
Multi-Receiver Online Bayesian Persuasion

Matteo Castiglioni¹ Alberto Marchesi¹ Andrea Celli¹ Nicola Gatti¹

Abstract

Bayesian persuasion studies how an informed sender should partially disclose information to influence the behavior of a self-interested receiver. Classical models make the stringent assumption that the sender knows the receiver’s utility. This can be relaxed by considering an *online learning* framework in which the sender repeatedly faces a receiver of an unknown, adversarially selected type. We study, for the first time, an online Bayesian persuasion setting with *multiple receivers*. We focus on the case with *no externalities* and binary actions, as customary in offline models. Our goal is to design *no-regret* algorithms for the sender with polynomial per-iteration running time. First, we prove a negative result: for any $0 < \alpha \leq 1$, there is no polynomial-time no- α -regret algorithm when the sender’s utility function is *supermodular* or *anonymous*. Then, we focus on the case of *submodular* sender’s utility functions and we show that, in this case, it is possible to design a polynomial-time no- $(1 - \frac{1}{e})$ -regret algorithm. To do so, we introduce a general *online gradient descent* scheme to handle online learning problems with a finite number of possible loss functions. This requires the existence of an approximate projection oracle. We show that, in our setting, there exists one such projection oracle which can be implemented in polynomial time.

1. Introduction

Bayesian persuasion was originally introduced by Kamenica & Gentzkow (2011) to model multi-agent settings where an informed *sender* tries to influence the behavior of a self-interested *receiver* through the strategic provision of payoff-relevant information. Agents’ payoffs are determined by the receiver’s action and some exogenous parameters collectively termed the *state of nature*, whose value is drawn from

a common prior distribution and observed by the sender only. Then, the sender decides how much of her/his private information has to be revealed to the receiver, according to a public randomized policy known as *signaling scheme*. From the sender’s perspective, this begets a decision-making problem that is essentially about controlling “who gets to know what”. This kind of problems are ubiquitous in application domains such as auctions and online advertising (Bro Miltersen & Sheffet, 2012; Emek et al., 2014; Badanidiyuru et al., 2018), voting (Alonso & Câmara, 2016; Cheng et al., 2015; Castiglioni et al., 2020a; Castiglioni & Gatti, 2021), traffic routing (Vasserman et al., 2015; Bhaskar et al., 2016; Castiglioni et al., 2021), recommendation systems (Mansour et al., 2016), security (Rabinovich et al., 2015; Xu et al., 2016), and product marketing (Babichenko & Barman, 2017; Candogan, 2019).¹

The classical Bayesian persuasion model by Kamenica & Gentzkow (2011) makes the stringent assumption that the sender knows the receiver’s utility exactly. This is unreasonable in practice. Recently, Castiglioni et al. (2020b) propose to relax the assumption by framing Bayesian persuasion into an online learning framework, focusing on the basic *single-receiver* problem.² In their model, the sender repeatedly faces a receiver whose type during each iteration—determining her/his utility function—is unknown and adversarially selected beforehand. In this work, we extend the model by Castiglioni et al. (2020b) to *multi-receiver* settings, where the (unknown) type of each receiver is adversarially selected before each iteration of the repeated interaction. We consider the case in which the sender has a *private* communication channel towards each receiver, which is commonly studied in multi-receiver models (see, e.g., (Babichenko & Barman, 2016)). Dealing with multiple receivers introduces the additional challenge of correlating information disclosure across them and requires different

¹Politecnico di Milano, Milan, Italy. Correspondence to: Matteo Castiglioni <matteo.castiglioni@polimi.it>.

¹Persuasion was famously attributed to a quarter of the GDP in the United States by McCloskey & Klamer (1995), with a more recent estimate placing this figure at 30% (Antioch et al., 2013).

²A recent work by Babichenko et al. (2021) relaxes the assumption in the offline setting. In that work, the goal is minimizing the sender’s regret over a single iteration, and the authors provide positive results for the case in which the sender knows the ordinal preferences of the receiver over states of nature. The authors study the case of a single receiver with a binary action space, and an arbitrary (unknown) utility function.

techniques from those used in the single-receiver setting.

As customary when studying multi-receiver Bayesian persuasion problems (Dughmi & Xu, 2017; Xu, 2020), we address the case in which there are *no inter-agent externalities*, where each receiver’s utility does *not* depend on the actions of the other receivers, but only on her/his own action and the state of nature. Moreover, we focus on the commonly-studied setting with *binary actions* (Babichenko & Barman, 2016; Arieli & Babichenko, 2019), and we analyze different scenarios depending on whether the sender’s utility function is supermodular, submodular, or anonymous. Despite its simplicity, this basic model encompasses several real-world scenarios. For instance, think of a marketing problem in which a firm (sender) wants to persuade some potential buyers (receivers) to buy one of its products. Each buyer has to take a binary decision as to whether to buy a unit of the product or not, while the firm’s goal is to strategically disclose information about the product to the buyers, so as to maximize the number of units sold. In this example, the sender’s utility is anonymous, since it only depends on the number of buyers who decide to purchase (and not on their identities). Moreover, submodular sender’s utilities represent diminishing returns in the number of items sold, while supermodular ones encode decreasing production costs.

1.1. Original Contributions

Our goal is to design online algorithms for the sender that recommend a signaling scheme at each iteration of the repeated interaction, guaranteeing a sender’s expected utility close to that of the best-in-hindsight signaling scheme. In particular, we look for no- α -regret algorithms, which collect an overall utility that is close to a fraction α of what can be obtained by the best-in-hindsight signaling scheme. In this work, we assume *full-information* feedback, which means that, after each iteration, the sender observes each receiver’s type during that iteration. Moreover, we are interested in no- α -regret algorithms having a per-iteration running time polynomial in the size of the problem instance. To this end, we assume that the number of possible types of each receiver is fixed, otherwise polynomial-time no- α -regret algorithms cannot be obtained even in the degenerate case of only one receiver (Castiglioni et al., 2020b).

