

Bayesian Persuasion Meets Mechanism Design: Going Beyond Intractability with Type Reporting

Matteo Castiglioni
Politecnico di Milano
Milan, Italy
matteo.castiglioni@polimi.it

Alberto Marchesi
Politecnico di Milano
Milan, Italy
alberto.marchesi@polimi.it

Nicola Gatti
Politecnico di Milano
Milan, Italy
nicola.gatti@polimi.it

ABSTRACT

Bayesian persuasion studies how an informed *sender* should partially disclose information so as to influence the behavior of self-interested *receivers*. In the last years, a growing attention has been devoted to relaxing the assumption that the sender perfectly knows receiver’s payoffs. The first crucial step towards such an achievement is to study settings where each receiver’s payoffs depend on their unknown *type*, which is randomly determined by a known finite-supported probability distribution. This begets considerable computational challenges, as computing a sender-optimal signaling scheme is inapproximable up to within any constant factor, even in basic settings with a single receiver. In this work, we circumvent this issue by leveraging ideas from mechanism design. In particular, we introduce a *type reporting* step in which the receiver is asked to report their type to the sender, after the latter has committed to a *menu* defining a signaling scheme for each possible receiver’s type. Surprisingly, we prove that, with a single receiver, the addition of this type reporting stage makes the sender’s computational problem tractable. Then, we extend our Bayesian persuasion framework with type reporting to settings with multiple receivers, focusing on the widely-studied case of *no inter-agent externalities* and binary actions. In such setting, we show that it is possible to find a sender-optimal solution in polynomial-time by means of the ellipsoid method, given access to a suitable polynomial-time separation oracle. This can be implemented for supermodular and anonymous sender’s utility functions. As for the case of submodular sender’s utility functions, we first approximately cast the sender’s problem into a linearly-constrained mathematical program whose objective function is the multi-linear extension of the sender’s utility. Then, we show how to find in polynomial-time an approximate solution to the program by means of a continuous greedy algorithm. This provides a $(1 - \frac{1}{e})$ -approximation to the problem, which is tight.

KEYWORDS

Algorithmic Bayesian Persuasion; Equilibrium Computation; Mechanism Design.

ACM Reference Format:

Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. 2022. Bayesian Persuasion Meets Mechanism Design: Going Beyond Intractability with Type Reporting. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAAMAS*, 9 pages.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Bayesian persuasion [26] studies the problem faced by an informed agent (the *sender*) trying to influence the behavior of other self-interested agents (the *receivers*) via the partial disclosure of payoff-relevant information. Agents’ payoffs are determined by the actions played by the receivers and by an exogenous parameter represented as a *state of nature*, which is drawn by a known prior probability distribution and observed by the sender only. The sender commits to a public, randomized information-disclosure policy, which is customarily called *signaling scheme*. In particular, it defines how the sender should send private signals to the receivers, essentially deciding “*who gets to know what*”. These kinds of problems are ubiquitous in applications such as auctions and online advertising [5, 6, 8, 17, 23], voting [1, 11, 12, 15, 18], traffic routing [7, 14, 30], recommendation systems [27], security [28, 33], and product marketing [3, 10].

In the classical Bayesian persuasion model by Kamenica and Gentzkow [26], the sender perfectly knows the payoffs of the receivers. This assumption is unreasonable in practice. Recently, some works tried to relax such an assumption. Castiglioni et al. [13] do that by framing the problem in an online learning framework, focusing on the single-receiver setting. They study the problem in which the sender repeatedly faces a receiver whose type during each iteration is unknown and selected beforehand by an adversary. They design no-regret learning algorithms under full-information and partial-information feedback. However, these algorithms require exponential running time, since even the offline problem in which the receiver’s type is randomly selected according to a known finite-supported probability distribution is NP-hard to approximate up to within any constant factor. Castiglioni et al. [16] consider the problem with multiple receivers, focusing on the classical model with binary actions and no-inter-agent-externalities [2, 3], where each receiver’s payoffs depend only on their action and the state of nature. In the restricted setting in which each receiver has a *constant* number of possible types, they show that the problem is intractable for supermodular and anonymous sender’s utilities and design a no- $(1 - \frac{1}{e})$ -regret polynomial-time algorithm for submodular sender’s utilities. Let us remark that Castiglioni et al. [13] and Castiglioni et al. [16] show that, in their respective settings, the design of polynomial-time no-regret algorithms is impossible due to the NP-hardness of the underlining offline optimization problems in which the distribution over types is known. Hence, the design of *efficient* algorithms for the offline problem is the bottleneck to the design of efficient online learning algorithms.

In this work, we show how to circumvent this problem by leveraging ideas from mechanism design. For the single-receiver setting, we introduce a *type reporting* step in which the receiver is asked to report their type to the sender, after the latter has committed

to a *menu* defining a signaling scheme for each possible receiver’s type. Moreover, we extend the framework to accommodate multiple receivers with binary actions. In such setting, we take advantage of the *no-inter-agent-externalities* assumption to design a type reporting step that is independent among the receivers. In particular, we introduce a type reporting step for each receiver, in which they are asked to report their type to the sender, after the latter has committed to a menu defining a *marginal* signaling scheme for each possible receiver’s type. Then, the sender commits to a signaling scheme that is consistent with all the marginal signaling schemes. By the no-inter-agent-externalities assumption, receivers’ payoffs do *not* depend on such a signaling scheme, and, thus, each receiver’s decision problem in the type reporting step is well defined.

1.1 Original Contribution

In this work, we show that the introduction of a type reporting step makes the sender computational problem tractable.

For the single-receiver case, our main result is to show the existence of an optimal menu of *direct* and *persuasive* signaling schemes. In the classical model in which the sender perfectly knows the receiver’s payoffs, a signaling scheme is direct if signals represent action recommendations, while it is persuasive if the receiver is incentivized to follow recommendations. We extend this definition to menus of signaling schemes. In particular, a menu is direct if the signals used by all the signaling schemes are action recommendations, whereas it is persuasive if the receiver has an incentive to follow action recommendations when they reported their true type. By using this result, an optimal menu of signaling schemes can be computed efficiently by a *linear program* (LP) of polynomial size.

In the multi-receiver setting, we focus on classes of sender’s utility functions that are commonly studied in the literature, namely *supermodular*, *submodular*, and *anonymous* functions [2, 3, 22, 32]. As in the single-receiver case, we show that there always exists an optimal sender’s strategy using a menu of direct and persuasive marginal signaling schemes for each receiver. This allows us to show that an optimal sender’s strategy can be computed by solving an LP with polynomially-many constraints and exponentially-many variables. This is possible in polynomial time by means of the *ellipsoid method*, given access to a suitable polynomial-time separation oracle. Such an oracle can be implemented for supermodular and anonymous sender’s utility functions. In the submodular case, the problem cannot be approximated within any factor better than $1 - \frac{1}{e}$, since our problem generalizes the one without types, which is NP-hard to approximate up to within any factor better than $1 - \frac{1}{e}$ [3]. However, we provide a polynomial-time algorithm that provides a tight $(1 - \frac{1}{e})$ -approximation. To do so, we show how to build a linearly-constrained mathematical program whose objective is the *multi-linear extension* of the sender’s utility, having optimal value arbitrary close to that of an optimal sender’s strategy. Moreover, we show that, from a solution to this program, we can recover in polynomial time a sender’s strategy having in expectation almost the same utility as the optimal value of the program. Finally, we show how to find in polynomial time an approximate solution to the program by means of a *continuous greedy* algorithm. This provides a $(1 - \frac{1}{e})$ -approximation to the problem, which is tight.

1.2 Related Works

Most of the computational works on Bayesian persuasion study models in which the sender knows the receiver’s utility function exactly. Dughmi and Xu [21] initiate these studies with the single-receiver case, while Arieli and Babichenko [2] extend their work to multiple receivers without inter-agent externalities, with a focus on private signaling. In particular, they focus on settings with binary actions for the receivers and a binary space of states of nature. They provide a characterization of an optimal signaling scheme in the case of supermodular, anonymous submodular, and supermajority sender’s utility functions. Arieli and Babichenko [2] extend this latter work by providing tight $(1 - \frac{1}{e})$ -approximate signaling schemes for monotone submodular sender’s utilities and showing that an optimal private signaling scheme for anonymous utility functions can be found efficiently. Dughmi and Xu [22] generalize the previous model to settings with an arbitrary number of states of nature. There are also some works focusing on public signaling with no inter-agent externalities, see, e.g., [22] and [32].

