmic problems. However, our techniques are different from those in [9–11].

Next, we consider a sender with a monotone submodular objective, and efficiently compute a private signaling scheme which is $(1 - \frac{1}{e})$ -approximately optimal for the sender, modulo an arbitrarily small additive ϵ . This generalizes the result of [4] from two states of nature to many states, modulo the additive loss. The techniques used in [4] do not appear to help in the presence of many states of nature, since they are tied to several structural properties which only hold in the case of a binary state. Therefore, our algorithm uses a different approach, and crucially relies on a new structural property of (approximately) optimal private signaling schemes. Specifically, we prove that there always exists an ϵ -optimal "simple" private signaling scheme which is a uniform mixture of polynomially many deterministic schemes, i.e., a scheme that deterministically sends a signal to each receiver upon receiving a state of nature. Notably, this property only depends on the monotonicity of the sender's utility function, and does not rely on submodularity. Using this property, we then use ideas from the literature on submodular function optimization to compute a $(1 - \frac{1}{e})$ -approximation to the best "simple" scheme.

We note that our algorithmic results from Sections 3 and 4 can be approximately extended to the sample oracle model, in which our algorithm is only given sample access to the prior distribution, using ideas from [18]. The resulting schemes suffer an arbitrarily small additive loss in the sender's objective and in persuasiveness of their recommendations, and this loss is unavoidable for information theoretic reasons.

In Section 5, we examine private signaling in the special case with two states of nature and a monotone submodular objective. We give a simple and explicit construction of a polynomial-time private signaling scheme that serves as a $(1 - \frac{1}{e})$ -approximation to the optimal scheme, which is the best possible assuming $P \neq NP$ as shown by [4]. This result simplifies, and slightly strengthens, a result of [4]. Moreover, the constructed private scheme has the following distinctive properties: (i) it is *independent* in that it signals independently to each receiver in each of the two states of nature; (ii) it is *oblivious* in the sense that it does not depend on the sender's utility function, and the approximation ratio is guaranteed so long as the function is monotone submodular. To obtain this result, we first use the idea of the *correlation gap* [1] to argue that there always exists an independent signaling scheme which is a $(1 - \frac{1}{e})$ approximation to the optimal private scheme. We

¹We mention that the reduction from our persuasion problem to the set function maximization problem does not require monotonicity, but the opposite reduction does.

then exploit the fact that there are two states of nature to argue that our scheme is the optimal independent scheme, simultaneously for all monotone objectives. We then show two respects in which this result cannot generalize to many states of nature. First, we show that it is NP-hard to obtain a $(1 - \frac{1}{e})$ -approximation to the best independent signaling scheme in multi-state settings. Second, we show that oblivious schemes cannot guarantee more than a $\frac{1}{\sqrt{m-1}}$ fraction of the optimal sender utility where m > 1 is the number of states of nature.

Finally, in Section 6, we consider public signaling in our model, and present two negative results. First, we show via a simple example that the optimal private scheme can outperform the optimal public scheme, in terms of maximizing the sender's objective, by a polynomial factor. Second, we employ a reduction from an NP-hard graph coloring promise problem to rule out an algorithm for approximating the optimal public scheme to within any constant factor, and also to rule out an additive PTAS for the problem.

2 PRELIMINARIES

2.1 Basic Setup

As in [3], we consider the special case of multi-agent persuasion with binary actions, no inter-agent externalities, and a monotone objective function. Here, we adopt the perspective of *sender* facing *n receivers*. Each receiver has two actions, which we denote by 0 and 1. The receiver's payoff depends only on his own action and a random *state of nature* θ supported on Θ . In particular, we use $u_i(\theta, 1)$ and $u_i(\theta, 0)$ to denote receiver *i*'s utility for action 1 and action 0, respectively, at the state of nature θ ; as shorthand, we use $u_i(\theta) = u_i(\theta, 1) - u_i(\theta, 0)$ to denote how much receiver *i* prefers action 1 over action 0 given state of nature θ . Note that $u_i(\theta)$ may be negative. The sender's utility (our objective) is a function of all the receivers' actions and the state of nature θ . We use $f_{\theta}(S)$ to denote the sender's utility when the state of nature is θ and *S* is the set of receivers who choose action 1. We assume throughout the paper that f_{θ} is a monotone non-decreasing set function for every θ . For convenience in stating our approximation guarantees, we assume without loss of generality that f_{θ} is normalized so that $f_{\theta}(\emptyset) = 0$ and $f_{\theta}(S) \in [0, 1]$ for all $\theta \in \Theta$ and $S \subseteq [n]$.

As is typical in information structure design, we assume that θ is drawn from a common prior distribution λ , that the sender has access to the realized state of nature, and that the sender can publicly *commit* to a policy— termed a *signaling scheme* — for mapping the realized state of nature to a *signal* for each receiver. The signaling scheme may be randomized, and hence reveals noisy partial information regarding the state of nature. The order of events is as follows: (1) The sender commits to a signaling scheme φ ; (2) Nature draws $\theta \sim \lambda$; (3) Signals $(\sigma_1, \ldots, \sigma_n) \sim \varphi(\theta)$ are drawn, and each receiver *i* receives the signal σ_i ; (4) Receivers select their actions.

A general signaling scheme permits sending different signals to different receivers through a private communication channel — we term these *private signaling schemes* to emphasize this generality. We also study the special case of *public signaling schemes* — these are restricted to a public communication channel, and hence send the same signal to all receivers. We discuss these two signaling models in Sections 2.2 and 2.3, including the equilibrium concept and the induced sender optimization problem for each. In both cases, we are primarily interested in the optimization problem faced by the sender in step (1), the goal of which is to maximize the sender's expected utility. When φ yields expected sender utility within an additive [multiplicative] ϵ of the best possible, we say it is ϵ -optimal [ϵ -approximate] in the additive [multiplicative] sense.

2.2 Private Signaling Schemes

A private signaling scheme φ is a randomized map from the set of states of nature Θ to a set of signal profiles $\Sigma = \Sigma_1 \times \Sigma_2 \dots \times \Sigma_n$, where Σ_i is the signal set of receiver *i*. We use $\varphi(\theta, \sigma)$ to denote the

probability of selecting the signal profile $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$ given a state of nature θ . Therefore, $\sum_{\sigma \in \Sigma} \varphi(\theta, \sigma) = 1$ for every θ . With some abuse of notation, we use $\varphi(\theta)$ to denote the random signal profile selected by the scheme φ given the state θ . Moreover, for each $\theta \in \Theta$, $i \in [n]$, and $\sigma_i \in \Sigma_i$, we use $\varphi_i(\theta, \sigma_i) = \Pr[\varphi_i(\theta) = \sigma_i]$ to denote the marginal probability that receiver *i* receives signal σ_i in state θ . An algorithm *implements* a signaling scheme φ if it takes as input a state of nature θ , and samples the random variable $\varphi(\theta)$.

Given a signaling scheme φ , each signal $\sigma_i \in \Sigma_i$ for receiver *i* is realized with probability $\Pr(\sigma_i) = \sum_{\theta \in \Theta} \lambda(\theta) \varphi_i(\theta, \sigma_i)$. Upon receiving σ_i , receiver *i* performs a Bayesian update and infers a posterior belief over the state of nature, as follows: the realized state is θ with posterior probability $\lambda(\theta)\varphi_i(\theta, \sigma_i)/\Pr(\sigma_i)$. Receiver *i* then takes the action maximizing his posterior expected utility. In case of indifference, we assume ties are broken in favor of the sender (i.e., in favor of action 1). Therefore, receiver *i* chooses action 1 if

$$\frac{1}{\Pr(\sigma_i)}\sum_{\theta\in\Theta}\lambda(\theta)\varphi_i(\theta,\sigma_i)u_i(\theta,1)\geq \frac{1}{\Pr(\sigma_i)}\sum_{\theta\in\Theta}\lambda(\theta)\varphi_i(\theta,\sigma_i)u_i(\theta,0),$$

or equivalently

$$\sum_{\theta \in \Theta} \lambda(\theta) \varphi_i(\theta, \sigma_i) u_i(\theta) \ge 0,$$

where $u_i(\theta) = u_i(\theta, 1) - u_i(\theta, 0)$.