In Section 4, we prove a negative result: for any $0 < \alpha \leq 1$, there is no polynomial-time no- α -regret algorithm when the sender’s utility function is *supermodular* or *anonymous*. Thus, in the rest of the work, we focus on the case in which the sender’s utility function is *submodular*, where we provide a polynomial-time no- $(1 - \frac{1}{e})$ -regret algorithm.³

³Our result is tight, as there is no poly-time no- α -regret algorithm with $\alpha > 1 - \frac{1}{e}$. Indeed, it is NP-hard to approximate the sender’s optimal utility within a factor $> 1 - \frac{1}{e}$, even in the basic (offline) multi-receiver model of Babichenko & Barman (2016).

As a first step in building our algorithm, in Section 5 we introduce a general *online gradient descent* (OGD) scheme to handle online learning problems with a finite number of possible loss functions. This can be applied to our setting, as we have a sender’s utility function (or, equivalently, negative loss function) for every combination of receivers’ types obtained as feedback. The OGD scheme works in a modified decision space whose dimensionality is the number of observed loss functions, and it is *not* affected by the dimensionality of the original space. This is crucial in our setting, as it avoids dealing with the set of sender’s signaling schemes, whose dimensionality grows exponentially in the number of receivers. Any OGD algorithm requires a projection oracle. Since in our setting an exact oracle cannot be implemented in polynomial time, we build our OGD scheme so that it works having access to a suitably-defined *approximate projection oracle*, which, as we show later, can be implemented in polynomial time in our model.

In Section 6, we build a polynomial-time approximate projection oracle. First, we formulate the projection problem as a convex linearly-constrained quadratic program, which has exponentially-many variables and polynomially-many constraints. Next, we show how to compute in polynomial time an approximate solution to this program by applying the *ellipsoid algorithm* to its dual. Since the dual has polynomially-many variables and exponentially-many constraints, the algorithm needs access to a particular (problem-dependent) polynomial-time separation oracle. Unfortunately, we do not have this in our setting, and, thus, our algorithm must rely on an *approximate separation oracle*. In general, running the ellipsoid method with an approximate separation oracle does *not* give any guarantee on the approximation quality of the returned solution. In order to make the ellipsoid algorithm return the desired approximate solution by only using an approximate separation oracle, we employ some *ad-hoc* technical tools suggested by a non-trivial primal-dual analysis. As a preparatory step towards our main result, at the beginning of Section 6, we use a derivation similar to that described so far to design a polynomial-time approximation algorithm for the offline version of our multi-receiver Bayesian persuasion problem, which may be of independent interest.

In Section 7, we conclude the construction of the no- $(1 - \frac{1}{e})$ -regret algorithm by showing how to implement in polynomial time an $(1 - \frac{1}{e})$ -approximate separation oracle for settings in which the sender’s utility is submodular.

All the proofs omitted from the paper are in the Appendix.

1.2. Related Works

Most of the computational works on Bayesian persuasion study (offline) models in which the sender knows the receiver’s utility function exactly. Dughmi & Xu (2016) initi-

ate these studies with the single receiver case, while Arieli & Babichenko (2019) 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 the optimal signaling scheme in the case of supermodular, anonymous submodular, and super-majority sender’s utility functions. Babichenko & Barman (2016) 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 & Xu (2017) 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, among others, (Dughmi & Xu, 2017) and (Xu, 2020).

The only computational work on Bayesian persuasion in an online learning framework is that of Castiglioni et al. (2020b), which, however, is restricted to the single-receiver case. The results and techniques in Castiglioni et al. (2020b) are different from those in our paper. In particular, they show that there are no polynomial-time no- α -regret algorithms even in settings with a single receiver, when the number of receiver’s types is arbitrary. In contrast, we focus on settings in which the number of receivers’ types is fixed. Moreover, the main goal of Castiglioni et al. (2020b) is to design a (necessarily exponential-time) no-regret algorithm in the partial-information feedback setting in which the sender only observes the actions played by the receiver (and *not* her/his types). This is accomplished by providing slightly-biased estimators of the sender’s utilities for different signaling schemes. In our work, we assume full-information feedback, and, thus, our main focus is dealing with multiple receivers. This also introduces the additional challenge of correlating information disclosure across the receivers and working with an exponential number of possible feedbacks (tuples specifying a type for each receiver).

Our work is also related to the research line on *online linear optimization with approximation oracles*. In such setting, Kakade et al. (2009) show how to design a no- α -regret algorithm relying on an α -approximate linear optimization oracle, while Garber (2017) and Hazan et al. (2018) obtain analogous results with a better query complexity. The approach of these works to design approximate projection oracles is fundamentally different from ours, since they have access to a linear optimization oracle working in the learner’s decision space. On the other hand, our OGD scheme works on a modified decision space, and the approximate projection oracle can only rely on an approximate separation oracle dealing with the original sender’s decision space.