A recent line of research relaxed the assumption that the sender perfectly knows the receivers’ utilities. Castiglioni et al. [13] and Castiglioni et al. [16] study online problems with a single receiver and multiple receivers, respectively. Babichenko et al. [4] study a game with a single receiver and binary actions in which the sender does not know the receiver utility, focusing on the problem of designing a signaling scheme that performs well for any possible receiver’s utility function. Zu et al. [34] relax the perfect knowledge assumption, assuming that the sender and the receiver do not know the prior distribution over the states of nature. They study the problem of computing a sequence of persuasive signaling schemes that achieve small regret with respect to an optimal signaling scheme with knowledge of the prior distribution.

Our problem is also related to *automated mechanism design* [19, 24, 31]. The closest to our work is [20], which studies a mechanism design problem between a mechanism designer and an agent. The agent has a finite number of types and both the agent and the mechanism designer have a utility function that depends on the agent’s type and on an outcome that the designer chooses from a finite set. Moreover, the mechanism designer can commit to a menu specifying an outcome for each reported type. The mechanism designer knows the receiver’s probability distribution over types and their goal is to design an incentive compatible menu in order to maximize their utility. The authors show that it is NP-hard to design an optimal menu, while if the mechanism is allowed to use randomization the problem can be solved in polynomial time.

2 FORMAL MODEL

We formally introduce the Bayesian persuasion framework with type reporting that we study in the rest of this work. In particular, in Subsection 2.1, we describe the model with a single receiver, while in Subsection 2.2 we extend it to multi-receiver settings.

2.1 Model with a Single Receiver

The receiver has a finite set $A := \{a_i\}_{i=1}^{\ell}$ of ℓ available actions and a type chosen from a finite set $K := \{k_i\}_{i=1}^m$ of m possible types. For each type $k \in K$, the receiver’s payoff function is $u^k : A \times \Theta \rightarrow [0, 1]$,

where $\Theta := \{\theta_i\}_{i=1}^d$ is a finite set of d states of nature. We denote by $u_\theta^k(a) \in [0, 1]$ the payoff obtained by the receiver of type $k \in K$ when the state of nature is $\theta \in \Theta$ and they play action $a \in A$. The sender's payoffs are described by the functions $u_\theta^s : A \rightarrow [0, 1]$ for $\theta \in \Theta$. As it is customary in Bayesian persuasion, we assume that the state of nature is drawn from a common prior distribution $\mu \in \text{int}(\Delta_\Theta)$, which is explicitly known to both the sender and the receiver.¹ The sender commits to a *signaling scheme* ϕ , which is a randomized mapping from states of nature to *signals* for the receiver. Formally, $\phi : \Theta \rightarrow \Delta_S$, where S is a set of available signals. For convenience, we let ϕ_θ be the probability distribution employed by the sender to draw signals when the state of nature is $\theta \in \Theta$ and we denote by $\phi_\theta(s)$ the probability of sending signal $s \in S$. Moreover, we slightly abuse the notation and use ϕ to also denote the probability distribution over signals induced by the signaling scheme ϕ and the prior distribution μ .

In the classical Bayesian persuasion framework by Kamenica and Gentzkow [26] (without type reporting), the interaction between the sender and the receiver goes on as follows: (i) the sender commits to a signaling scheme ϕ and the receiver is informed about it; (ii) the sender observes the realized state of nature $\theta \sim \mu$; (iii) the sender draws a signal $s \in S$ according to ϕ_θ and communicates it to the receiver; (iv) the receiver observes s and rationally updates their prior belief over Θ according to the *Bayes rule*; (v) the receiver selects an action maximizing their expected utility.

In step (iv), after observing a signal $s \in S$, the receiver infers a posterior belief $\xi^s \in \Delta_\Theta$ over the states of nature such that the component of ξ^s corresponding to state $\theta \in \Theta$ is:²

$$\xi_\theta^s := \frac{\mu_\theta \Pr_{s' \sim \phi_\theta} \{s' = s\}}{\Pr_{s' \sim \phi} \{s' = s\}}. \quad (1)$$

For the ease of notation, we let $\Xi := \Delta_\Theta$ be the set of receiver's posterior beliefs over states of nature. After computing ξ^s , the receiver plays an action maximizing their utility in ξ^s . As it is customary in the literature [13, 16], we assume that the receiver breaks ties in favor of the sender. In the following, letting $\mathcal{B}_\xi^k := \arg \max_{a \in A} \sum_{\theta \in \Theta} \xi_\theta u_\theta^k(a)$ be the set of actions that maximize the expected utility of the receiver of type $k \in K$ in any posterior $\xi \in \Xi$, we denote by $b_\xi^k \in \arg \max_{a \in \mathcal{B}_\xi^k} \sum_{\theta \in \Theta} \xi_\theta u_\theta^k(a)$ the action in \mathcal{B}_ξ^k that is actually played by the receiver of type k in posterior ξ .

In our *Bayesian persuasion framework with type reporting*, the sender asks the receiver to report their type before observing the realized state of nature. This enables the sender to increase their expected utility. In particular, before the receiver reports their type, the sender proposes to the receiver a *menu* $\Phi = \{\phi^k\}_{k \in K}$ of signaling schemes, committing to send signals according to the signaling scheme ϕ^k if the receiver reports their type to be $k \in K$. In details, the interaction goes on as follows: (i) the sender proposes a menu $\Phi = \{\phi^k\}_{k \in K}$ to the receiver; (ii) the receiver reports a type $k \in K$ that maximizes their expected utility given the proposed menu;

¹ $\text{int}(X)$ is the *interior* of set X and Δ_X is the set of all probability distributions over X . Vectors are highlighted in bold. For any vector \mathbf{x} , the value of its i -th component is x_i .

²We omit the dependency of ξ^s from ϕ as the signaling scheme that is actually used to compute the posterior will be clear from context. Moreover, for the ease of presentation, when we use notation $\Pr \{\cdot\}$ we assume that the set S is finite, so that $\Pr \{\cdot\}$ is well defined. The notation can be easily generalized to the case of infinite sets S .

(iii) the sender observes the realized state of nature $\theta \sim \mu$; (iv) the sender draws a signal $s \in S$ according to ϕ_θ^k and communicates it to the receiver; finally, the interaction terminates with steps (iv) and (v) of the classical setting described above.

Notice that, in step (ii), the receiver of type $k \in K$ can compute their expected utility for each signaling scheme $\phi^{k'}$ in the menu as

$$\sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{s \sim \phi_\theta^{k'}} \left[u_\theta^k \left(b_{\xi^s}^k \right) \right],$$

and, then, they can report a type $k' \in K$ whose corresponding signaling scheme $\phi^{k'}$ maximizes their expected utility.

We focus on menus of signaling schemes that are *incentive compatible* (IC), *i.e.*, in which the receiver of type k is incentivized to report their true type, for any $k \in K$.³ Formally, a menu $\Phi = \{\phi^k\}_{k \in K}$ is IC if, for every type $k \in K$, the following constraints are satisfied:

$$\sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{s \sim \phi_\theta^k} \left[u_\theta^k \left(b_{\xi^s}^k \right) \right] \geq \sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{s \sim \phi_\theta^{k'}} \left[u_\theta^k \left(b_{\xi^s}^k \right) \right] \quad \forall k' \neq k. \quad (2)$$

We say that a signaling scheme is *direct* if $S = A$, which means that signals correspond to action recommendations for the receiver. Moreover, we say that a direct signaling scheme is *persuasive* if the receiver has an incentive to follow the action recommendations that they receive as signals, when they report their true type. It is easy to check that a menu $\Phi = \{\phi^k\}_{k \in K}$ of direct and persuasive signaling schemes is IC if

$$\sum_{a \in A} \sum_{\theta \in \Theta} \mu_\theta \phi_\theta^k(a) u_\theta^k(a) \geq \sum_{a \in A} \max_{a' \in A} \sum_{\theta \in \Theta} \mu_\theta \phi_\theta^{k'}(a) u_\theta^k(a') \quad \forall k' \neq k. \quad (3)$$

In the rest of this work, we will use the well-known equivalence between signaling schemes and distributions over receiver's posteriors (see [25] for further details). In particular, a signaling scheme ϕ is equivalent to a probability distribution $\gamma \in \Delta_\Xi$ over posteriors such that $\mathbb{E}_{\xi \sim \gamma} [\xi] = \mu$, so that the expected utility of the receiver of type $k \in K$ under the signaling scheme can be written as $\mathbb{E}_{\xi \sim \gamma} \left[\sum_{\theta \in \Theta} \xi_\theta u_\theta^k \left(b_\xi^k \right) \right]$. Finally, when the distribution $\gamma \in \Delta_\Xi$ has finite support, we denote by γ_ξ the probability of $\xi \in \Xi$ in γ .