A simple revelation-principle style argument [3, 22] shows that there exist an optimal private signaling scheme which is *direct* and *persuasive*. By *direct* we mean that signals correspond to actions — in our setting $\Sigma_i = \{0, 1\}$ for each receiver i — and can be interpreted as action recommendations. A direct scheme is *persuasive* if the strategy profile where all receivers follow their recommendations forms an equilibrium of the resulting Bayesian game.² Due to the absence of inter-receiver externalities in our setting, such an equilibrium would necessarily also satisfy the stronger property of being a dominant-strategy equilibrium — i.e., each receiver *i* maximizes his posterior expected utility by following the recommendation, regardless of whether other receivers follow their recommendations.

When designing private signaling schemes, we restrict attention (without loss) to direct and persuasive schemes. Here, a signal profile can be equivalently viewed as a set $S \subseteq [n]$ of receivers — namely, the set of receivers who are recommended action 1. Using this alternative representation, a scheme can be specified by variables $\varphi(\theta, S)$ for all $\theta \in \Theta, S \subseteq [n]$. We can now encode the sender's optimization problem of computing the optimal scheme using the following exponentially-large linear program; note the use of auxiliary variables $x_{\theta,i}$ to denote the marginal probability of recommending action 1 to receiver *i* in state θ .

$$\begin{array}{ll} \text{maximize} & \sum_{\theta \in \Theta} \lambda(\theta) \sum_{S \subseteq [n]} \varphi(\theta, S) f_{\theta}(S) \\ \text{subject to} & \sum_{S:i \in S} \varphi(\theta, S) = x_{\theta,i}, & \text{for } i \in [n], \theta \in \Theta. \\ & \sum_{\theta \in \Theta} \lambda(\theta) x_{\theta,i} u_i(\theta) \ge 0, & \text{for } i = 1, ..., n. \\ & \sum_{S \subseteq [n]} \varphi(\theta, S) = 1, & \text{for } \theta \in \Theta. \\ & \varphi(\theta, S) \ge 0, & \text{for } \theta \in \Theta; S \subseteq [n]. \end{array}$$

$$(1)$$

The second set of constraints in LP (1) are *persuasiveness constraints*, and state that each receiver *i* should maximize his utility by taking action 1 whenever that action 1 is recommended. Note that the persuasiveness constraints for action 0, which can be written as $\sum_{\theta \in \Theta} \lambda(\theta)(1 - x_{\theta,i})u_i(\theta) \le 0$ for each $i \in [n]$, are intentionally omitted from this LP. This omission is without loss when f_{θ} is a

²Persuasiveness has also been called *incentive compatibility* or *obedience* in prior work.

non-decreasing set function for each θ : any solution to the LP in which a receiver prefers action 1 when recommended action 0 can be improved by always recommending action 1 to that receiver.

2.3 Public Signaling Schemes

A public signaling scheme π can be viewed as a special type of private signaling schemes in which each receiver must receive the same signal, i.e., only a public signal is sent. Overloading the notation of Section 2.2, we use Σ to denote the set of public signals and $\sigma \in \Sigma$ to denote a public signal. A public signaling scheme π is fully specified by $\{\pi(\theta, \sigma)\}_{\theta,\sigma}$, where $\pi(\theta, \sigma)$ denotes the probability of sending signal σ at state θ . Upon receiving a signal σ , each receiver performs the same Bayesian update and infers a posterior belief over the state of nature, as follows: the realized state is θ with probability $\lambda(\theta)\pi(\theta, \sigma)/\Pr(\sigma)$, where $\Pr(\sigma) = \sum_{\theta \in \Theta} \pi(\theta, \sigma)$. This induces a subgame for each signal σ , one in which all receivers share the same belief regarding the state of nature.

Whereas in more general settings than ours receivers may play a mixed Nash equilibrium in each subgame, our restriction to a setting with no externalities removes this complication. Given a posterior distribution on states of nature (say, one induced by a signal σ), our receivers face disjoint single-agent decision problems, each of which admits an optimal pure strategy. We assume that receivers break ties in favor of the sender (specifically, in favor of action 1), which distinguishes a unique pure response for each receiver. Therefore, our solution concept here distinguishes a unique action profile for each posterior distribution, and hence for each signal. A simple revelation-principle style argument then allows us to conclude that there is an optimal public signaling schemes which is *direct*, meaning that the public signals are action profiles, and *persuasive*, meaning that in the subgame induced by signal $\sigma = (\sigma_1, \ldots, \sigma_n)$ each receiver *i*'s optimal decision problem (which breaks ties in favor of action 1) solves to action σ_i .

Restricting attention to direct and persuasive public signaling schemes, each signal can also be viewed as a subset $S \subseteq [n]$ of receivers taking action 1. The sender's optimization problem can then be written as the following exponentially-large linear program.

$$\begin{array}{ll} \text{maximize} & \sum_{\theta \in \Theta} \lambda(\theta) \sum_{S \subseteq [n]} \pi(\theta, S) f_{\theta}(S) \\ \text{subject to} & \sum_{\theta \in \Theta} \lambda(\theta) \pi(\theta, S) \cdot u_{i}(\theta) \geq 0, & \text{for } S \subseteq [n] \text{ with } i \in S. \\ & \sum_{S \subseteq [n]} \pi(\theta, S) = 1, & \text{for } \theta \in \Theta. \\ & \pi(\theta, S) \geq 0, & \text{for } \theta \in \Theta; S \subseteq [n]. \end{array}$$

$$(2)$$

The first set of constraints are persuasiveness constraints corresponding to action 1. Note that the persuasiveness constraints for action 0, which can be written as $\sum_{\theta \in \Theta} \lambda(\theta) \pi(\theta, S) \cdot u_i(\theta) \leq 0$ for each $S \subseteq [n]$ and $i \notin S$, are intentionally omitted from this LP. This omission is without loss when f_{θ} is non-decreasing for each state θ : if signal *S* with $i \notin S$ is such that receiver *i* prefers action 1 in the resulting subgame, then we can replace it with the signal $S \cup i$ without degrading the sender's utility. We remark that LP (2) and LP (1) only differ in their persuasiveness constraints.

2.4 Input Models

We distinguish two input models for describing persuasion instances in this paper. The first is the *explicit* model, in which the prior distribution λ is given explicitly as a probability vector. The second is the *sample oracle* model, where Θ and λ are provided implicitly through sample access to λ . In both models, we assume that given a state of nature θ , we can efficiently evaluate $u_i(\theta)$ for each $i \in [n]$ and $f_{\theta}(S)$ for each $S \subseteq [n]$. In the explicit input model, by *computing* a signaling scheme φ (whether private or public) we mean that we explicitly list the non-zero variables $\varphi(\theta, S)$ on which the scheme is supported. In the implicit model, computing a signaling scheme φ (whether private or public) amounts to providing an algorithm which takes as input a state of nature θ , and samples the random variable $\varphi(\theta)$.

2.5 Set Functions and Submodularity

Given a finite ground set X, a set function is a map $f : 2^X \to \mathbb{R}$. Such a function is nonnegative if $f(S) \ge 0$ for all $S \subseteq X$, monotone non-decreasing (or monotone for short) if $f(S) \le f(T)$ for all $S \subseteq T$. Most importantly, f is submodular if for any $S, T \subseteq X$, we have $f(S \cup T) + f(S \cap T) \le f(S) + f(T)$.

We also consider continuous functions G from the solid hypercube $[0, 1]^X$ to the real numbers. Such a function is *nonnegative* if $G(x) \ge 0$ for all x, *monotone non-decreasing* (or *monotone* for short) if $G(x) \le G(y)$ whenever $x \le y$ (coordinate wise), and *smooth submodular* (in the sense of [12]) if its second partial derivatives exist and are non-positive everywhere.