2. Preliminaries

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 (let $m := \max_{r \in \mathcal{R}} m_r$). 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 $r \in \mathcal{R}$.⁴ 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. Sender and receivers’ payoffs depend on a random state of nature, which is selected from a finite set $\Theta := \{\theta_i\}_{i=1}^d$ of d states. The payoff of a receiver also depends on the action played by her/him, while it does *not* depend on the actions played by the other receivers, since there are *no inter-agent externalities*. Formally, a receiver $r \in \mathcal{R}$ of type $k \in \mathcal{K}_r$ has a utility $u^{r,k} : \mathcal{A}_r \times \Theta \rightarrow [0, 1]$. For the ease of notation, we let $u_\theta^{r,k} := u^{r,k}(a_1, \theta) - u^{r,k}(a_0, \theta)$ be the payoff difference for a receiver r of type k when the state of nature is $\theta \in \Theta$. The sender’s utility depends on the actions played by all the receivers, and it is defined by $u^s : \mathcal{A} \times \Theta \rightarrow [0, 1]$. For the ease of presentation, for every state $\theta \in \Theta$, we introduce the function $f_\theta : 2^{\mathcal{R}} \rightarrow [0, 1]$ such that $f_\theta(R)$ represents the sender’s utility when the state of nature is θ and all the receivers in $R \subseteq \mathcal{R}$ play action a_1 , while the others play a_0 .

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 receivers.⁵ The sender can commit to a *signaling scheme* ϕ , which is a randomized mapping from states of nature to signals for the receivers. In this work, we focus on *private* signaling, where each receiver has her/his own signal that is privately communicated to her/him. Formally, there is a finite set \mathcal{S}_r of possible signals for each receiver $r \in \mathcal{R}$. Then, $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$, 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 with ϕ_θ the probability distribution employed by ϕ when the state of nature is $\theta \in \Theta$, with $\phi_\theta(\mathbf{s})$ being the probability of sending a signal profile $\mathbf{s} \in \mathcal{S}$. The *one-shot* interaction between the sender and the receivers goes on as follows: (i) the sender commits to a publicly known signaling scheme ϕ ; (ii) she/he observes the realized state of nature $\theta \sim \mu$; (iii) she/he draws a signal profile $\mathbf{s} \sim \phi_\theta$ and communicates to each receiver $r \in \mathcal{R}$ signal s_r ; and (iv) each receiver $r \in \mathcal{R}$ rationally updates her/his prior belief over Θ according to the *Bayes rule* and selects an action maximizing her/his expected utility. We remark that, given a signaling scheme ϕ , a receiver $r \in \mathcal{R}$ of type $k \in \mathcal{K}_r$

⁴All vectors and tuples are denoted by bold symbols. For any vector (tuple) \mathbf{x} , the value of its i -th component is denoted by x_i .

⁵ $\text{int}(X)$ is the *interior* of a set X , while Δ_X is the set of all the probability distributions over a set X .

observing a private signal $s \in \mathcal{S}_r$ experiences an expected utility $\sum_{\theta \in \Theta} \mu_\theta \sum_{s \in \mathcal{S}: s_r = s} \phi_\theta(s) u^{r,k}(a, \theta)$ (up to a normalization constant) when playing action $a \in \mathcal{A}_r$. Assuming the receivers' type profile is $\mathbf{k} \in \mathcal{K}$, the goal of the sender is to commit to an *optimal* signaling scheme ϕ , which is one maximizing her/his expected utility $f(\phi, \mathbf{k}) := \sum_{\theta \in \Theta} \mu_\theta \sum_{s \in \mathcal{S}} \phi_\theta(s) f_\theta(R_s^{\mathbf{k}})$, where we let $R_s^{\mathbf{k}} \subseteq \mathcal{R}$ be the set of receivers who play a_1 after observing their private signal s_r in \mathbf{s} , under signaling scheme ϕ .

Assumptions In the rest of this work, we assume that the sender's utility is *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. Moreover, we assume that the number of types m_r of each receiver $r \in \mathcal{R}$ is fixed; in other words, the value of m cannot grow arbitrarily large.⁶

Direct Signaling Schemes By well-known revelation-principle-style arguments (Kamenica & Gentzkow, 2011; Arieli & Babichenko, 2019), we can restrict our attention to signaling schemes that are direct and persuasive. In words, a signaling scheme is *direct* if signals correspond to recommendations of playing actions, while it is *persuasive* if the receivers do not have any incentive to deviate from the recommendations prescribed by the signals they receive. In our setting, a direct signal sent to a receiver specifies an action recommendation for each receiver's type; thus, we let $\mathcal{S}_r := 2^{\mathcal{K}_r}$ for every $r \in \mathcal{R}$. A signal $s \in \mathcal{S}_r$ for a receiver $r \in \mathcal{R}$ is encoded by a subset of her/his types, namely $s \subseteq \mathcal{K}_r$. Intuitively, s can be interpreted as the recommendation to play action a_1 when the receiver has type $k \in \mathcal{K}_r$ such that $k \in s$, while a_0 otherwise. Given a direct and persuasive signaling scheme ϕ , for a signal profile $\mathbf{s} \in \mathcal{S}$ and a type profile $\mathbf{k} \in \mathcal{K}$, the set $R_s^{\mathbf{k}}$ appearing in the definition of the sender's expected utility $f(\phi, \mathbf{k})$ can be formally expressed as $R_s^{\mathbf{k}} := \{r \in \mathcal{R} \mid k_r \in s_r\}$.

Set Functions and Matroids In Section 7, we show how to implement our approximate separation oracle by optimizing functions f_θ over suitably defined matroids (representing signals). Next, we introduce the necessary definitions on set functions and matroids. For the ease of presentation, we consider a generic function $f : 2^{\mathcal{G}} \rightarrow [0, 1]$ for a finite set \mathcal{G} . We say that f is *submodular*, respectively *supermodular*, if for $I, I' \subseteq \mathcal{G}$: $f(I \cap I') + f(I \cup I') \leq f(I) + f(I')$, respectively $f(I \cap I') + f(I \cup I') \geq f(I) + f(I')$. The function f is *anonymous* if $f(I) = f(I')$ for all $I, I' \subseteq \mathcal{G} : |I| = |I'|$. A matroid $\mathcal{M} := (\mathcal{G}, \mathcal{I})$ is defined by a finite ground set

⁶The monotonicity assumption is w.l.o.g. for this work, since our main positive result (Theorem 7) relies on it. Instead, assuming a fixed number of types is necessary, since, even in single-receiver settings, designing no-regret algorithms with running time polynomial in m is intractable (Castiglioni et al., 2020b).