2.2 Model with Multiple Receivers

In a multi-receiver setting, there is a finite set $\mathcal{R} := \{r_i\}_{i=1}^n$ of n receivers, and each receiver $r \in \mathcal{R}$ has a type chosen from a finite set $\mathcal{K}_r := \{k_{r,i}\}_{i=1}^{m_r}$ of m_r different types. We introduce $\mathcal{K} := \times_{r \in \mathcal{R}} \mathcal{K}_r$ as the set of type profiles, which are tuples $\mathbf{k} \in \mathcal{K}$ defining a type $k_r \in \mathcal{K}_r$ for each receiver $r \in \mathcal{R}$. Each receiver $r \in \mathcal{R}$ has two actions available, defined by $\mathcal{A}_r := \{a_0, a_1\}$. We let $\mathcal{A} := \times_{r \in \mathcal{R}} \mathcal{A}_r$ be the set of action profiles specifying an action for each receiver. The payoff of a receiver depends on the action played by them, while it does *not* depend on the actions played by the other receivers, since we assume that there are *no inter-agent externalities*. Formally, a receiver $r \in \mathcal{R}$ of type $k \in \mathcal{K}_r$ has a payoff function $u^{r,k} : \mathcal{A} \times \Theta \rightarrow [0, 1]$. The sender's payoffs depend on the actions played by all the receivers, and they are defined by $u^s : \mathcal{A} \times \Theta \rightarrow [0, 1]$. For the ease of presentation, for every state of nature $\theta \in \Theta$, we introduce the function $f_\theta : 2^{\mathcal{R}} \rightarrow [0, 1]$ such that $f_\theta(R)$ represents the sender's payoff when the state of nature is θ

³Notice that, by a revelation-principle-style argument (see the book by Shoham and Leyton-Brown [29] for some examples of these kind of arguments), focusing on IC menus of signaling schemes is w.l.o.g. when looking for a sender-optimal menu.

and all the receivers in $R \subseteq \mathcal{R}$ play action a_1 , while the others play a_0 . In the rest of this work, we assume that the sender's payoffs are *monotone non-decreasing* in the set of receivers playing a_1 . Formally, for each state $\theta \in \Theta$, we let $f_\theta(R) \leq f_\theta(R')$ for every $R \subseteq R' \subseteq \mathcal{R}$, while $f_\theta(\emptyset) = 0$ for the ease of presentation. As it is customary, we focus on three families of functions: *submodular*, *supermodular*, and *anonymous*. We say that f_θ is submodular, respectively supermodular, if for $R, R' \subseteq \mathcal{R}$: $f_\theta(R \cap R') + f_\theta(R \cup R') \leq f_\theta(R) + f_\theta(R')$, respectively $f_\theta(R \cap R') + f_\theta(R \cup R') \geq f_\theta(R) + f_\theta(R')$. The function f_θ is anonymous if $f_\theta(R) = f_\theta(R')$ for all $R, R' \subseteq \mathcal{R} : |R| = |R'|$.

With multiple receivers, the sender must send a signal to each of them. In this work, we focus on *private* signaling, where each receiver has their own signal that is privately communicated to them. Formally, there is a set \mathcal{S}_r of possible signals for each receiver $r \in \mathcal{R}$. Then, $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$ is a signaling scheme, where $\mathcal{S} := \times_{r \in \mathcal{R}} \mathcal{S}_r$ is the set of signal profiles, which are tuples $\mathbf{s} \in \mathcal{S}$ defining a signal $s_r \in \mathcal{S}_r$ for each receiver $r \in \mathcal{R}$. We denote by ϕ_θ the probability distribution over signal profiles corresponding to state $\theta \in \Theta$, while we let $\phi_\theta(\mathbf{s})$ be the probability of sending $\mathbf{s} \in \mathcal{S}$. Given a signaling scheme ϕ , we define the resulting *marginal signaling scheme* for receiver $r \in \mathcal{R}$ as $\phi^r : \Theta \rightarrow \mathcal{S}_r$. Formally, for every $s \in \mathcal{S}_r$, it holds that $\phi_\theta^r(s) = \Pr_{\mathbf{s} \sim \phi_\theta} \{s_r = s\}$. Notice that receiver r 's posterior beliefs and expected utilities only depend on the marginal signaling scheme ϕ^r .

The interaction between the sender and the receivers goes on as follows: (i) the sender proposes to each receiver $r \in \mathcal{R}$ a menu of marginal signaling schemes $\Phi^r = \{\phi^{r,k}\}_{k \in \mathcal{K}_r}$; (ii) each receiver $r \in \mathcal{R}$ reports a type $k_r \in \mathcal{K}_r$ such that ϕ^{r,k_r} is the marginal signaling scheme maximizing their expected utility; (iii) the sender commits to a signaling scheme ϕ whose resulting marginal signaling schemes ϕ^r are such that $\phi^r := \phi^{r,k_r}$ for all $r \in \mathcal{R}$; (iv) the sender observes the realized state of nature $\theta \sim \mu$ and draws a signal profile $\mathbf{s} \sim \phi_\theta$; (v) each receiver $r \in \mathcal{R}$ observes their signal s_r , rationally updates their prior belief over Θ according to the *Bayes* rule, and selects an action maximizing their expected utility. Notice that the sender only needs to propose marginal signaling schemes to the receivers (rather than general ones), since the expected utility of each receiver only depends on their marginal signaling scheme, and *not* on the others. Thus, the sender can delay the choice of the (general) signaling scheme after types have been reported.

As customary, we assume that the receivers break ties in favor of the sender. Since functions f_θ are monotone, this amounts to play a_1 whenever indifferent between the two actions. Moreover, we say that a signaling scheme is *direct* and *persuasive* if $S = \mathcal{A}$ and the receivers are better off playing recommended actions. We denote with $R \subseteq \mathcal{R}$ the direct signal profile in which it is recommended to play a_1 to all the receiver in R and a_0 to all the receiver in $\mathcal{R} \setminus R$.

Similarly to the single-receiver case, we restrict the attention to IC menu of marginal signaling schemes. Thus, in a multi-receiver setting, a sender's strategy is composed by an IC menu of marginal signaling scheme $\Phi^r = \{\phi^{r,k}\}_{k \in \mathcal{K}_r}$ for each receiver $r \in \mathcal{R}$, and a set of signaling schemes $\{\phi^k\}_{k \in \mathcal{K}}$ (one per type profile possibly reported by the receivers) such that the resulting marginal signaling schemes satisfy $\phi^{k,r} = \phi^{r,k_r}$ for all $k \in \mathcal{K}$ and $r \in \mathcal{R}$.

2.3 Sender's Computational Problems

We consider the computational problem in which, given the probability distribution over the receivers' types, the sender wants to maximize their expected utility. In the single-receiver case, the receiver's type $k \in K$ is drawn from a known distribution $\lambda \in \Delta_K$. We call MENU-SINGLE the problem of computing an IC menu of signaling schemes $\Phi = \{\phi^k\}_{k \in K}$ that maximizes the sender's expected utility, given a probability distribution $\lambda \in \Delta_K$ as input. In the multi-receiver case, the types profiles $\mathbf{k} \in \mathcal{K}$ are drawn from a known distribution $\lambda \in \Delta_{\mathcal{K}}$, where $\mathcal{K} \subseteq \mathcal{K}$ is a subset of possible types vectors, *i.e.*, the support of λ . We call MENU-MULTI the problem of computing a sender's strategy—made by an IC menu of marginal signaling schemes $\Phi^r = \{\phi^{r,k}\}_{k \in \mathcal{K}_r}$ for each receiver $r \in \mathcal{R}$ and a set of signaling schemes $\{\phi^k\}_{k \in \mathcal{K}}$ —that maximizes the sender's expected utility, given $\lambda \in \Delta_{\mathcal{K}}$ as input.⁴

3 SINGLE-RECEIVER PROBLEM

We show how to solve MENU-SINGLE in polynomial time.