The Multilinear Extension of a Set Function. Given any set function $f : 2^X \to \mathbb{R}$, the *multilinear extension* of f is the continuous function $F : [0, 1]^X \to \mathbb{R}$ defined as follows:

$$F(x) = \sum_{S \subseteq X} f(S) \prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i),$$
(3)

Notice that, F(x) can be viewed as the expectation of f(S) when the random set *S* independently includes each element *i* with probability x_i . In particular, let p_x^I denote the *independent distribution* with marginals *x*, defined by $p_x^I(S) = \prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i)$, then $F(x) = \mathbb{E}_{S \sim p_x^I} f(S)$. If *f* is nonnegative/monotone then so is *F*. Moreover, if *f* is submodular then *F* is smooth submodular. For our results, we will need to maximize F(x) subject to a set of linear constraints on *x*. This problem is NP-hard in general, yet can be approximated by the *continuous greedy process* of Calinescu et al. [12] for fairly general families of constraints. Note that though we cannot exactly evaluate F(x) in polynomial time, it is sufficient to approximate F(x) within a good precision in order to apply the continuous greedy process. By an additive FPTAS evaluation oracle for *F*, we mean an algorithm that evaluates F(x) within additive error ϵ in poly $(n, \frac{1}{\epsilon})$ time.

THEOREM 2.1 (ADAPTED FORM [12]). Let $F : [0, 1]^n \to [0, 1]$ be a non-negative, monotone, smooth submodular function. Let $\mathcal{P} \subseteq [0, 1]^n$ be a down-monotone polytope³, specified explicitly by its linear constraints. Given an additive FPTAS evaluation oracle for F, there is a poly $(n, \frac{1}{\epsilon})$ time algorithm that outputs $\overline{x} \in \mathcal{P}$ such that $F(\overline{x}) \ge (1 - \frac{1}{\epsilon})OPT - \epsilon$, where $OPT = \max_{x \in \mathcal{P}} F(x)$.

Correlation Gap. A general definition of the correlation gap can be found in [1]. For our results, the following simple definition will suffice. Specifically, for any $x \in [0, 1]^X$, let D(x) be the set of all distributions p over 2^X with fixed marginal probability $\Pr_{S \sim p}(i \in S) = x_i$ for all i. Let p_x^I , as defined above, be the independent distribution with marginal probabilities x. Note that $p_x^I \in D(x)$. For any set function f(S), the correlation gap κ is defined as follows:

$$\kappa = \max_{x \in [0,1]^X} \max_{p \in D(x)} \frac{\mathbf{E}_{S \sim p} f(S)}{\mathbf{E}_{S \sim p_x^I} f(S)}.$$
(4)

Loosely speaking, the correlation gap upper bounds the "loss" of the expected function value over a random set by ignoring the correlation in the randomness.

THEOREM 2.2 ([1]). The correlation gap κ is upper bounded by $\frac{e}{e-1}$ for any non-negative monotone non-decreasing submodular function.

³A polytope $\mathcal{P} \subseteq \mathbb{R}^{n}_{+}$ is called *down-monotone* if for all $x, y \in \mathbb{R}^{n}_{+}$, if $y \in P$ and $x \leq y$ (coordinate-wise) then $x \in P$.

3 EQUIVALENCE BETWEEN PRIVATE SIGNALING AND OBJECTIVE MAXIMIZATION

In this section, we relate the computational complexity of private persuasion to the complexity of maximizing the sender's objective function, and show that the optimal private signaling scheme can be computed efficiently for a broad class of sender objectives. Let \mathcal{F} denote any collection of monotone set functions. We use $I(\mathcal{F})$ to denote the class of all persuasion instances in our model in which the sender utility function f_{θ} is in \mathcal{F} for all states of nature θ . We restrict attention to the explicit input model for most of this discussion, though discuss how to extend our results to the sample oracle model, modulo an arbitrarily small additive loss in both the sender's objective and the persuasiveness constraints, at the end of this section.

The following theorem establishes the polynomial-time equivalence between computing the optimal private signaling scheme and the problem of maximizing the objective function plus an additive function. Note that although the number of variables in LP (1) is exponential in the number of receivers, a vertex optimal solution of this LP is supported on $O(n|\Theta|)$ variables.

THEOREM 3.1. Let \mathcal{F} be any collection of monotone set functions. There is a polynomial-time algorithm which computes the optimal private signaling scheme given any instance in $\mathcal{I}(\mathcal{F})$ if and only if there is a polynomial time algorithm for maximizing $f(S) + \sum_{i \in S} w_i$ given any $f \in \mathcal{F}$ and any set of weights $w_i \in \mathbb{R}$.

PROOF. We first reduce optimal private signaling to maximizing the objective function plus an additive function, via linear programming duality. In particular, consider the following dual program of LP (1) with variables $w_{\theta,i}$, α_i , y_{θ} .

$$\begin{array}{ll} \text{minimize} & \sum_{\theta \in \Theta} y_{\theta} \\ \text{subject to} & \sum_{i \in S} w_{\theta, i} + y_{\theta} \ge \lambda(\theta) f_{\theta}(S), & \text{for } S \subseteq [n], \theta \in \Theta. \\ & w_{\theta, i} + \alpha_{i} \lambda(\theta) u_{i}(\theta) = 0, & \text{for } i = 1, ..., n. \\ & \alpha_{i} \ge 0, & \text{for } i \in [n]. \end{array}$$

$$(5)$$

We can obtain a separation oracle for LP (5) given an algorithm for maximizing $f_{\theta}(S)$ plus an additive function. Given any variables $w_{\theta,i}, \alpha_i, y_{\theta}$, separation over the first set of constraints reduces to maximizing the set function $g_{\theta}(S) = f_{\theta}(S) - \frac{1}{\lambda(\theta)} \sum_{i \in S} w_{\theta,i}$ for each $\theta \in \Theta$. The other constraints can be checked directly in linear time. Given the resulting separation oracle, we can use the Ellipsoid method to obtain a vertex optimal solution to both LP (5) and its dual LP (1) in polynomial time [20].

We now prove the converse. Namely, we construct a polynomial-time Turing reduction from the problem of maximizing f plus an additive function to a private signaling problem in $I(\mathcal{F})$. At a high level, we first reduce the set function maximization problem to a certain linear program, and then prove that solving the dual of the linear program reduces to optimal private signaling for a set of particularly constructed instances in $I(\mathcal{F})$.

Given $f \in F$ and weights w, our reduction concerns the following linear program, parameterized by $\mathbf{a} = (a_1, ..., a_n)$ and b, with variables $\mathbf{z} = (z_1, ..., z_n)$ and v.

$$\begin{array}{ll} \text{minimize} & \sum_{i \in [n]} a_i z_i + bv\\ \text{subject to} & \sum_{i \in S} z_i + v \ge f(S), \quad \text{for } S \subseteq [n]. \end{array}$$
(6)

Let \mathcal{P} denote the feasible region of LP (6). As the first step of our reduction, we reduce maximizing the set function $g_w(S) = f(S) + \sum_{i \in S} w_i$ to the separation problem for \mathcal{P} . Let $z_i = -w_i$ for each *i*. Notice that (z, v) is feasible (i.e., in \mathcal{P}) if and only if $v \ge \max_{S \subseteq [n]} f(S) - \sum_{i \in S} z_i$. Therefore, we can binary search for a value \tilde{v} such that (z, \tilde{v}) is almost feasible, but not quite. More precisely, let *B* denote the bit complexity of the f(S)'s and the w_i 's, then binary search returns the exact optimal value of the set function maximization problem after O(B) steps. We then set \tilde{v} to equal that value minus 2^{-B} . Feeding (z, \tilde{v}) to the separation oracle, we obtain a violated constraint which must correspond to the maximizer of $f(S) + \sum_{i \in S} w_i$.

As the second step of our reduction, we reduce the separation problem for \mathcal{P} to solving LP (6) for every choice of objective coefficients **a** and *b*. This polynomial-time Turing reduction follows from the equivalence of separation and optimization [20].