\mathcal{G} and a collection \mathcal{I} of independent sets, i.e., subsets of \mathcal{G} satisfying some characterizing properties (see (Schrijver, 2003) for a detailed formal definition). We denote by $\mathcal{B}(\mathcal{M})$ the set of the *bases* of \mathcal{M} , which are the maximal sets in \mathcal{I} .

3. Multi-Receiver Online Bayesian Persuasion

We consider a multi-receiver generalization of the online setting introduced by Castiglioni et al. (2020b). The sender plays a repeated game in which, at each iteration $t \in [T]$, she/he commits to a signaling scheme ϕ^t , observes the realized state of nature $\theta^t \sim \mu$, and privately sends signals determined by $\mathbf{s}^t \sim \phi_{\theta^t}^t$ to the receivers.⁷ Then, each receiver (whose type is unknown to the sender) selects an action maximizing her/his expected utility given the observed signal (in the *one-shot* interaction at iteration t).

We focus on the problem of computing a sequence $\{\phi^t\}_{t \in [T]}$ of signaling schemes maximizing the sender's expected utility when the sequence of receivers' types $\{\mathbf{k}^t\}_{t \in [T]}$, with $\mathbf{k}^t \in \mathcal{K}$, is adversarially selected beforehand. After each iteration $t \in [T]$, the sender gets payoff $f(\phi^t, \mathbf{k}^t)$ and receives a *full-information feedback* on her/his choice at t , which is represented by the type profile \mathbf{k}^t . Therefore, after each iteration, the sender can compute the expected utility $f(\phi, \mathbf{k}^t)$ guaranteed by any signaling scheme ϕ she/he could have chosen during that iteration.

We are interested in an algorithm computing ϕ^t at each iteration $t \in [T]$. We measure the performance of one such algorithm using the α -regret R_α^T . Formally, for $0 < \alpha \leq 1$,

$$R_\alpha^T := \alpha \max_{\phi} \sum_{t \in [T]} f(\phi, \mathbf{k}^t) - \mathbb{E} \left[\sum_{t \in [T]} f(\phi^t, \mathbf{k}^t) \right],$$

where the expectation is on the randomness of the algorithm. The classical notion of regret is obtained for $\alpha = 1$.

Ideally, we would like an algorithm that returns a sequence $\{\phi^t\}_{t \in [T]}$ with the following properties:

- the α -regret is sublinear in T for some $0 < \alpha \leq 1$;
- the number of computational steps it takes to compute ϕ^t at each iteration $t \in [T]$ is $\text{poly}(T, n, d)$, that is, it is a polynomial function of the parameters T , n , and d

An algorithm satisfying the first property is called a *no- α -regret algorithm* (it is *no-regret* if it does so for $\alpha = 1$). In this work, we focus on the weaker notion of α -regret since, as we discuss next, requiring no-regret is oftentimes too limiting in our setting (from a computational perspective).

⁷Throughout the paper, the set $\{1, \dots, x\}$ is denoted by $[x]$.

4. Hardness of Being No- α -Regret

We start with a negative result. We show that designing no- α -regret algorithms with polynomial per-iteration running time is an intractable problem (formally, it is impossible unless $\text{NP} \subseteq \text{RP}$) when the sender's utility is such that functions f_θ are *supermodular* or *anonymous*. This hardness result is deeply connected with the intractability of the offline version of our multi-receiver Bayesian persuasion problem that we formally define in the following Section 4.1. Then, Section 4.2 collects all the hardness results.

4.1. Offline Multi-Receiver Bayesian Persuasion

We consider an offline setting where the receivers' type profile $\mathbf{k} \in \mathcal{K}$ is drawn from a known probability distribution (rather than being selected adversarially at each iteration). Given a subset of possible type profiles $K \subseteq \mathcal{K}$ and a distribution $\lambda \in \text{int}(\Delta_K)$, we call BAYESIAN-OPT-SIGNAL the problem of computing a signaling scheme that maximizes the sender's expected utility. This can be achieved by solving the following LP of exponential size.⁸

$$\max_{\phi} \sum_{\mathbf{k} \in \mathcal{K}} \lambda_{\mathbf{k}} \sum_{\theta \in \Theta} \mu_{\theta} \sum_{\mathbf{s} \in \mathcal{S}} \phi_{\theta}(\mathbf{s}) f_{\theta}(R_{\mathbf{s}}^{\mathbf{k}}) \quad (1a)$$

$$\text{s.t.} \sum_{\theta \in \Theta} \mu_{\theta} \sum_{\mathbf{s} \in \mathcal{S}: s_r = s} \phi_{\theta}(\mathbf{s}) u_{\theta}^{r, \mathbf{k}} \geq 0 \quad \forall r \in \mathcal{R}, \forall s \in \mathcal{S}_r, \forall \mathbf{k} \in \mathcal{K}_r : \mathbf{k} \in \mathcal{S} \quad (1b)$$

$$\sum_{\mathbf{s} \in \mathcal{S}} \phi_{\theta}(\mathbf{s}) = 1 \quad \forall \theta \in \Theta \quad (1c)$$

$$\phi_{\theta}(\mathbf{s}) \geq 0 \quad \forall \theta \in \Theta, \forall \mathbf{s} \in \mathcal{S}. \quad (1d)$$

4.2. Hardness Results

First, we study the computational complexity of finding an approximate solution to BAYESIAN-OPT-SIGNAL. In particular, given $0 < \alpha \leq 1$, we look for an α -approximate solution in the multiplicative sense, *i.e.*, a signaling scheme providing at least a fraction α of the sender's optimal expected utility (the optimal value of LP (1)). Theorem 1 provides our main hardness result, which is based on a reduction from the *promise-version* of LABEL-COVER (see Appendix A for its definition and the proof of the theorem).