By using the equivalence between signaling schemes and distributions over posteriors (see Section 2.1), it is easy to check that an optimal menu of signaling schemes can be computed by the following LP 4 with an *infinite* number of variables, namely $\gamma^k \in \Delta_{\Xi}$ for $k \in K$. In LP 4, the objective is the sender's expected utility assuming the receiver reports their true type, the first set of constraints encodes IC conditions, while the last one ensures that the distributions over posteriors correctly represent signaling schemes.

$$\begin{aligned} \max_{\gamma} \quad & \sum_{k \in K} \lambda_k \mathbb{E}_{\xi \sim \gamma^k} \sum_{\theta \in \Theta} \xi_\theta u_\theta^s \left(b_\xi^k \right) \quad \text{s.t.} \quad (4) \\ & \mathbb{E}_{\xi \sim \gamma^k} \left[\sum_{\theta \in \Theta} \xi_\theta u_\theta^k \left(b_\xi^k \right) \right] \geq \mathbb{E}_{\xi \sim \gamma^{k'}} \left[\sum_{\theta \in \Theta} \xi_\theta u_\theta^k \left(b_\xi^k \right) \right] \quad \forall k \neq k' \in K \\ & \mathbb{E}_{\xi \sim \gamma^k} [\xi_\theta] = \mu_\theta \quad \forall \theta \in \Theta, \forall k \in K \\ & \gamma^k \in \Delta_{\Xi} \quad \forall k \in K. \end{aligned}$$

As a first step, we show that there always exists an optimal solution to LP 4 in which the probability distributions $\gamma^k \in \Delta_{\Xi}$ have finite support. This allows us to compute an optimal menu of signaling schemes by solving an LP with a *finite* number of variables. In the following, for every $k \in K$ and $a \in A$, let $\Xi^{k,a} := \{\xi \in \Xi : a \in \mathcal{B}_\xi^k\}$ and $\hat{\Xi}^{k,a} := \{\xi \in \Xi : a = b_\xi^k\}$. Moreover, for every $\mathbf{a} \in \times_{k \in K} A$, let $\Xi^{\mathbf{a}} := \bigcap_{k \in K} \Xi^{k,a_k}$ and $\hat{\Xi}^{\mathbf{a}} := \bigcap_{k \in K} \hat{\Xi}^{k,a_k}$, where a_k is the k -th component of \mathbf{a} . Finally, let Ξ^* be such that $\Xi^* := \bigcup_{\mathbf{a} \in \times_{k \in K} A} V(\Xi^{\mathbf{a}})$, where $V(\Xi^{\mathbf{a}})$ denotes the set of vertices of the polytope $\Xi^{\mathbf{a}}$. The following Lemma 1 shows that there always exists an optimal menu of signaling schemes that can be encoded as probability distributions over Ξ^* . Formally, the lemma is proved by showing that the following LP 5 is equivalent to LP 4.

$$\max_{\gamma} \quad \sum_{k \in K} \lambda_k \sum_{\xi \in \Xi^*} \gamma_\xi^k \sum_{\theta \in \Theta} \xi_\theta u_\theta^s \left(b_\xi^k \right) \quad \text{s.t.} \quad (5a)$$

$$\sum_{\xi \in \Xi^*} \gamma_\xi^k \sum_{\theta \in \Theta} \xi_\theta u_\theta^k \left(b_\xi^k \right) \geq \sum_{\xi \in \Xi^*} \gamma_\xi^{k'} \sum_{\theta \in \Theta} \xi_\theta u_\theta^k \left(b_\xi^k \right) \quad \forall k \neq k' \in K \quad (5b)$$

⁴A polynomial-time algorithm for MENU-MULTI must run in time polynomial in the size of the instance and in the size of the support of the distribution λ . Notice that, in general, the latter may be exponential in the number of receivers n .

$$\sum_{\xi \in \Xi^*} \gamma_{\xi}^k \xi_{\theta} = \mu_{\theta} \quad \forall k \in K, \forall \theta \in \Theta \quad (5c)$$

$$\sum_{\xi \in \Xi^*} \gamma_{\xi}^k = 1 \quad \forall k \in K. \quad (5d)$$

Intuitively, the result is shown by noticing that, once fixed the receiver's best responses to $\mathbf{a} \in \times_{k \in K} A$, the sums over Θ in the objective and the constraints of LP 4 are linear in the posterior ξ , which allows to apply Carathéodory theorem to replace each posterior with a probability distributions over the vertices of Ξ^* .

Lemma 1. *In single-receiver instances, there always exists a sender-optimal menu of signaling schemes that can be encoded as probability distributions over the finite set of posteriors Ξ^* .*

Next, we show that there always exists an optimal menu of direct and persuasive signaling schemes, and that it can be computed in polynomial time by solving a polynomially-sized LP obtained by further simplifying LP 5 (Theorem 1). Notice that, in a Bayesian persuasion problem without type reporting, an optimal signaling scheme must employ a signal for each action profile $\mathbf{a} \in \times_{k \in K} A$. Since these profiles are exponentially many, an optimal direct and persuasive signaling scheme cannot be computed in polynomial time by linear programming. Indeed, without typer reporting, the problem has been shown to be NP-hard [13].

An intuition behind the proof of Theorem 1 is provided in the following. Fix type $k \in K$ and action $a \in A$. Suppose that an optimal menu of signaling schemes employs $\gamma^k \in \Delta_{\Xi^*}$ for the type k , and that γ^k has in the support two posteriors $\xi^1, \xi^2 \in \Xi^{k,a}$ with probabilities $\gamma_{\xi^1}^k$ and $\gamma_{\xi^2}^k$. Consider a new signaling scheme that replaces the two posteriors ξ^1 and ξ^2 with their convex combination $\xi^* \in \Delta_{\Xi^*}$, so that

$$\xi_{\theta}^* = \frac{\gamma_{\xi^1}^k \xi_{\theta}^1 + \gamma_{\xi^2}^k \xi_{\theta}^2}{\gamma_{\xi^1}^k + \gamma_{\xi^2}^k} \text{ for every } \theta \in \Theta \text{ and } \gamma_{\xi^*}^k = \gamma_{\xi^1}^k + \gamma_{\xi^2}^k.$$

Both ξ^1 and ξ^2 induce the same best response of the receiver of type k , and Objective (5a) and Constraints (5c) are linear in ξ . Hence, replacing the two posteriors with their convex combination ξ^* preserves the value of the objective, while maintaining the constraints satisfied. The same does *not* hold for Constraints (5b), which are linear in the posterior only if we fix the best responses of all the receiver's types. For Constraints (5b), if we consider an inequality in which γ^k appears in the left hand side, the sum over Θ is linear in ξ and

$$\gamma_{\xi^1}^k \sum_{\theta \in \Theta} \xi_{\theta}^1 u_{\theta}^k(a) + \gamma_{\xi^2}^k \sum_{\theta \in \Theta} \xi_{\theta}^2 u_{\theta}^k(a) = \gamma_{\xi^*}^k \sum_{\theta \in \Theta} \xi_{\theta}^* u_{\theta}^k(a).$$

Instead, if γ^k appears in the right hand side, by the convexity of the max operator it holds:

$$\begin{aligned} & \gamma_{\xi^1}^k \max_{a' \in A} \sum_{\theta \in \Theta} \xi_{\theta}^1 u_{\theta}^k(a') + \gamma_{\xi^2}^k \max_{a' \in A} \sum_{\theta \in \Theta} \xi_{\theta}^2 u_{\theta}^k(a') \\ & \geq \max_{a' \in A} \left[\gamma_{\xi^1}^k \sum_{\theta \in \Theta} \xi_{\theta}^1 u_{\theta}^k(a') + \gamma_{\xi^2}^k \sum_{\theta \in \Theta} \xi_{\theta}^2 u_{\theta}^k(a') \right] \\ & = \max_{a' \in A} \sum_{\theta \in \Theta} \xi_{\theta}^* u_{\theta}^k(a'). \end{aligned}$$

Therefore, if we replace two posteriors that induce the same receiver's best responses with their convex combination, the left hand side of Constraints (5b) is preserved, while the value of the right hand side can only decrease, guaranteeing that Constraints (5b) remain satisfied. By using this idea, we can join all the posteriors that induce the same best responses. Finally, by resorting to the equivalence between signaling schemes and distributions over, we obtain the following LP 6 of polynomial size. Hence, an optimal menu of signaling schemes can be computed in polynomial time.

$$\max_{\phi, l} \sum_{k \in K} \lambda_k \sum_{\theta \in \Theta} \mu_{\theta} \sum_{a \in A} \phi_{\theta}^k(a) u_{\theta}^s(a) \quad \text{s.t.} \quad (6a)$$

$$\sum_{a \in A} \sum_{\theta \in \Theta} \mu_{\theta} \phi_{\theta}^k(a) u_{\theta}^k(a) \geq \sum_{a \in A} l_a^{k,k'} \quad \forall k \neq k' \in K \quad (6b)$$

$$l_a^{k,k'} \geq \sum_{\theta \in \Theta} \mu_{\theta} \phi_{\theta}^k(a) u_{\theta}^k(a') \quad \forall k \neq k' \in K, \forall a, a' \in A \quad (6c)$$

$$\sum_{\theta \in \Theta} \mu_{\theta} \phi_{\theta}^{k'}(a) u_{\theta}^k(a) \geq \sum_{\theta \in \Theta} \mu_{\theta} \phi_{\theta}^{k'}(a) u_{\theta}^k(a') \quad \forall k \in K, \forall a, a' \in A \quad (6d)$$

$$\sum_{a \in A} \phi_{\theta}^k(a) = 1 \quad \forall k \in K, \forall \theta \in \Theta. \quad (6e)$$

Notice that Constraints (6b) and (6c) are equivalent to the IC constraints for direct and persuasive signaling schemes, which are those specified in Equation (3), where $\max_{a' \in A} \sum_{\theta \in \Theta} \mu_{\theta} \phi_{\theta}^{k'}(a) u_{\theta}^k(a')$ is the best response of the receiver of type $k \in K$ to the direct signal a for the receiver of type $k' \in K$. Moreover, Constraints (6d) force the signaling schemes to be persuasive.