Third, we reduce solving LP (6) for arbitrary **a** and *b* to the special case where $\mathbf{a} \in [0, 1]^n$ and b = 1. The reduction involves a case analysis. (a) If any of the objective coefficients are negative, then the fact that \mathcal{P} is upwards closed implies that LP (6) is unbounded. (b) If b = 0 and $a_i > 0$ for some *i*, then the LP is unbounded since we can make *v* arbitrarily small and z_i arbitrarily large. Normalizing by dividing by *b*, we have reduced the problem to the case when b = 1 and $a \ge 0$ (coordinate-wise). (c) Now suppose that $a_i > 1 = b$ for some *i*; the LP is unbounded by making z_i arbitrarily small and v arbitrarily large. This analysis leaves the case of b = 1 and $\mathbf{a} \in [0, 1]^n$.

Fourth, we reduce LP (6) with parameters $\mathbf{a} \in [0, 1]^n$ and b = 1 to its dual shown below, with variables p_S for $S \subseteq [n]$.

$$\begin{array}{ll} \text{maximize} & \sum_{S \subseteq [n]} p_S f(S) \\ \text{subject to} & \sum_{S:i \in S} p_S \leq a_i, \quad \text{for } i \in [n]. \\ & \sum_{S \subseteq [n]} p_S = 1 \\ & p_S \geq 0, \qquad \text{for } S \subseteq [n]. \end{array}$$
(7)

We note that LP (7) is not the standard dual of LP (6); in particular the first set of constraints are inequality rather than equality constraints. It is easy to see that LP (7) is equivalent to the standard dual when f is monotone non-decreasing, and that an optimal solution to one of the two duals can be easily converted to an optimal solution of the other.

The fifth and final step of our reduction will reduce LP (7) to a private signaling problem in $I(\mathcal{F})$. There are *n* receivers and two states of nature θ_0 , θ_1 with $\lambda(\theta_0) = \lambda(\theta_1) = 1/2$. Define $u_i(\theta_1) = 1$ and $u_i(\theta_0) = -\frac{1}{a_i}$ ($-\infty$ if $a_i = 0$) for all *i*. The sender's utility function satisfies $f_{\theta_1} = f_{\theta_0} = f$. Let φ^* be an optimal signaling scheme, in particular an optimal solution to the instantiation of LP (1) for our instance. Note that all receivers prefer action 1 in state θ_1 ; therefore, it is easy to see that φ^* can be weakly improved, without violating the persuasiveness constraints, by modifying it to always recommend action 1 to all receivers when in state θ_1 . After this modification, φ^* is an optimal solution to the following LP, which optimizes over all signaling schemes satisfying $\varphi(\theta_1, [n]) = 1$.

$$\begin{array}{ll} \text{maximize} & \frac{1}{2}f([n]) + \frac{1}{2}\sum_{S \subseteq [n]}\varphi(\theta_0, S)f(S) \\ \text{subject to} & \sum_{S:i \in S}\varphi(\theta_0, S) = x_{\theta_0, i}, & \text{for } i \in [n]. \\ & x_{\theta_0, i} \leq a_i, & \text{for } i = 1, ..., n. \\ & \sum_{S \subseteq [n]}\varphi(\theta_0, S) = 1 \\ & \varphi(\theta_0, S) \geq 0, & \text{for } \theta \in \Theta; S \subseteq [n]. \end{array}$$

$$(8)$$

It is now easy to see that setting $p_S = \varphi^*(\theta_0, S)$ yields an optimal solution to LP (7)

As an immediate corollary of Theorem 3.1, the optimal private signaling scheme can be computed efficiently when the sender's objective function is supermodular or anonymous. Recall that a set function $f : 2^{[n]} \to \mathbb{R}$ is anonymous if there exists a function $g : \mathbb{Z} \to \mathbb{R}$ such that f(S) = g(|S|).

COROLLARY 3.2. There is a polynomial time algorithm for computing the optimal private signaling scheme when the sender objective functions are either supermodular or anonymous.

PROOF. Since a supermodular function plus an additive function is still supermodular, and the problem of unconstrained supermodular maximization can be solved in polynomial time. Theorem 3.1 implies that the optimal private signaling scheme can also be computed in polynomial time. As for anonymous objectives, there is a simple algorithm for maximizing an anonymous set function

plus an additive function. In particular, consider the problem of maximizing $f(S) + \sum_{i \in S} w_i$ where f(S) = g(|S|). Observe that fixing |S| = k, the optimal set S_k corresponds to the k highest-weight elements in w. Enumerating all k and choosing the best S_k yields the optimal set. \Box

Finally, we make two remarks on Theorem 3.1, particularly on the reduction from optimal private signaling to set function maximization. First, we note that the assumption of monotonicity is not necessary to the reduction from signaling to optimization. In other words, even without the monotonicity assumption for the sender's objective function, one can still efficiently compute the optimal private signaling scheme for instances in $I(\mathcal{F})$ given access to an oracle for maximizing $f(S) + \sum_{i \in S} w_i$ for any $f \in \mathcal{F}$ and weight vector w. This can be verified by adding back the persuasiveness constraints for action 0 to LP (1) and examining the corresponding dual, which has similar structure to LP (5). We omit the trivial details here. Consequently, Corollary 3.2 applies to non-monotone supermodular or anonymous functions as well.

Second, our reduction assumes that the prior distribution λ over the state of nature is explicitly given. This can be generalized to the sample oracle model. In particular, when our only access to λ is through random sampling, we can implement an ϵ -optimal and ϵ -persuasive⁴ private signaling scheme in $poly(n, \frac{1}{\epsilon})$ time using an idea from [18] (assuming $u_i(\theta) \in [-1, 1]$). The algorithm is as follows: given any input state θ , we first take $poly(n, \frac{1}{\epsilon})$ samples from λ , and then solve LP (1) on the empirical distribution of the samples plus θ , with relaxed (by ϵ) persuasiveness constraints. Finally, we signal for θ as the solution to the LP suggests. The analysis of this algorithm is very similar to that in [18], thus is omitted here. Moreover, the bi-criteria loss is inevitable in this oracle model due to information theoretic reasons [18].

4 PRIVATE SIGNALING WITH SUBMODULAR OBJECTIVES

In this section we consider optimal private signaling for submodular sender objectives, and show that there is a polynomial time $(1 - \frac{1}{e})$ -approximation scheme, modulo an additive loss of ϵ . This is almost the best possible: Babichenko and Barman [4] show that even in the special case of two states of nature, it is NP-hard to approximate the optimal private signaling scheme within a factor better than $(1 - \frac{1}{e})$ for monotone submodular sender objectives.

THEOREM 4.1. Consider private signaling with monotone submodular sender objectives. Let OPT denote the optimal sender utility. For any $\epsilon > 0$, a private signaling scheme achieving expected sender utility at least $(1 - \frac{1}{e})OPT - \epsilon$ can be implemented in $poly(n, |\Theta|, \frac{1}{\epsilon})$ time.

The main technical challenge in proving Theorem 4.1 is that a private signaling scheme may have exponentially large support, as apparent from linear program (1). To overcome this difficulty, we prove a structural characterization of (approximately) optimal persuasive private schemes, i.e., solutions to LP (1). Roughly speaking, we show that LP (1) always has an approximately optimal solution with polynomial-sized support and nicely structured distributions. This greatly narrows down the solution space we need to search over. Recall that for any θ , $\varphi(\theta)$ is a random variable supported on $2^{[n]}$. We say $\varphi(\theta)$ is *K*-uniform if it follows a uniform distribution on a multiset of size *K*. The following lemma exhibits a structural property regarding (approximately) optimal solutions to LP (1). Notably, this property only depends on monotonicity of the sender's objective functions and does not depend on submodularity. Its proof is postponed to the end of this section.

⁴This is the natural relaxation of persuasiveness. In our setting, a receiver approximately maximizes his posterior expected payoff, to within an additive ϵ , by following the scheme's recommendation. This is regardless of the actions of other receivers.

LEMMA 4.2. Let f_{θ} be monotone for each θ . For any $\epsilon > 0$, there exists an ϵ -optimal persuasive private signaling scheme $\overline{\varphi}$ such that $\overline{\varphi}(\theta)$ is K-uniform for every θ , where $K = \frac{108n \log(2n|\Theta|)}{\epsilon^3}$.