Theorem 1. *For every $0 < \alpha \leq 1$, it is NP-hard to compute an α -approximate solution to BAYESIAN-OPT-SIGNAL, even when the sender's utility is such that, for every $\theta \in \Theta$, $f_{\theta}(R) = 1$ iff $|R| \geq 2$, while $f_{\theta}(R) = 0$ otherwise.*

Notice that Theorem 1 holds for problem instances in which

⁸Constraints (1b) encode persuasiveness for the signals recommending to play a_1 . The analogous constraints for a_0 can be omitted. Indeed, by assuming that each f_{θ} is non-decreasing in the set of receivers who play a_1 , any signaling scheme in which the sender recommends a_0 when the state is θ and the receiver prefers a_1 over a_0 can be improved by recommending a_1 instead.

functions f_{θ} are anonymous. Moreover, the reduction can be easily modified so that functions f_{θ} are supermodular and satisfy $f_{\theta}(R) = \max\{0, |R| - 1\}$ for $R \subseteq \mathcal{R}$. Thus:

Corollary 1.1. *For $0 < \alpha \leq 1$, it is NP-hard to compute an α -approximate solution to BAYESIAN-OPT-SIGNAL, even when the sender's utility is such that functions f_{θ} are supermodular or anonymous for every $\theta \in \Theta$.*

By using arguments similar to those employed in the proof of Theorem 6.2 by Roughgarden & Wang (2019), the hardness of computing an α -approximate solution to the offline problem can be extended to designing no- α -regret algorithms in the online setting. Then:

Theorem 2. *For every $0 < \alpha \leq 1$, there is no polynomial-time no- α -regret algorithm for the multi-receiver online Bayesian persuasion problem, unless $\text{NP} \subseteq \text{RP}$, even when functions f_{θ} are supermodular or anonymous for all $\theta \in \Theta$.*

In the rest of the work, we show how to design a polynomial-time no- $(1 - \frac{1}{e})$ -regret algorithm for the case in which the sender's utility is such that functions f_{θ} are submodular.

5. An Online Gradient Descent Scheme with Approximate Projection Oracles

As a first step in building our polynomial-time algorithm, we introduce our OGD scheme with an *approximate projection oracle*. Intuitively, it works by transforming the multi-receiver online Bayesian persuasion setting into an equivalent online learning problem whose decision space does not need to explicitly deal with signaling schemes (thus avoiding the burden of having an exponential number of possible signal profiles). The OGD algorithm is then applied on this new domain. In our setting, we do *not* have access to a polynomial-time (exact) projection oracle, and, thus, we design and analyze the algorithm assuming access to an approximate one only. As we show later in Sections 6 and 7, such approximate projection oracle can be implemented in polynomial time when the functions f_{θ} are submodular.

Let us recall that the OGD scheme that we describe in this section is general and applies to any online learning problem with a finite number of possible loss functions.

5.1. A General Approach

Consider an online learning problem in which the learner takes a decision $y^t \in \mathcal{Y}$ at each iteration $t \in [T]$. Then, the learner observes a feedback $e^t \in \mathcal{E}$, where \mathcal{E} is a finite set of p possible feedbacks. The reward (or negative loss) of a decision $y \in \mathcal{Y}$ given feedback $e \in \mathcal{E}$ is defined by $u(y, e)$ for a given function $u : \mathcal{Y} \times \mathcal{E} \rightarrow [0, 1]$. Thus, the learner is awarded $u(y^t, e^t)$ for decision y^t at iteration t , while she/he would have achieved $u(y, e^t)$ for any other choice $y \in \mathcal{Y}$.

We transform this general online learning problem to a new one in which the learner's decision set is $\mathcal{X} \subseteq [0, 1]^p$ with:

$$\mathcal{X} := \bigcup_{y \in \mathcal{Y}} \left\{ \mathbf{x} \in [0, 1]^p \mid x_e \leq u(y, e) \quad \forall e \in \mathcal{E} \right\}. \quad (2)$$

Intuitively, the set \mathcal{X} contains all the vectors whose components x_e (one for each feedback $e \in \mathcal{E}$) are the learner's rewards $u(y, e)$ for some decision $y \in \mathcal{Y}$ in the original problem. Moreover, the inequality " \leq " in the definition of \mathcal{X} also includes all the reward vectors that are dominated by those corresponding to some decision in \mathcal{Y} . At each iteration $t \in [T]$, the learner takes a decision $\mathbf{x}^t \in \mathcal{X}$ and observes a feedback $e^t \in \mathcal{E}$. The reward of decision $\mathbf{x} \in \mathcal{X}$ at iteration t is the e^t -th component of \mathbf{x} , namely x_{e^t} . It is sometimes useful to write it as $\mathbf{1}_{e^t}^\top \mathbf{x}$, where $\mathbf{1}_{e^t} \in \{0, 1\}^p$ is a vector whose e^t -th component is 1, while all the others are 0. Thus, the learner's reward at iteration t is x_{e^t} . Notice that the size of the decision set \mathcal{X} of the new online learning setting does *not* depend on the dimensionality of the original decision set \mathcal{Y} (which, in our setting, would be exponential), but only on the number of feedbacks p .