Theorem 1. *In single-receiver instances, there always exists an optimal menu of direct and persuasive signaling schemes. Moreover, it can be computed in polynomial time.*

4 MULTI-RECEIVER PROBLEM

In this section, we switch the attention to MENU-MULTI. As we will show in the following (Theorem 2), given any multi-receiver instance, there always exists an optimal sender's strategy that uses menus of direct and persuasive marginal signaling schemes. This allows us to formulate the sender's problem as the following LP 7, which will be crucial for the results in the rest of this section.

Since $\phi_{\theta}^k(a_0) = 1 - \phi_{\theta}^k(a_1)$ for every $r \in \mathcal{R}$, $k \in \mathcal{K}_r$, and $\theta \in \Theta$, by letting $x_{\theta}^{r,k} = \phi_{\theta}^k(a_1)$ we can formulate the following LP:

$$\max_{\phi \geq 0, x \geq 0} \sum_{\theta \in \Theta} \mu_{\theta} \sum_{\mathbf{k} \in \mathcal{K}} \lambda_{\mathbf{k}} \sum_{R \subseteq \mathcal{R}} \phi_{\theta}^{\mathbf{k}}(R) f_{\theta}(R) \quad \text{s.t.} \quad (7a)$$

$$\sum_{R \subseteq \mathcal{R}: r \in R} \phi_{\theta}^{\mathbf{k}}(R) = x_{\theta}^{r,k_r} \quad \forall \mathbf{k} \in \mathcal{K}, \forall r \in \mathcal{R}, \forall \theta \in \Theta \quad (7b)$$

$$\begin{aligned} & \sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k} u_{\theta}^{r,k}(a_1) + \sum_{\theta \in \Theta} \mu_{\theta} (1 - x_{\theta}^{r,k}) u_{\theta}^{r,k}(a_0) \\ & \geq l_{a_1}^{r,k,k'} + l_{a_0}^{r,k,k'} \quad \forall r \in \mathcal{R}, \forall k \neq k' \in \mathcal{K}_r \quad (7c) \end{aligned}$$

$$\begin{aligned} & l_{a_1}^{r,k,k'} \geq \sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k'} u_{\theta}^{r,k}(a) \\ & \quad \forall r \in \mathcal{R}, \forall a \in \mathcal{A}_r, \forall k \neq k' \in \mathcal{K}_r \quad (7d) \end{aligned}$$

$$\begin{aligned} & l_{a_0}^{r,k,k'} \geq \sum_{\theta \in \Theta} \mu_{\theta} (1 - x_{\theta}^{r,k'}) u_{\theta}^{r,k}(a) \\ & \quad \forall r \in \mathcal{R}, \forall a \in \mathcal{A}_r, \forall k \neq k' \in \mathcal{K}_r \quad (7e) \end{aligned}$$

$$\sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k} \left[u_{\theta}^{r,k}(a_1) - u_{\theta}^{r,k}(a_0) \right] \geq 0 \quad \forall r \in \mathcal{R}, \forall k \in \mathcal{K}_r \quad (7f)$$

$$\sum_{\theta \in \Theta} \mu_{\theta} \left(1 - x_{\theta}^{r,k} \right) \left[u_{\theta}^{r,k}(a_0) - u_{\theta}^{r,k}(a_1) \right] \geq 0 \quad \forall r \in \mathcal{R}, \forall k \in \mathcal{K}_r \quad (7g)$$

$$\sum_{R \subseteq \mathcal{R}} \phi_{\theta}^k(R) = 1 \quad \forall k \in \bar{\mathcal{K}}, \forall \theta \in \Theta. \quad (7h)$$

In the LP, Constraints (7b) represent consistency conditions ensuring that the general signaling scheme ϕ^k results in the marginal signaling schemes $\phi^{k,r}$, which are defined by means of variables $x_{\theta}^{r,k}$. Constraints (7c), (7d), and (7e) represent IC constraints for the menus of marginal signaling schemes, where, as in LP 6, we use Constraints (7d) and (7e) with variables $l_{a_0}^{r,k,k'}$, $l_{a_1}^{r,k,k'}$ to compute receivers' expected utilities of playing a best response. Finally, Constraints (7f) and (7g) encode the persuasiveness conditions, while Constraints (7h) require the signaling scheme be well defined.

Next, we prove our main existence result supporting LP 7.

Theorem 2. *In multi-receiver instances, there always exists an optimal sender's strategy that uses menus of direct and persuasive marginal signaling schemes.*

4.1 Supermodular/Anonymous Sender's Utility

LP 7 has an exponential number of variables and a polynomial number of constraints. Nevertheless, we show that it is possible to apply the ellipsoid algorithm to its dual formulation in polynomial time, provided access to a suitably-defined separation oracle.

Theorem 3. *Given access to an oracle that solves $\max_{R \subseteq \mathcal{R}} f_{\theta}(R) + \sum_{r \in R} w_r$ for any $\mathbf{w} \in \mathbb{R}^n$, there exists a polynomial-time algorithm that finds an optimal sender's strategy in any multi-receiver instance.*

An oracle that solves $\max_{R \subseteq \mathcal{R}} f_{\theta}(R) + \sum_{r \in R} w_r$ can be implemented in polynomial time for supermodular and anonymous functions, as shown by Dughmi and Xu [22]. As a consequence, we obtain the following corollary.

Corollary 1. *In multi-receiver instances with supermodular or anonymous sender's utility functions, there exists a polynomial-time algorithm that computes an optimal sender's strategy.*

4.2 Submodular Sender's Utility

In this section, we show how to obtain in polynomial time a $\left(1 - \frac{1}{e}\right)$ -approximation to an optimal sender's strategy in instances with submodular utility functions, modulo an additive loss $\epsilon > 0$. This is the best approximation result that can be achieved in polynomial time, since, as it follows from results in the literature, it is NP-hard to obtain an approximation factor better than $1 - \frac{1}{e}$. Indeed, if we consider settings without types, *i.e.*, in which $|\mathcal{K}_r| = 1$ for all $r \in \mathcal{R}$, the problem reduces to computing an optimal signaling scheme when the sender knows receivers' utilities. Then, in the restricted case in which there are only two states of nature, Babichenko and Barman [3] show that, for each $\epsilon > 0$, it is NP-hard to provide a $\left(1 - \frac{1}{e} + \epsilon\right)$ -approximation of an optimal signaling scheme.

Then, the following theorem provides a tight approximation algorithm that runs in polynomial time.