By Lemma 4.2, we can, without much loss, restrict our design of $\varphi(\theta)$ to the special class of *K*-uniform distributions. Note that a *K*-uniform distribution $\varphi(\theta)$ can be described by variables $x_{\theta,i}^j \in \{0,1\}$ for $i \in [n], j \in [K]$, where $x_{\theta,i}^j$ denotes the recommended action to receiver *i* in the *j*'th profile in the support of $\varphi(\theta)$. Relaxing our variables to lie in [0, 1], this leads to optimization problem (9), where $F_{\theta}(x) = \sum_{S \subseteq [n]} f_{\theta}(S) \prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i)$ is the multi-linear extension of f_{θ} .

$$\begin{array}{ll} \text{maximize} & \sum_{\theta \in \Theta} \frac{\lambda(\theta)}{K} \sum_{j=1}^{K} F_{\theta}(x_{\theta}^{j}) \\ \text{subject to} & \sum_{\theta \in \Theta} \frac{\lambda(\theta)}{K} \sum_{j=1}^{K} x_{\theta,i}^{j} u_{i}(\theta) \geq 0, \quad \text{for } i = 1, ..., n. \\ & 0 \leq x_{\theta,i}^{j} \leq 1, \qquad \qquad \text{for } i = 1, ..., n; \theta \in \Theta. \end{array}$$

$$(9)$$

As a high level, our algorithm first approximately solves Program (9) and then signals according to its solution. Details are in Algorithm 1, which we instantiate with $\epsilon > 0$ and $K = \frac{108n \log(2n|\Theta|)}{\epsilon^3}$. Since $F_{\theta}(x) = \mathbb{E}_{S \sim p_x^I} f(S)$ where p_x^I is the independent distribution over $2^{[n]}$ with marginal probability x, the expected sender utility induced by the signaling scheme in Algorithm 1 is precisely the objective value of Program (9) at the obtained solution. Theorem 4.1 then follows from two claims: 1. The optimal objective value of Program (9) is ϵ -close to the optimal sender utility (Claim 4.3); 2. The continuous greedy process [12] can be applied to Program (9) to efficiently compute a (1 - 1/e)-approximate solution, modulo a small additive loss (Claim 4.4). We remark that Theorem 4.1 can be generalized to the sample oracle model, but with an additional ϵ -loss in persuasiveness constraints (assuming $u_i(\theta) \in [-1, 1]$), using the idea from [18].

Algorithm 1 Private Signaling Scheme for Submodular Sender Objectives

Parameter: $\epsilon > 0$ **Input:** Prior distribution λ supported on Θ **Input:** $u_i(\theta)$'s and value oracle access to the sender utility $f_{\theta}(S)$ **Input:** State of nature θ **Output:** A set $S \subseteq [n]$ indicating the set of receivers who will be recommended action 1. 1: Approximately solve Program (9). Let $\{\tilde{x}_{\theta,i}^{j}\}_{\theta\in\Theta,i\in[n],j\in[K]}$ be the returned solution. 2: Choose *i* from [K] uniformly at random: For each receiver *i* add *i* to S independent

- 2: Choose *j* from [K] uniformly at random; For each receiver *i*, add *i* to *S* independently with probability $\tilde{x}_{\theta,i}^{j}$.
- 3: Return S.

CLAIM 4.3. When $K = \frac{108n \log(2n|\Theta|)}{\epsilon^3}$, the optimal objective value of Program (9) is at least OPT - ϵ , where OPT is the optimal sender utility in private signaling.

PROOF. By Lemma 4.2, there exists a private signaling scheme $\overline{\varphi}$ such that: (i) $\overline{\varphi}$ achieves sender utility at least $OPT - \epsilon$; (ii) for each θ , there exists K sets $S^1_{\theta}, ..., S^K_{\theta} \subseteq [n]$ such that $\overline{\varphi}_{\theta}$ is a uniform distribution over $\{S^1_{\theta}, ..., S^K_{\theta}\}$. Utilizing $\overline{\varphi}$, we can construct a feasible solution \overline{x} to Program (9) with objective value at least $OPT - \epsilon$. In particular, let $\overline{x}^j_{\theta} \in \{0, 1\}^n$ be the indicator vector of the set S^j_{θ} , formally defined as follows: $\overline{x}^j_{\theta,i} = 1$ if and only if $i \in S^j_{\theta}$. By referring to the feasibility of $\overline{\varphi}$ to LP (1), it is easy to check that $\overline{x}^j_{\theta,i}$'s are feasible to Program (9). Moreover, since $F_{\theta}(\overline{x}^j_{\theta}) = f_{\theta}(S^j_{\theta})$, the objective value of Program (9) at the solution \overline{x} equals the objective value of Program (1) at the solution $\overline{\varphi}$, which is at least $OPT - \epsilon$. Therefore, the optimal objective value of Program (9) is at least $OPT - \epsilon$, as desired.

CLAIM 4.4. There is an algorithm that runs in poly $(n, |\Theta|, K, \frac{1}{\epsilon})$ time and computes a (1 - 1/e)-approximate solution, modulo an additive loss of ϵ/e , to Program (9).

PROOF. The objective function of Program (9) is a linear combination, with non-negative coefficients, of multilinear extensions of monotone submodular functions, thus is smooth, monotone and submodular. Moreover, the function value can be evaluated within error ϵ by $poly(n, \frac{1}{\epsilon})$ random samples, thus in $poly(n, \frac{1}{\epsilon})$ time. To apply Theorem 2.1, we only need to prove that the feasible region is a down-monotone polytope. Observe that there always exists an optimal solution to Program (9) such that $x_{\theta,i} = 1$ for any θ, i such that $u_i(\theta) \ge 0$. Therefore, w.l.o.g., we can pre-set these variables to be 1 and view the program as an optimization problem over $x_{\theta,i}$'s for all θ, i such that $u_i(\theta) < 0$. It is easy to check that these $x_{\theta,i}$'s form a down-monotone polytope determined by polynomially many constraints, as desired.

Proof of Lemma 4.2

Our proof is based on the probabilistic method. Recall that the optimal private signaling scheme can be computed by solving the exponentially-large LP (1). Roughly speaking, given any optimal private scheme φ^* , we will take polynomially many samples from $\varphi^*(\theta)$ for each θ , and prove that with strictly positive probability the corresponding empirical distributions form a solution to LP (1) that is close to optimality. However, the sampling approach usually suffers from ϵ -loss in both the objective and persuasiveness constraints. It turns out that the ϵ -loss to persuasiveness constraints can be avoided in our setting with carefully designed pre-processing steps.

At a high level, to get rid of the ϵ -loss in persuasiveness constraints, there are two main technical barriers. The first is to handle the estimation error in the receiver's utilities, which is inevitable due to sampling. We address this by adjusting the φ^* to strengthen the persuasiveness constraints so that a small estimation error would still preserve the original persuasiveness constraints. The second barrier arises when some $x^*_{\theta,i}$'s are smaller than inverse polynomial of the precision ϵ , then $poly(\frac{1}{\epsilon})$ samples cannot guarantee a good multiplicative estimate of $x^*_{\theta,i}$. We deal with this issue by making "honest" recommendation, i.e., action 0, at these cases, and show that such modification will not cause much loss to our objective.

We first introduce some convenient notations. For any receiver *i*, let set $\Theta_i^+ = \{\theta : u_i(\theta) \ge 0\}$ be the set of states at which receiver *i* (weakly) prefers action 1; Similarly, $\Theta_i^- = \{\theta : u_i(\theta) < 0\}$ be the set of states at which receiver *i* prefers action 0. Moreover, for any state of nature θ , let $I_{\theta}^+ = \{i : u_i(\theta) \ge 0\}$ be the set of receivers who (weakly) prefer action 1 at state θ . It would be convenient to think of $\{\Theta_i^+\}_{i \in [n]}$ and $\{I_{\theta}^+\}_{\theta \in \Theta}$ as two different partitions of the set $\{(\theta, i) : u_i(\theta) \ge 0\}$.