If \mathcal{Y} and u are such that \mathcal{X} is compact and convex, then we can minimize the α -regret R_α^T in the original problem by doing that in the new setting. Let us introduce the set $\alpha\mathcal{X} := \{\alpha\mathbf{x} \mid \mathbf{x} \in \mathcal{X}\}$ for any $0 < \alpha \leq 1$. Given a sequence of feedbacks $\{e^t\}_{t \in [T]}$ and a sequence of decisions $\{\mathbf{x}^t\}_{t \in [T]}$, with $e^t \in \mathcal{E}$ and $\mathbf{x}^t \in \mathcal{X}$, we have that:

$$\begin{aligned} R_\alpha^T &:= \max_{\mathbf{x} \in \alpha\mathcal{X}} \sum_{t \in [T]} \mathbf{1}_{e^t}^\top (\mathbf{x} - \mathbf{x}^t) \\ &\geq \alpha \max_{y \in \mathcal{Y}} \sum_{t \in [T]} u(y, e^t) - \sum_{t \in [T]} u(y^t, e^t), \end{aligned}$$

where $\{y^t\}_{t \in [T]}$ is a sequence of decisions $y^t \in \mathcal{Y}$ for the original problem such that $x_{e^t}^t \leq u(y^t, e)$ for $e \in \mathcal{E}$.

We assume to have access to an approximate projection oracle for $\alpha\mathcal{X}$, which we define in the following. By letting $E \subseteq \mathcal{E}$ be a subset of feedbacks, we define $\tau_E : \mathcal{X} \rightarrow [0, 1]^p$ as the function mapping any vector $\mathbf{x} \in \mathcal{X}$ to another one that is equal to \mathbf{x} in all the components corresponding to feedbacks $e \in E$, while it is 0 everywhere else. Moreover, we let $\mathcal{X}_E := \{\tau_E(\mathbf{x}) \mid \mathbf{x} \in \mathcal{X}\}$ be the image of \mathcal{X} through τ_E , while $\alpha\mathcal{X}_E := \{\alpha\mathbf{x} \mid \mathbf{x} \in \mathcal{X}_E\}$ for $0 < \alpha \leq 1$.

Definition 1 (Approximate projection oracle). *Consider a subset of feedbacks $E \subseteq \mathcal{E}$, a vector $\mathbf{y} \in [0, 2]^p$ such that $y_e = 0$ for all $e \notin E$, and an approximation error $\epsilon \in \mathbb{R}_+$. Then, for any $0 < \alpha \leq 1$, an approximate projection oracle $\varphi_\alpha(E, \mathbf{y}, \epsilon)$ is an algorithm returning a vector $\mathbf{x} \in \mathcal{X}_E$ and a decision $y \in \mathcal{Y}$ with $x_e \leq u(y, e)$ for all $e \in \mathcal{E}$, such that:*

$$\|\mathbf{x}' - \mathbf{x}\|^2 \leq \|\mathbf{x}' - \mathbf{y}\|^2 + \epsilon \quad \forall \mathbf{x}' \in \alpha\mathcal{X}_E.$$

Intuitively, φ_α returns a vector $\mathbf{x} \in \mathcal{X}_E$ that is an approximate projection of \mathbf{y} onto the subspace $\alpha\mathcal{X}_E$. The vector

\mathbf{x} can be outside of $\alpha\mathcal{X}_E$. However, it is "better" than a projection onto $\alpha\mathcal{X}_E$, since, ignoring the ϵ error, \mathbf{x} is closer than \mathbf{y} to any vector in $\alpha\mathcal{X}_E$. Moreover, φ_α also gives a decision $y \in \mathcal{Y}$ that corresponds to the returned vector \mathbf{x} . Notice that, if $\alpha = 1$ and $\epsilon = 0$, this is equivalent to find an exact projection onto the subspace \mathcal{X}_E .

5.2. A Particular Setting: Multi-Receiver Online Bayesian Persuasion

Our setting can be easily cast into the general learning framework described so far. The possible feedbacks are type profiles, namely $\mathcal{E} := \mathcal{K}$, while the receivers' type profile $\mathbf{k}^t \in \mathcal{K}$ is the feedback observed at iteration $t \in [T]$, namely $e^t := \mathbf{k}^t$. Notice that the number of possible feedbacks is p is m^n , which is exponential in the number of receivers. The decision set of the learner (sender) \mathcal{Y} is the set of all the possible signaling schemes ϕ , with $y^t := \phi^t$ being the one chosen at iteration t . The rewards observed by the sender are the utilities $f(\phi, \mathbf{k})$; formally, for every signaling scheme ϕ and type profile $\mathbf{k} \in \mathcal{K}$, which define a pair $y \in \mathcal{Y}$ and $e \in \mathcal{E}$ using the generic notation, we let $u(y, e) := f(\phi, \mathbf{k})$. Then, the new decision set $\mathcal{X} \subseteq [0, 1]^{|\mathcal{K}|}$ is defined as in Equation (2). Notice that \mathcal{X} is a compact and convex set, since it can be defined by a set of linear inequalities. In the following, we overload the notation and, for any subset $K \subseteq \mathcal{K}$ of types profiles, we let $\mathcal{X}_K := \mathcal{X}_E$ for $E \subseteq \mathcal{E} : E = K$.

5.3. OGD with Approximate Projection Oracle

Algorithm 1 is an OGD scheme that operates in the \mathcal{X} domain by having access to an approximate projection oracle φ_α (we call the algorithm OGD-APO).

The procedure in Algorithm 1 keeps track of the set $E^t \subseteq \mathcal{E}$ of different feedbacks observed up to each iteration $t \in [T]$. Moreover, it works on the subspace \mathcal{X}_{E^t} , whose vectors are zero in all the components corresponding to feedbacks $e \notin E^t$. Since it is the case that $|E^t| \leq t$, the procedure in Algorithm 1 attains a per-iteration running time that is independent of the number of possible feedbacks p .