Theorem 4. *For each $\epsilon > 0$, there exists an algorithm with running time polynomial in the instance size and $\frac{1}{\epsilon}$ that returns a sender's strategy with utility at least $\left(1 - \frac{1}{e}\right)OPT - \epsilon$ in expectation, where OPT is the sender's expected utility in an optimal strategy.*

In order to prove the result, we reduce the problem of computing the desired (approximate) sender's strategy to solving the following linearly-constrained mathematical program (Program 8). The program exploits the fact that, as we will show next, there always exists an "almost" optimal sender's strategy in which the sender employs signaling schemes ϕ^k (for $k \in \mathcal{K}$) such that the distributions ϕ_{θ}^k are q -uniform over the set $2^{\mathcal{R}}$. In particular, we say that a distribution is q -uniform if it follows a uniform distribution on a multiset of size q , where we denote by $[q]$ the set $\{1, \dots, q\}$. Then, the mathematical program reads as follows.

$$\max_x \sum_{\theta \in \Theta} \mu_{\theta} \sum_{k \in \mathcal{K}} \lambda_k \frac{1}{q} \sum_{j \in [q]} F_{\theta}(x^{j,k,\theta}) \quad \text{s.t.} \quad (8a)$$

$$\sum_{j \in [q]} \frac{1}{q} x_r^{j,k,\theta} \leq x_{\theta}^{r,k} \quad \forall r \in \mathcal{R}, \forall k \in \bar{\mathcal{K}}, \forall \theta \in \Theta \quad (8b)$$

$$\sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k} u_{\theta}^{r,k}(a_1) + \sum_{\theta \in \Theta} \mu_{\theta} \left(1 - x_{\theta}^{r,k} \right) u_{\theta}^{r,k}(a_0) \geq l_{a_1}^{k,k'} + l_{a_0}^{k,k'} \quad \forall r \in \mathcal{R}, \forall k \neq k' \in \mathcal{K}_r \quad (8c)$$

$$l_{a_1}^{k,k'} \geq \sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k'} u_{\theta}^{r,k'}(a) \quad \forall r \in \mathcal{R}, \forall a \in \mathcal{A}_r, \forall k \neq k' \in \mathcal{K}_r \quad (8d)$$

$$l_{a_0}^{k,k'} \geq \sum_{\theta \in \Theta} \mu_{\theta} \left(1 - x_{\theta}^{r,k'} \right) u_{\theta}^{r,k'}(a) \quad \forall r \in \mathcal{R}, \forall a \in \mathcal{A}_r, \forall k \neq k' \in \mathcal{K}_r \quad (8e)$$

$$\sum_{\theta \in \Theta} \mu_{\theta} x_{\theta}^{r,k} \left[u_{\theta}^{r,k}(a_1) - u_{\theta}^{r,k}(a_0) \right] \geq 0 \quad \forall r \in \mathcal{R}, k \in \mathcal{K}_r \quad (8f)$$

$$\sum_{\theta \in \Theta} \mu_{\theta} \left(1 - x_{\theta}^{r,k} \right) \left[u_{\theta}^{r,k}(a_0) - u_{\theta}^{r,k}(a_1) \right] \geq 0 \quad \forall r \in \mathcal{R}, \forall k \in \mathcal{K}_r \quad (8g)$$

$$0 \leq x_{\theta}^{r,k} \leq 1 \quad \forall r \in \mathcal{R}, \forall k \in \mathcal{K}_r, \forall \theta \in \Theta \quad (8h)$$

$$0 \leq x_r^{j,k,\theta} \leq 1 \quad \forall j \in [q], \forall r \in \mathcal{R}, \forall k \in \mathcal{K}_r, \forall \theta \in \Theta. \quad (8i)$$

In Program 8, each variable $x_{\theta}^{r,k}$ represents the probability $\phi_{\theta}^k(a_1)$ that the sender recommends action a_1 to receiver $r \in \mathcal{R}$ of type $k \in \mathcal{K}_r$ in state $\theta \in \Theta$. Constraints (8c)–(8h) force the marginal signaling schemes to be well defined, where Constraints (8c), (8d), and (8e) encode the IC conditions, Constraints (8f) and (8g) ensure the persuasiveness property, and Constraints (8h) require the marginal signaling schemes to be feasible, *i.e.*, $\phi_{\theta}^{r,k}(a_1) + \phi_{\theta}^{r,k}(a_0) = 1$ and $\phi_{\theta}^{r,k}(a) \geq 0$ for every $a \in \{a_0, a_1\}$. Moreover, the program uses variables $x_r^{j,k,\theta} \in \{0, 1\}$ to represent whether the recommended action to receiver $r \in \mathcal{R}$ is a_1 or a_0 in the j -th action profile in the support of ϕ_{θ}^k . Notice that we relaxed these variables to $x_r^{j,k,\theta} \in [0, 1]$ and use the multi-linear extension of the sender's utility functions f_{θ} , which, for every $\theta \in \Theta$, reads as

$$F_{\theta}(x) := \sum_{R \subseteq \mathcal{R}} f_{\theta}(R) \prod_{r \in R} x_r \prod_{r \notin R} (1 - x_r).$$

Moreover, we also relax the constraints ensuring the consistency of the marginal signaling schemes, namely Constraints (8b), by replacing the condition $\sum_{j \in [q]} \frac{1}{q} x_r^{j, \mathbf{k}, \theta} = x_\theta^{r, \mathbf{k}, \theta}$ for all $r \in \mathcal{R}, \mathbf{k} \in \bar{\mathcal{K}}, \theta \in \Theta$ with $\sum_{j \in [q]} \frac{1}{q} x_r^{j, \mathbf{k}, \theta} \leq x_\theta^{r, \mathbf{k}, \theta}$ for all $r \in \mathcal{R}, \mathbf{k} \in \bar{\mathcal{K}}, \theta \in \Theta$.

In order to reduce the problem of computing the desired sender's strategy to solving Program 8, we need the following two lemmas (Lemma 2 and Lemma 3). We show that the value of Program 8 for a suitably-defined q approximates the value of an optimal sender's strategy (i.e., an optimal solution to LP 7) and that, given a solution to Program (8), we can recover a sender's strategy with approximately the same expected utility for the sender. Our result is related to those in [22] for the case without types. However, Dughmi and Xu [22] use a probabilistic method to show the existence of an "almost" optimal signaling scheme that uses q -uniform distributions over the signals. This approach cannot be applied to our problem since it slightly modifies the receivers' utilities. In the case of persuasiveness constraints, they show how to maintain feasibility. However, this approach does *not* work for the IC constraints. We propose a different technique based on the fact that LP 7 has a polynomial number of constraints. Let β be the number of constraint of LP 8. Notice that β is polynomial in the size of the LP. We show that, for each $\epsilon > 0$, there exist a q such that LP 8 has value at least $OPT - \epsilon$, where OPT is the value of an optimal sender's strategy.

Lemma 2. *For each $\epsilon > 0$, the optimal value of Program 8 with $q = \left\lceil \frac{\beta}{\epsilon} \right\rceil$ is at least $OPT - \epsilon$, where OPT is the value of an optimal sender's strategy and β is the number of constraints of LP 7.*

Then, we show how to obtain a signaling scheme given a solution of Program 8. Dughmi and Xu [22] build a signaling scheme by using a technique whose generalization to our setting works as follows. Given a state of nature $\theta \in \Theta$ and a vector of types $\mathbf{k} \in \bar{\mathcal{K}}$, it selects a $j \in [q]$ uniformly at random and recommends signal a_1 to receiver $r \in \mathcal{R}$ with probability $x_r^{j, \mathbf{k}, \theta}$, while it recommends a_0 otherwise. By definition of multi-linear extension, using this technique the sender achieves expected utility equal to the value of the given solution to Program 8. However, this signaling scheme uses an exponential number of signal profiles, and, thus, it cannot be represented explicitly. In the following lemma, we show how to obtain a sender's strategy in which signaling schemes use a polynomial number of signal profiles.

Lemma 3. *Given a solution to Program 8 with value APX , for each $\iota > 0$, there exists an algorithm with running time polynomial in the instance size and ι that returns a sender's strategy with utility at least $APX - \frac{n}{q} - \iota$ in expectation. Moreover, such sender's strategy employs signaling schemes using polynomially-many signal profiles.*

PROOF. Let x be a solution to Program 8 with value APX . Next, we show how to obtain the desired sender's strategy.

First, we build a new solution to Program 8 such that Constraints (8b) are satisfied with equality. Since the functions F_θ are monotonic, we can simply obtain such solution by increasing the values of variables $x_r^{j, \mathbf{k}, \theta}$. It is easy to see that, by the monotonicity of F_θ , the objective function does *not* decrease.

Then, we obtain an "almost binary" solution by applying, for every $\theta \in \Theta$ and $\mathbf{k} \in \bar{\mathcal{K}}$, the procedure outlined in Algorithm 1. An

first operation, the algorithm iterates over the receivers, doing the operations described in the following for each receiver $r \in \mathcal{R}$.