Observe that by monotonicity there always exists an optimal signaling scheme φ^* such that $x_{\theta,i}^* = 1$ for every $\theta \in \Theta_i^+$. Let φ^* be such an optimal signaling scheme and *OPT* denote the optimal sender utility. We now adjust the scheme φ^* such that they do not degrade the objective value by much but is more suitable for applying concentration bounds for our probabilistic argument.

Adjustment 1: Always Recommend Action 0 When $x^*_{\theta,i} < \frac{\epsilon}{3n}$

Note that $x_{\theta,i}^* < \frac{\epsilon}{3n}$ only when $\theta \in \Theta_i^-$, i.e., action 0 is the best action for receiver *i* conditioned on θ . We first adjust ϕ^* to obtain a new scheme $\tilde{\varphi}$, as follows: $\tilde{\varphi}$ is the same as ϕ^* except that for every θ , *i* such that $x_{\theta,i}^* < \frac{\epsilon}{3n}$, $\tilde{\varphi}$ always recommends action 0 to receiver *i* given the state of nature θ . As a result, $\tilde{x}_{\theta,i}$ equals $x_{\theta,i}^*$ whenever $x_{\theta,i}^* \ge \frac{\epsilon}{3n}$ and equals 0 otherwise. Note that the signaling scheme still satisfies the persuasiveness constraints. Naturally, each adjustment above, corresponding to θ , *i* satisfying $x_{\theta,i}^* < \frac{\epsilon}{3n}$, could decrease the objective value since the marginal probability of recommending action 1 decreases. Nevertheless, this loss, denoted as $L(\theta, i)$, can be properly bounded as follows:

$$\begin{split} L(\theta, i) &= \lambda(\theta) \cdot \left[\sum_{S:i \in S} \varphi^*(\theta, S) f_{\theta}(S) - \sum_{S:i \in S} \varphi^*(\theta, S) f_{\theta}(S \setminus \{i\}) \right] \\ &\leq \lambda(\theta) \cdot \left[\sum_{S:i \in S} \varphi^*(\theta, S) \right] \\ &= \lambda(\theta) x_{\theta,i}^* \leq \frac{\lambda(\theta)\epsilon}{3n}. \end{split}$$

As a result, the aggregated loss of all the adjustments made in this step can be upper bounded by $\sum_{\theta \in \Theta} \sum_{i=1}^{n} \frac{\lambda(\theta)\epsilon}{3n} = \frac{\epsilon}{3}$. That is, the objective value of $\tilde{\varphi}$ is at least $OPT - \frac{\epsilon}{3}$.

Adjustment 2: Strengthen the Persuasive Constraints by Scaling Down $x_{\theta,i}$'s

We now strengthen the persuasiveness constraints by further adjusting the $\tilde{\varphi}$ obtained from above so that a small estimation error due to sampling will still maintain the original persuasiveness constraints. For any θ , we define $\varphi'(\theta, S) = \frac{3}{3+\epsilon} \tilde{\varphi}(\theta, S)$ for all $S \neq I_{\theta}^+$, and define $\varphi'(\theta, I_{\theta}^+) =$ $1 - \sum_{S \neq I_{\theta}^+} \varphi'(\theta, S)$. Obviously, φ'_{θ} is still a distribution over $2^{[n]}$. Moreover, we claim that $x'_{\theta,i} =$ $\mathbf{E}_{S \sim \varphi'_{\theta}} \mathbb{I}(i \in S) = 1$ whenever $\tilde{x}_{\theta,i} = 1$, i.e., $\theta \in \Theta_i^+$. That is, given state θ , any receiver $i \in I_{\theta}^+$ will still aways be recommended action 1. This is because, to construct φ'_{θ} , we moved some probability mass from all other sets S to the set I_{θ}^+ , therefore the marginal probability of recommending action 1 to any receiver $i \in I_{\theta}^+$ will not decrease. However, this marginal probability is originally 1 in the solution of $\tilde{\varphi}$. Therefore, $x'_{\theta,i}$ still equals 1 for any $i \in I_{\theta}^+$, or equivalently, for any $\theta \in \Theta_i^+$. Similarly, we also have $x'_{\theta,i} = 0$ whenever $\tilde{x}_{\theta,i} = 0$.

Let $Val(\varphi)$ denote the objective value of a scheme φ . We claim that $Val(\varphi') \ge OPT - \frac{2\epsilon}{3}$ and φ' satisfies $x'_{\theta,i} = \frac{3}{3+\epsilon} \tilde{x}_{\theta,i}$ for every $\theta \in \Theta_i^-$. For any $i \in [n], \theta \in \Theta_i^-$ (which means $i \notin I_{\theta}^+$), we have

$$x'_{\theta,i} = \sum_{S:i \in S} \varphi'(\theta, S) = \frac{3}{3+\epsilon} \sum_{S:i \in S} \widetilde{\varphi}(\theta, S) = \frac{3}{3+\epsilon} \widetilde{x}_{\theta,i},$$

since the summation excludes the term $\varphi'(\theta, I_{\theta}^+)$. We now prove the guarantee of the objective value. Observe that $\varphi'(\theta, I_{\theta}^+) \geq \frac{3}{3+\epsilon} \widetilde{\varphi}(\theta, I_{\theta}^+)$ also holds in our construction. Therefore, we have

$$\begin{aligned} Val(\varphi') &= \sum_{\theta \in \Theta} \lambda(\theta) \sum_{S \subseteq [n]} \varphi'(\theta, S) f_{\theta}(S) \\ &\geq \frac{3}{3 + \epsilon} \sum_{\theta \in \Theta} \lambda(\theta) \sum_{S \subseteq [n]} \widetilde{\varphi}(\theta, S) f_{\theta}(S) \\ &= \frac{3}{3 + \epsilon} \cdot Val(\widetilde{\varphi}) \\ &\geq OPT - \frac{2\epsilon}{3}, \end{aligned}$$

where we used the upper bound $Val(\tilde{\varphi}) \leq 1$.

Existence of An ϵ -Optimal Solution of Small Support.

The above two steps of adjustment result in a feasible $\frac{2\epsilon}{3}$ -optimal solution φ' to LP (1) that satisfies the following properties: (i) $x'_{\theta,i} = x^*_{\theta,i} = 1$ whenever $u_i(\theta) \ge 0$; (ii) $x'_{\theta,i} = \frac{3}{3+\epsilon} \tilde{x}_{\theta,i} = \frac{3}{3+\epsilon} x^*_{\theta,i} \ge \frac{\epsilon}{4n}$ when $x^*_{\theta,i} \ge \frac{\epsilon}{3n}$ and $\theta \in \Theta^-_i$; (iii) $x'_{\theta,i} = 0$ when $x^*_{\theta,i} < \frac{\epsilon}{3n}$ and $\theta \in \Theta^-_i$. Utilizing such a φ' we show that there exists an ϵ -optimal solution $\overline{\varphi}$ to LP (1) such that the distribution $\overline{\varphi}_{\theta}$ is a *K*-uniform distribution for every θ , where $K = \frac{108n \log(2n|\Theta|)}{\epsilon^3}$.

Our proof is based on the probabilistic method. For each θ , independently take $K = \frac{108n \log(2n|\Theta|)}{\epsilon^3}$ samples from random variable $\varphi'(\theta)$, and let $\overline{\varphi}_{\theta}$ denote the corresponding empirical distribution. Obviously, $\overline{\varphi}_{\theta}$ is a *K*-uniform distribution. We claim that with strictly positive probability over the randomness of the samples, $\overline{\varphi}$ is feasible to LP (1) and achieves utility at least $Val(\varphi') - \frac{\epsilon}{3} \ge OPT - \epsilon$.