Algorithm 1 OGD-APO

- Input:**
- approximate projection oracle φ_α
 - learning rate $\eta \in (0, 1]$
 - approximation error $\epsilon \in [0, 1]$
-

Initialize $y^1 \in \mathcal{Y}$, $E^0 \leftarrow \emptyset$, and $\mathbf{x}^1 \leftarrow \mathbf{0} \in \mathcal{X}_{E^1}$

for $t = 1, \dots, T$ **do**

Take decision y^t

Observe feedback $e^t \in \mathcal{E}$ and reward $u(y^t, e^t) = x_{e^t}^t$

$E^t \leftarrow E^{t-1} \cup \{e^t\}$

$\mathbf{y}^{t+1} \leftarrow \mathbf{x}^t + \eta \mathbf{1}_{e^t}$

$(\mathbf{x}^{t+1}, y^{t+1}) \leftarrow \varphi_\alpha(E^t, \mathbf{y}^{t+1}, \epsilon)$

end for

Next, we bound the α -regret incurred by Algorithm 1.

Theorem 3. *Given an oracle φ_α (as in Definition 1) for some $0 < \alpha \leq 1$, a learning rate $\eta \in (0, 1]$, and an approximation error $\epsilon \in [0, 1]$, Algorithm 1 has α -regret*

$$R_\alpha^T \leq \frac{|E^T|}{2\eta} + \frac{\eta T}{2} + \frac{\epsilon T}{2\eta},$$

with a per-iteration running time $\text{poly}(t)$.

By setting $\eta = \frac{1}{\sqrt{T}}$, $\epsilon = \frac{1}{T}$, we get $R_\alpha^T \leq \sqrt{T} \left(1 + \frac{|E^T|}{2}\right)$.

Notice that the bound only depends on the number of observed feedbacks $|E^T|$, while it is independent of the overall number of possible feedbacks p . This is crucial for the multi-receiver online Bayesian persuasion case, where p is exponential in the the number of receivers n . On the other hand, as T goes to infinity, we have $|E^T| \leq p$, so that the regret bound is sublinear in T .

6. Constructing a Poly-Time Approximate Projection Oracle

The crux of the OGD-APO algorithm (Algorithm 1) is being able to perform the approximate projection step. In this section, we show that, in the multi-receiver Bayesian persuasion setting, the approximate projection oracle φ_α required by OGD-APO can be implemented in polynomial time by an appropriately-engineered ellipsoid algorithm. This calls for an *approximate separation oracle* \mathcal{O}_α (see Definition 2).

We proceed as follows. In Section 6.1, we define an appropriate notion of approximate separation oracle, and show how to find, in polynomial time, an α -approximate solution to the offline problem BAYESIAN-OPT-SIGNAL. This is a preparatory step towards the understanding of our main result in this section, and it may be of independent interest. Then, in Section 6.2, we exploit some of the techniques introduced for the offline setting in order to build φ_α starting from an approximate separation oracle \mathcal{O}_α .

6.1. Warming Up: The Offline Setting

An approximate separation oracle \mathcal{O}_α finds a signal profile $s \in \mathcal{S}$ that approximately maximizes a weighted sum of the f_θ functions, plus a weight for each receiver which depends on the signal s_r sent to that receiver. Formally:

Definition 2 (Approximate separation oracle). *Consider a state $\theta \in \Theta$, a subset $K \subseteq \mathcal{K}$, a vector $\lambda \in \mathbb{R}_+^{|K|}$, weights $\mathbf{w} = (w_{r,s})_{r \in \mathcal{R}, s \in \mathcal{S}_r}$ with $w_{r,s} \in \mathbb{R}$ and $w_{r,\emptyset} = 0$ for all $r \in \mathcal{R}$, and an approximation error $\epsilon \in \mathbb{R}_+$. Then, for any $0 < \alpha \leq 1$, an approximation oracle $\mathcal{O}_\alpha(\theta, K, \lambda, \mathbf{w}, \epsilon)$ is*

an algorithm returning an $s \in \mathcal{S}$ such that:

$$\begin{aligned} & \sum_{\mathbf{k} \in K} \lambda_{\mathbf{k}} f_\theta(R_{\mathbf{s}}^{\mathbf{k}}) + \sum_{r \in \mathcal{R}} w_{r,s_r} \\ & \geq \max_{\mathbf{s}^* \in \mathcal{S}} \left\{ \alpha \sum_{\mathbf{k} \in K} \lambda_{\mathbf{k}} f_\theta(R_{\mathbf{s}^*}^{\mathbf{k}}) + \sum_{r \in \mathcal{R}} w_{r,s_r^*} \right\} - \epsilon, \end{aligned} \quad (3)$$

in time $\text{poly}\left(n, |K|, \max_{r,s} |w_{r,s}|, \max_{\mathbf{k}} \lambda_{\mathbf{k}}, \frac{1}{\epsilon}\right)$.

As a preliminary result, we show how to use an oracle \mathcal{O}_α to find in polynomial time an α -approximate solution to BAYESIAN-OPT-SIGNAL (see Section 4). This problem is interesting in its own right, and allows us to develop a line of reasoning that will be essential to prove Theorem 5.