For each $j \in [q]$, the algorithm computes an estimate of the following partial derivative

$$\frac{\partial F_\theta(x^{j, \mathbf{k}, \theta})}{\partial x_r^{j, \mathbf{k}, \theta}} = \sum_{R \subseteq \mathcal{R} \setminus \{r\}} [f_\theta(R \cup \{r\}) - f_\theta(R)] \prod_{r' \in R} x_{r'}^{j, \mathbf{k}, \theta} \prod_{r' \notin R, r' \neq r} (1 - x_{r'}^{j, \mathbf{k}, \theta}),$$

This is accomplished by drawing $\sigma = \frac{-8}{\epsilon^2} n^2 \log \frac{p}{2}$ samples of the random variable $f_\theta(\tilde{R} \cup \{r\}) - f_\theta(\tilde{R})$ (with $p = \frac{1}{2|\bar{\mathcal{K}}|dq n}$), where $\tilde{R} \subseteq \mathcal{R}$ is obtained by randomly picking each receiver $r' \in \mathcal{R} : r' \neq r$ independently with probability $x_{r'}^{j, \mathbf{k}, \theta}$. It is easy to see that the expected value of the random variable is exactly equal to value of the partial derivative above. Letting $\tilde{e}_r^{j, \mathbf{k}, \theta}$ be the empirical mean of the samples, by an Hoeffding bound, we get

$$\Pr \left\{ \left| \tilde{e}_r^{j, \mathbf{k}, \theta} - \frac{\partial F_\theta(x^{j, \mathbf{k}, \theta})}{\partial x_r^{j, \mathbf{k}, \theta}} \right| \geq \frac{1}{4n} \right\} \leq p.$$

Moreover, consider the event \mathcal{E} in which $\left| \tilde{e}_r^{j, \mathbf{k}, \theta} - \frac{\partial F_\theta(x^{j, \mathbf{k}, \theta})}{\partial x_r^{j, \mathbf{k}, \theta}} \right| \leq \frac{1}{4n}$ for all $j \in [q], \mathbf{k} \in \bar{\mathcal{K}}, \theta \in \Theta$, and $r \in \mathcal{R}$. By a union bound, the event \mathcal{E} holds with probability at least $1 - p|\bar{\mathcal{K}}|dq n$.

As a second step, the algorithm re-labels the indexes so that, if $j < j'$, then $\tilde{e}_r^{j, \mathbf{k}, \theta} \geq \tilde{e}_r^{j', \mathbf{k}, \theta}$. Notice that the value of the partial derivative with respect to $x_r^{j, \mathbf{k}, \theta}$ does *not* depend on its value. Hence, given two indexes $j < j'$, by "moving" a value t from $x_r^{j', \mathbf{k}, \theta}$ to $x_r^{j, \mathbf{k}, \theta}$, the sum $\sum_{j \in [q]} F_\theta(x_r^{j, \mathbf{k}, \theta})$ decreases at most of

$$t \left(\frac{\partial F_\theta(x^{j', \mathbf{k}, \theta})}{\partial x_r^{j', \mathbf{k}, \theta}} - \frac{\partial F_\theta(x^{j, \mathbf{k}, \theta})}{\partial x_r^{j, \mathbf{k}, \theta}} \right) \leq t \left(\frac{1}{2n} + \tilde{e}_r^{j', \mathbf{k}, \theta} - \tilde{e}_r^{j, \mathbf{k}, \theta} \right) \leq t \frac{1}{2n}.$$

Let $Q^{\mathbf{k}, \theta, r} = \{1, \dots, j^*\}$ be the set of the $j^* = \left\lfloor \sum_{j \in [q]} x_r^{j, \mathbf{k}, \theta} \right\rfloor$ smallest indexes in $[q]$. Then, the algorithm updates the solution x by setting $x_r^{j, \mathbf{k}, \theta} = 1$ for all indexes $j \in Q^{\mathbf{k}, \theta, r}$ and setting

$$x_r^{j^*, \mathbf{k}, \theta} = \sum_{j' \in Q^{\mathbf{k}, \theta, r}} x_r^{j', \mathbf{k}, \theta} - \left\lfloor \sum_{j' \in Q^{\mathbf{k}, \theta, r}} x_r^{j', \mathbf{k}, \theta} \right\rfloor.$$

After having iterated over all the receivers, the algorithm has built a new feasible solution \bar{x} to Program 8 such that

$$\sum_{j \in [q]} \left[F_\theta(\bar{x}_r^{j, \mathbf{k}, \theta}) - F_\theta(x_r^{j, \mathbf{k}, \theta}) \right] \geq -qt/2,$$

since the algorithm moved at most a value q from variables indexed by j' to variables indexed by $j < j'$. Moreover, each receiver $r \in \mathcal{R}$ has at most a non-binary element among variables $\bar{x}_r^{j, \mathbf{k}, \theta}$.

As a final step, the algorithm first builds a set $Q^{\mathbf{k}, \theta}$ of indexes $j \in [q]$ such that $\bar{x}_r^{j, \mathbf{k}, \theta}$ is a binary vector. Notice that there always exists one such set $Q^{\mathbf{k}, \theta}$ of size at least $q - n$. Then, the algorithm constructs a signaling scheme such that

$$\phi_\theta^{\mathbf{k}}(R) = \frac{1}{q} \left\{ \left\{ j \in Q^{\mathbf{k}, \theta} : x_r^{j, \mathbf{k}, \theta} = 1 \forall r \in R, x_r^{j, \mathbf{k}, \theta} = 0 \forall r \notin R \right\} \right\}.$$

Notice that $\sum_{R \in \mathcal{R}: r \in R} \phi_\theta^k(R) \leq x^{r, k_r}$ and, by the monotonicity assumption on f_θ , it is easy to build a signaling scheme such that $\sum_{R \in \mathcal{R}: r \in R} \phi_\theta^k(R) = x^{r, k_r}$ with greater sender's expected utility. Finally, the algorithm outputs the sender's strategy made by $\{\phi^k\}_{k \in \bar{\mathcal{K}}}$ and $\{x^{r, k}\}_{r \in \mathcal{R}, k \in \mathcal{K}_r}$, where the menus of marginal signaling schemes are those given as input.

To conclude the proof, we show that the utility of the sender's strategy described above is at least $APX - \frac{n}{q} - \iota$ in expectation. If the event \mathcal{E} holds, the utility of the solution is at least

$$\begin{aligned}
& \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \phi_\theta^k(R) f_\theta(R) \\
& \geq \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \frac{1}{q} \sum_{j \in Q^{k, \theta}} F_\theta(\bar{x}^{j, k, \theta}) \\
& \geq \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \frac{1}{q} \left[\sum_{j \in [q]} F_\theta(\bar{x}^{j, k, \theta}) - n \right] \\
& \geq \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \frac{1}{q} \left[\sum_{j \in [q]} F_\theta(x^{j, k, \theta}) - n - \frac{\iota q}{2} \right] \\
& = \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \left[\frac{1}{q} \sum_{j \in [q]} F_\theta(x^{j, k, \theta}) - \frac{n}{q} - \iota/2 \right] \\
& \geq \sum_{\theta \in \Theta} \mu_\theta \sum_{k \in \bar{\mathcal{K}}} \lambda_k \frac{1}{q} \sum_{j \in [q]} F_\theta(x^{j, k, \theta}) - \frac{n}{q} - \iota/2 \\
& = APX - \frac{n}{q} - \iota/2.
\end{aligned}$$

Hence, the sender's expected utility is at least

$$\begin{aligned}
\Pr\{\mathcal{E}\} \left(APX - \frac{n}{q} - \iota/2 \right) & \geq (1 - p|\bar{\mathcal{K}}|dq n) \left(APX - \frac{n}{q} - \frac{\iota}{2} \right) \\
& \geq APX - \frac{n}{q} - \frac{\iota}{2} - p|\bar{\mathcal{K}}|dq n \\
& \geq APX - \frac{n}{q} - \iota.
\end{aligned}$$

Since we the marginal signaling schemes do not change, all the persuasiveness and IC constraints are satisfied. Moreover, for every $k \in \bar{\mathcal{K}}$, $\theta \in \Theta$, and $r \in \mathcal{R}$, it holds

$$\sum_{R \subseteq \mathcal{R}: r \in R} \phi_\theta^k(R) = \frac{1}{q} \sum_{j \in [q]} \bar{x}_r^{j, k, \theta} = \frac{1}{q} \sum_{j \in [q]} x_r^{j, k, \theta} = x_\theta^{r, k_r},$$

while it is easy to see that $\sum_{R \subseteq \mathcal{R}} \phi_\theta^k(R) = 1$ for every $k \in \bar{\mathcal{K}}$ and $\theta \in \Theta$. This concludes the proof of the lemma. \square

Now, we can prove Theorem 4.