We first examine the objective value. Observe that the objective value $Val(\varphi')$ can be viewed as the expectation of the random variable $\sum_{\theta \in \Theta} \lambda(\theta) f_{\theta}(S_{\theta}) \in [0, 1]$, where S_{θ} follows the distribution of $\varphi'(\theta)$. Our sampling procedure generates K samples for the random variable $\{S_{\theta}\}_{\theta \in \Theta}$, therefore by the Hoeffding bound, with probability at least $1 - \exp(-2K\epsilon^2/9) > 1 - 1/(2n|\Theta|)$, the empirical mean is at least $Val(\varphi') - \epsilon/3$.

Now we only need to show that all the persuasiveness constraints are preserved with high probability. First, observe that if $x'_{\theta,i} = 0$, then $\overline{x}_{\theta,i}$ induced by $\overline{\varphi}$ also equals 0. This is because $x'_{\theta,i} = \mathbb{E}_{S \sim \varphi'(\theta)} \mathbb{I}(i \in S) = 0$ implies that *i* is not contained in any *S* from the support of $\varphi'(\theta)$, therefore, also not contained in any sample. Similarly, $x'_{\theta,i} = 1$ implies $\overline{x}_{\theta,i} = 1$. To show that all the persuasiveness constraints hold, we only need to argue that $\overline{x}_{\theta,i} \leq x^*_{\theta,i}$ for every $\theta \in \Theta_i^-$ satisfying $x^*_{\theta,i} \geq \frac{\epsilon}{3n}$. This holds with high probability by tail bounds. In particular, $x'_{\theta,i} = \mathbb{E}_{S \sim \varphi'(\theta)} \mathbb{I}(i \in S)$ and we take *K* samples from $\varphi'(\theta)$. By the Chernoff bound, with probability at least

$$1 - \exp(-\frac{K\epsilon^2 x'_{\theta,i}}{27}) \ge 1 - \exp(-\frac{K\epsilon^3}{108n}) > 1 - \frac{1}{2n|\Theta|},$$

the empirical mean $\overline{x}_{\theta,i}$ is at most $(1 + \epsilon/3)x'_{\theta,i} = x^*_{\theta,i}$.

Note that there are at most $n|\Theta|$ choices of such θ , *i*. By union bound, with probability at least $1 - (n|\Theta| + 1)/(2n|\Theta|) > 0$, $\overline{\varphi}$ satisfies all the persuasiveness constraints thus is feasible to LP (1), and achieves objective value at least $Val(\varphi') - \frac{\epsilon}{3} \ge OPT - \epsilon$. So there must exist a feasible ϵ -optimal solution $\overline{\varphi}$ to LP (1) such that $\overline{\varphi}_{\theta}$ is *K*-uniform for every θ . This concludes our proof of Lemma 4.2.

5 AN OBLIVIOUS PRIVATE SCHEME FOR BINARY-STATE SETTINGS

In this section, we consider the special case with two states of nature, denoted by θ_0 and θ_1 , and submodular sender utility functions. We show that a $(1 - \frac{1}{e})$ -approximate private signaling scheme can be *explicitly* constructed. The approximation ratio is tight by [4]. Moreover, the constructed signaling scheme has the following distinctive properties: (i) it signals independently to each receiver, which we term an *independent signaling scheme*; (ii) it is *oblivious* in the sense that it does not depend on the sender's utility function so long as it is monotone submodular. All proofs in this section are omitted due to space constraints, and can be found in the full version on arXiv.

In particular, we consider the independent signaling scheme φ_I that simply maximizes the probability of recommending action 1 to each receiver *i* independently. Such a scheme can be easily constructed by alway recommending action 1 to receiver *i* whenever $u_i(\theta) \ge 0$, and recommend action 1 with probability as high as possible subject to the persuasiveness constraints. We refer the reader to the Appendix for an explicit construction. Notice that φ_I is determined by the receiver's payoffs, and does *not* depend on the sender's payoff function $f_{\theta}(S)$. Nevertheless, the following theorem shows that φ_I is approximately optimal.

THEOREM 5.1. In the binary-state setting with monotone submodular sender objectives, the scheme φ_I is a $(1 - \frac{1}{e})$ -approximate private signaling scheme.

Theorem 5.1 follows from the following two lemmas. The proof of Lemma 5.2 employs the idea of the correlation gap from the robust stochastic optimization literature [1].

LEMMA 5.2. In the multiple-state setting with monotone submodular sender objectives, the optimal independent signaling scheme is a $(1 - \frac{1}{e})$ -approximation to the optimal private signaling scheme.

LEMMA 5.3. In the binary-state setting with monotone sender objectives, the φ_I defined above is the optimal independent signaling scheme.

We conclude this section with two negative results regarding generalizing Theorem 5.1 to many states of natures. We first show that it is NP-hard to approximate the optimal independent signaling scheme within a factor better than $(1 - \frac{1}{e})$ when there are multiple states of nature (Proposition 5.4). We then prove that the best oblivious scheme can perform poorly in terms of the sender utility when there are multiple states (Proposition 5.5).

PROPOSITION 5.4. In the multiple-state setting with monotone submodular sender objectives, it is NP-hard to compute an independent signaling scheme that (1 - 1/e)-approximates the optimal independent scheme.

PROPOSITION 5.5. For any integer m > 1, there exists an instance with m states of nature such that any oblivious private scheme can achieve at most $\frac{1}{|\sqrt{m-1}|}$ fraction of the optimal sender utility.

6 INEFFICACY AND HARDNESS OF PUBLIC PERSUASION

In this section, we turn our attention to the design of optimal public signaling schemes, and show a stark contrast with private signaling, both in terms of their efficacy at optimizing the sender's utility, and in terms of their computational complexity.

We start with an example illustrating how the restriction to public signaling can drastically reduce the sender's expected utility. The example is notably simple: two states of nature, and a binary sender utility function which is independent of the state of nature. We show a multiplicative gap of $\Omega(n)$, and an additive gap of $1 - \frac{1}{\Omega(n)}$, between the expected sender utility from the optimal private and public signaling schemes, where *n* is the number of receivers.

Example 6.1 (Inefficacy of Public Signaling Schemes). Consider an instance with *n* identical receivers and two states of nature $\Theta = \{H, L\}$. Each receiver has the same utility function, defined as follows: $u_i(H) = 1$ and $u_i(L) = -1$, for all *i*. The state of nature H occurs with probability $\frac{1}{n+1}$, and L occurs with probability $\frac{n}{n+1}$. The sender's utility function is $f_{\theta}(S) = f(S) = \min(|S|, 1)$. In other words, the sender gets utility 1 precisely when at least one receiver takes action 1.

The persuasiveness constraints imply that each receiver can take action 1 with probability no more than $\frac{2}{n+1}$. This is achievable by always recommending action 1 to the receiver in state H, and recommending action 0 with probability $\frac{1}{n}$ in state L. The sender's expected utility depends on how these recommendations are correlated.

The optimal private scheme anti-correlates the receivers' recommendations in order to guarantee that at least one receiver takes action 1 always, which achieves an expected sender utility of 1, the maximum possible. Specifically, in state H the scheme always recommends action 1 to every receiver, and in state L the scheme chooses one receiver uniformly at random and recommends action 1 to that receiver, and action 0 to the other receivers.

We argue that no public scheme can achieve sender utility more than $\frac{2}{n+1}$. Indeed, since receivers are identical, our solution concept implies that they choose the same action for every realization

of a public signal. Therefore, the best that a public scheme can do is to recommend action 1 to all receivers simultaneously with probability $\frac{2}{n+1}$ in aggregate, and recommend action 0 with the remaining probability, yielding an expected sender utility of $\frac{2}{n+1}$. This is achievable: in state H the scheme always recommends action 1 to every receiver, and in state L the scheme recommends action 1 to all receivers with probability $\frac{1}{n}$, and action 0 to all receivers with probability $1 - \frac{1}{n}$.

Our next result illustrates the computational barrier to obtaining the optimal public signaling scheme, even for additive sender utility functions. Our proof is inspired by a reduction in [14] for proving the hardness of computing the best posterior distribution over Θ , a problem termed *mixture selection* in [14], in a voting setting. That reduction is from the maximum independent set problem. Since a public signaling scheme is a combination of posterior distributions, one for each signal, we require a more involved reduction from a graph-coloring problem to prove our result.