PROOF OF THEOREM 4. By Lemmas 2 and 3, we only need to provide an algorithm that approximates the optimal solution of LP 8. The objective is a linear combination with non-negative coefficients of the multi-linear extension of monotone submodular functions. Hence, it is smooth, monotone and submodular. Moreover, since we relaxed Constraints (8b), the feasible region is a down-monotone polytope⁵ and it is defined by polynomially-many constraints. For

⁵A polytope $\mathcal{P} \in \mathbb{R}_+^n$ is down-monotone if $0 \leq x \leq y$ coordinate-wise and $y \in \mathcal{P}$ imply $x \in \mathcal{P}$.

Algorithm 1 Algorithm in Lemma 3

Input: N. of samples $\sigma > 0$; Solution x to Program 8; $k \in \bar{\mathcal{K}}$; $\theta \in \Theta$

- 1: **for** $r \in \mathcal{R}$ **do**
- 2: Compute $\hat{e}_r^{j, k, \theta}$ estimating $\frac{\partial F_\theta(x^{j, k, \theta})}{\partial x_r^{j, k, \theta}}$ with σ samples
- 3: Re-label indexes $j \in [q]$ in decreasing order of $\hat{e}_r^{j, k, \theta}$
- 4: $j^* \leftarrow \left\lfloor \sum_{j \in [q]} x_r^{j, k, \theta} \right\rfloor$
- 5: $x_r^{j^*+1, k, \theta} \leftarrow \sum_{j \in [q]} x_r^{j, k, \theta} - j^*$
- 6: $Q^{k, \theta, r} \leftarrow \{1, \dots, j^*\}$
- 7: **for** $j \in Q^{k, \theta, r}$ **do**
- 8: $x_r^{j, k, \theta} \leftarrow 1$
- 9: **for** $j \geq j^* + 2$ **do**
- 10: $x_r^{j, k, \theta} \leftarrow 0$
- 11: Construct ϕ_θ^k as follows:
- 12: $\phi_\theta^k(R) = \frac{1}{q} \left| \left\{ j \in [q] : x_r^{j, k, \theta} = 1 \forall r \in R, x_r^{j, k, \theta} = 0 \forall r \notin R \right\} \right| \forall R \subseteq \mathcal{R}$
- 13: Update ϕ_θ^k to make it consistent with the menus of marginal signaling schemes $\{x_\theta^{r, k_r}\}_{r \in \mathcal{R}}$
- 14: **return** ϕ_θ^k

each $\delta > 0$, this problem admits a $\left(1 - \frac{1}{e}\right) OPT - \delta$ -approximation in time polynomial in the instance size and δ , see the continuous greedy algorithm in [9] and [22] for a formulation in a similar problem.⁶ Finally, we can obtain an arbitrary good approximation choosing an arbitrary large value for q and an arbitrary small value for δ and ι . \square

5 CONCLUSIONS AND FUTURE WORKS

We proposed to extend the Bayesian persuasion framework with a type reporting step. We proved that, with a single receiver, the addition of this type reporting step makes the sender's computational problem tractable. Moreover, we extended the framework to settings with multiple receivers, focusing on the widely-studied case with no inter-agent-externalities and binary actions. We showed that an optimal sender's strategy can be computed in polynomial time when the sender's utility function is supermodular or anonymous. Moreover, when the sender's utility function is submodular, we designed a polynomial-time algorithm that provides a tight $\left(1 - \frac{1}{e}\right)$ -approximation.

In the future, it would be interesting to study the setting in which the sender has access only to samples from the distribution of the receiver's types. Another interesting direction is to explore how the type reporting step can be used to provide polynomial-time no-regret algorithms in an online learning framework.

ACKNOWLEDGMENTS

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Market".

⁶The bound holds only for arbitrary large probability. This reduces the total expected utility by an arbitrary small factor.

REFERENCES

- [1] Ricardo Alonso and Odilon Câmara. 2016. Persuading voters. *American Economic Review* 106, 11 (2016), 3590–3605.
- [2] Itai Arieli and Yakov Babichenko. 2019. Private bayesian persuasion. *Journal of Economic Theory* 182 (2019), 185–217.
- [3] Y. Babichenko and S. Barman. 2017. Algorithmic aspects of private bayesian persuasion. In *Innovations in Theoretical Computer Science Conference*.
- [4] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. 2021. Regret-Minimizing Bayesian Persuasion. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, Péter Biró, Shuchi Chawla, and Federico Echenique (Eds.). ACM, 128. <https://doi.org/10.1145/3465456.3467574>
- [5] Francesco Bacchiocchi, Matteo Castiglioni, Alberto Marchesi, Giulia Romano, and Nicola Gatti. 2022. Public Signaling in Bayesian Ad Auctions. arXiv:2201.09728 [cs.GT]
- [6] Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. 2018. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. 2545–2563.
- [7] Umang Bhaskar, Yu Cheng, Young Kun Ko, and Chaitanya 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*. 479–496.
- [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*. 234–247.
- [9] 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. <https://doi.org/10.1137/080733991> arXiv:<https://doi.org/10.1137/080733991>
- [10] Ozan Candogan. 2019. Persuasion in networks: Public signals and k-cores. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 133–134.
- [11] Matteo Castiglioni, Andrea Celli, and Nicola Gatti. 2020. Persuading Voters: It's Easy to Whisper, It's Hard to Speak Loud. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*. 1870–1877.
- [12] Matteo Castiglioni, Andrea Celli, and Nicola Gatti. 2020. Public Bayesian Persuasion: Being Almost Optimal and Almost Persuasive. arXiv:2002.05156 [cs.GT]
- [13] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. 2020. Online Bayesian Persuasion. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 16188–16198. <https://proceedings.neurips.cc/paper/2020/file/ba5451d3c91a0f982f103cde249bc78-Paper.pdf>
- [14] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. 2021. Signaling in Bayesian Network Congestion Games: the Subtle Power of Symmetry. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- [15] Matteo Castiglioni and Nicola Gatti. 2021. Persuading Voters in District-based Elections. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- [16] Matteo Castiglioni, Alberto Marchesi, Andrea Celli, and Nicola Gatti. 2021. Multi-Receiver Online Bayesian Persuasion. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 1314–1323. <http://proceedings.mlr.press/v139/castiglioni21a.html>
- [17] Matteo Castiglioni, Giulia Romano, Alberto Marchesi, and Nicola Gatti. 2022. Signaling in Posted Price Auctions. arXiv:2201.12183 [cs.GT]
- [18] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. 2015. Mixture Selection, Mechanism Design, and Signaling. In *56th Annual Symposium on Foundations of Computer Science*. 1426–1445.
- [19] Vincent Conitzer and Tuomas Sandholm. 2002. Complexity of mechanism design. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*. 103–110.
- [20] Vincent Conitzer and Tuomas Sandholm. 2003. Automated mechanism design: Complexity results stemming from the single-agent setting. In *Proceedings of the 5th international conference on Electronic commerce*. 17–24.
- [21] S. Dughmi and H. Xu. 2016. Algorithmic bayesian persuasion. In *ACM STOC*. 412–425.
- [22] S. Dughmi and H. Xu. 2017. Algorithmic persuasion with no externalities. In *ACM EC*. 351–368.
- [23] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tennenholtz. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation* 2, 2 (2014), 1–19.
- [24] Mingyu Guo and Vincent Conitzer. 2010. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 24.
- [25] Emir Kamenica. 2019. Bayesian Persuasion and Information Design. *Annual Review of Economics* 11, 1 (2019), 249–272. <https://doi.org/10.1146/annurev-economics-080218-025739> arXiv:<https://doi.org/10.1146/annurev-economics-080218-025739>
- [26] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [27] Yishay Mansour, Aleksandr Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. 2016. Bayesian Exploration: Incentivizing Exploration in Bayesian Games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 661–661.
- [28] Zinovi Rabinovich, Albert Xin Jiang, Manish Jain, and Haifeng Xu. 2015. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. 645–653.
- [29] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.
- [30] Shoshana Vasserman, Michal Feldman, and Avinatan Hassidim. 2015. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 660–666.
- [31] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P Wellman. 2006. Empirical mechanism design: Methods, with application to a supply-chain scenario. In *Proceedings of the 7th ACM conference on Electronic commerce*. 306–315.
- [32] Haifeng Xu. 2020. On the Tractability of Public Persuasion with No Externalities. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2708–2727.
- [33] Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. 2016. Signaling in Bayesian Stackelberg Games. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*. 150–158.
- [34] You Zu, Krishnamurthy Iyer, and Haifeng Xu. 2021. Learning to Persuade on the Fly: Robustness Against Ignorance. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, Péter Biró, Shuchi Chawla, and Federico Echenique (Eds.). ACM, 927–928. <https://doi.org/10.1145/3465456.3467593>