THEOREM 6.2. Consider public signaling in our model, with sender utility function $f_{\theta}(S) = f(S) = \frac{|S|}{n}$. It is NP-hard to approximate the optimal sender utility to within any constant multiplicative factor. Moreover, there is no additive PTAS for evaluating the optimal sender utility, unless P = NP.

PROOF. We prove by reducing from the following NP-hard problem. In particular, [23] prove that for any positive integer k, any integer q such that $q \ge 2^k + 1$, and an arbitrarily small constant $\epsilon > 0$, given an undirected graph G, it is NP-hard to distinguish between the following two cases:

- **Case 1**: There is a *q*-colorable induced subgraph of *G* containing a (1ϵ) fraction of all vertices, where each color class contains a $\frac{1-\epsilon}{q}$ fraction of all vertices.
- **Case 2**: Every independent set in *G* contains less than a $\frac{1}{q^{k+1}}$ fraction of all vertices.

Given graph *G* with vertices $[n] = \{1, ..., n\}$ and edges *E*, we will construct a public persuasion instance so that the desired algorithm for approximating the optimal sender utility can be used to distinguish these two cases. Our construction is similar to that in [14]. We let there be *n* receivers, and let $\Theta = [n]$. In other words, both receivers and states of nature correspond to vertices of the graph. We fix the uniform prior distribution over states of nature - i.e., the realized state of nature is a uniformly-drawn vertex in the graph. We define the receiver utilities as follows: $u_i(\theta) = \frac{1}{2}$ if $i = \theta$; $u_i(\theta) = -1$ if $(i, \theta) \in E$; and $u_i(\theta) = -\frac{1}{4n}$ otherwise. We define the sender's utility function, with range [0, 1], to be $f_{\theta}(S) = f(S) = \frac{|S|}{n}$. The following claim is proven in [14].

CLAIM 6.3 ([14]). For any distribution $x \in \Delta_{\Theta}$, the set $S = \{i \in [n] : \sum_{\theta \in \Theta} x_{\theta}u_i(\theta) \ge 0\}$ is an independent set of G.

Claim 6.3 implies that upon receiving any public signal with any posterior distribution *x* over Θ , the players who take action 1 always form an independent set of *G*. Therefore, if the graph *G* is from **Case 2**, the sender's expected utility in any public signaling scheme is at most $\frac{1}{a^{k+1}}$.

Now supposing that *G* is from **Case 1**, we fix the corresponding coloring of $(1 - \epsilon)n$ vertices with colors $k = 1, \ldots, q$, and we use this coloring to construct a public scheme achieving expected sender utility at least $\frac{(1-\epsilon)^2}{q}$. The scheme uses q + 1 signals, and is as follows: if θ has color k then deterministically send the signal k, and if θ is uncolored then deterministically send the signal 0. Given signal k > 0, the posterior distribution on states of nature is the uniform distribution over the vertices with color k – an independent set S_k of size $\frac{1-\epsilon}{q}n$. It is easy to verify that receivers $i \in S_k$ prefer action 1 to action 0, since $\sum_{\theta \in S_k} \frac{1}{|S_k|}u_i(\theta) = \frac{1-\epsilon}{|S_k|}(\frac{1}{2} - \frac{|S_k|-1}{4n}) > \frac{1}{4|S_k|} \ge 0$. Therefore, the sender's utility is $f(S_k) = \frac{|S_k|}{n} = \frac{1-\epsilon}{q}$ whenever k > 0. Since signal 0 has probability ϵ , we conclude that the sender's expected utility is at least $\frac{(1-\epsilon)^2}{q}$, as needed.

Since distinguishing **Case 1** and **Case 2** is NP-hard for arbitrarily large constants k and q, we conclude that it is NP-hard to approximate the optimal sender utility to within any constant factor. Moreover, by setting k = 1, q = 3, we conclude that the sender's utility cannot be approximated additively to within $(1 - \epsilon)^2/3 - 1/3^2 > 1/9$, and thus there is no additive PTAS, unless P=NP.

REFERENCES

- Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. 2010. Correlation Robust Stochastic Optimization. In Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '10). Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1087–1096.
- [2] Ricardo Alonso and Odilon Câmara. 2016. Persuading voters. The American Economic Review 106, 11 (2016), 3590–3605.
- [3] Itai Arieli and Yakov Babichenko. 2016. Private Bayesian persuasion. Available at SSRN 2721307 (2016).
- [4] Yakov Babichenko and Siddharth Barman. 2017. Algorithmic Aspects of Private Bayesian Persuasion. In Proceedings of the 2017 ACM Conference on Innovations in Theoretical Computer Science (ITCS).
- [5] Arjada Bardhi and Yingni Guo. 2016. Modes of Persuasion Toward Unanimous Consent. Working paper (2016).
- [6] Dirk Bergemann and Stephen Morris. 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11, 2 (2016), 487–522.
- [7] U. Bhaskar, Y. Cheng, Y. Kun Ko, and C. Swamy. 2016. Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games. In Proceedings of the 2016 ACM Conference on Economics and Computation (EC). ACM, 18.
- [8] Peter Bro Miltersen and Or Sheffet. 2012. Send Mixed Signals: Earn More, Work Less. In Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12). ACM, New York, NY, USA, 234–247.
- [9] Y. Cai, C. Daskalakis, and S M. Weinberg. 2012. An algorithmic characterization of multi-dimensional mechanisms. In Proceedings of the forty-fourth annual ACM symposium on Theory of computing. ACM, 459–478.
- [10] Y. Cai, C. Daskalakis, and S M. Weinberg. 2012. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS). 130–139.
- [11] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. 2013. Understanding incentives: Mechanism design becomes algorithm design. In Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on. 618–627.
- [12] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. 2011. Maximizing a monotone submodular function subject to a matroid constraint. SIAM J. Comput. 40, 6 (2011), 1740–1766.
- [13] Jimmy Chan, Seher Gupta, Fei Li, and Yun Wang. 2016. Pivotal Persuasion. Available at SSRN (2016).
- [14] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. 2015. Mixture selection, mechanism design, and signaling. In IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 1426–1445.
- [15] Shaddin Dughmi. 2014. On the Hardness of Signaling. In Proceedings of the 55th Symposium on Foundations of Computer Science (FOCS '14). IEEE Computer Society.
- [16] S. Dughmi. 2017. Algorithmic Information Structure Design: A Survey. ACM SIGecom Exchanges 15 (2017), 2–24.
- [17] S. Dughmi, N. Immorlica, and A. Roth. 2014. Constrained Signaling in Auction Design. In Proceedings of the Twenty-five Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). Society for Industrial and Applied Mathematics.
- [18] Shaddin Dughmi and Haifeng Xu. 2016. Algorithmic Bayesian Persuasion. In Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing (STOC '16). ACM, 412–425.
- [19] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2012. Signaling Schemes for Revenue Maximization. In Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12). ACM, 514–531.
- [20] Martin Grötschel, László Lovász, and Alexander Schrijver. 2012. Geometric algorithms and combinatorial optimization. Vol. 2. Springer Science & Business Media.
- [21] Mingyu Guo and Argyrios Deligkas. 2013. Revenue Maximization via Hiding Item Attributes. In Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13). AAAI Press, 157–163.
- [22] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. The American Economic Review 101, 6 (2011), 2590–2615.
- [23] Subhash Khot and Rishi Saket. 2012. Hardness of finding independent sets in almost q-colorable graphs. In IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 380–389.
- [24] Luis Rayo and Ilya Segal. 2010. Optimal Information Disclosure. Journal of Political Economy 118, 5 (2010), pp. 949-987.
- [25] Aviad Rubinstein. 2017. Honest signaling in zero-sum games is hard...and lying is even harder!. In Proceedings of the 44th international colloquium conference on Automata, Languages, and Programming. Springer-Verlag.
- [26] Ina A Taneva. 2015. Information design. (2015).
- [27] Yun Wang. 2015. Bayesian persuasion with multiple receivers. Available at SSRN 2625399 (2015).