Theorem 4. *Given $\epsilon \in \mathbb{R}_+$ and an approximate separation oracle \mathcal{O}_α , with $0 < \alpha \leq 1$, there exists a polynomial-time approximation algorithm for BAYESIAN-OPT-SIGNAL returning a signaling scheme with sender's utility at least $\alpha \text{OPT} - \epsilon$, where OPT is the value of an optimal signaling scheme. Moreover, the algorithm works in time $\text{poly}\left(\frac{1}{\epsilon}\right)$.*

Proof Overview. The dual of LP (1) has a polynomial number of variables and an exponential number of constraints, and a natural way to prove polynomial-time solvability would be via the ellipsoid method (see, e.g., (Khachiyan, 1980; Grötschel et al., 1981)). However, in our setting, we can only rely on an approximate separation oracle, which renders the traditional ellipsoid method unsuitable for our problem. We show that it is possible to exploit a binary search scheme on the dual problem to find a value $\gamma^* \in [0, 1]$ such that the dual problem with objective γ^* is feasible, while the dual with objective $\gamma^* - \beta$, $\beta \geq 0$, is infeasible. That algorithm runs in $\log(\beta)$ steps. At each iteration of the algorithm, we solve a feasibility problem through the ellipsoid method equipped with an appropriate approximate separation oracle which we design. In order to build a poly-time separation oracle we have to carefully manage all the settings in which \mathcal{O}_α would not run in polynomial time, according to Definition 2. Specifically, we need to properly manage large values of the weights \mathbf{w} , since \mathcal{O}_α is polynomial in $\max_{r,s} |w_{r,s}|$. Once we do that, the approximate separation oracle is guaranteed to find a violated constraint, or to certify that all constraints are approximately satisfied. Finally, we show that the approximately feasible solution computed via bisection allows one to recover an approximate solution to the original problem \square

6.2. From an Approximate Separation Oracle to an Approximate Projection Oracle

Now, we show how to design a polynomial-time approximate projection oracle φ_α using an approximate separation oracle \mathcal{O}_α . The proof employs a convex linearly-constrained

quadratic program that computes the optimal projection on \mathcal{X} , the ellipsoid method, and a careful primal-dual analysis.

Theorem 5. *Given a subset $K \subseteq \mathcal{K}$, a vector $\mathbf{y} \in [0, 2]^{|\mathcal{K}|}$ such that $y_{\mathbf{k}} = 0$ for all $\mathbf{k} \notin K$, and an approximation error $\epsilon \in \mathbb{R}_+$, for any $0 < \alpha \leq 1$, the approximate projection oracle $\varphi_\alpha(K, \mathbf{y}, \epsilon)$ can be computed in polynomial time by querying the approximate separation oracle \mathcal{O}_α .*

Proof Overview. We start by defining a convex minimization problem, which we denote by \textcircled{P} , for computing the projection of \mathbf{y} on \mathcal{X}_K . Then, we work on the dual of \textcircled{P} , which we suitably simplify by reasoning over the KKT conditions of the problem. As in the proof of Theorem 4, we proceed by repeatedly applying the ellipsoid method on a feasibility problem obtained from the dual, decreasing the required objective γ^* by a small additive factor β . The ellipsoid method is equipped with the approximate separation oracle that employs the oracle in Definition 2 and carefully manages the cases in which \mathcal{O}_α would not run in polynomial time. In this case, the problem is complicated by the fact that we have to determine an approximate projection over $\alpha\mathcal{X}_K$, rather than an approximate solution to \textcircled{P} . We found two dual problems such that one dual problem with objective γ^* is feasible, while the second one with objective $\gamma^* + \beta$ is infeasible. From these problems, we define a new convex optimization problem that is a modified version of \textcircled{P} and has value at least γ^* . Then, we show that a solution to this problem is close to a projection on a set which includes $\alpha\mathcal{X}_K$. Finally, we restrict \textcircled{P} to the primal variables corresponding to the set of (polynomially-many) violated dual constraints determined during the last application of the ellipsoid method that returns unfeasible, *i.e.*, where the ellipsoid method for feasibility problem is run with objective $\gamma^* + \beta$. We conclude the proof by showing that a solution to this restricted problem is precisely an approximate projection on a superset of $\alpha\mathcal{X}_K$. \square

7. A Poly-Time No- α -Regret Algorithm for Submodular Sender's Utilities

In this section, we conclude the construction of our polynomial-time no- $(1 - \frac{1}{e})$ -regret algorithm for settings in which sender's utilities are submodular. The last component that we need to design is an approximate separation oracle \mathcal{O}_α (see Definition 2) running in polynomial time. Next, we show how to obtain this by exploiting the fact that functions f_θ are submodular in the set of receivers playing action a_1 .

First, we establish a relation between direct signals \mathcal{S} and matroids. We define a matroid $\mathcal{M}_\mathcal{S} := (\mathcal{G}_\mathcal{S}, \mathcal{I}_\mathcal{S})$ such that:

- the ground set is $\mathcal{G}_\mathcal{S} := \{(r, s) \mid r \in \mathcal{R}, s \in \mathcal{S}_r\}$;
- a subset $I \subseteq \mathcal{G}_\mathcal{S}$ belongs to $\mathcal{I}_\mathcal{S}$ if and only if I contains at most one pair for each receiver $r \in \mathcal{R}$.

The elements of the ground set $\mathcal{G}_\mathcal{S}$ represent receiver, signal pairs. However, sets $I \in \mathcal{I}_\mathcal{S}$ do *not* characterize signal profiles, as they may not define a signal for each receiver. Indeed, direct signal profiles are captured by the basis set $\mathcal{B}(\mathcal{M}_\mathcal{S})$ of the matroid $\mathcal{M}_\mathcal{S}$. Let us recall that $\mathcal{B}(\mathcal{M}_\mathcal{S})$ contains all the maximal sets in $\mathcal{I}_\mathcal{S}$, and, thus, a subset $I \subseteq \mathcal{I}_\mathcal{S}$ belongs to $\mathcal{B}(\mathcal{M}_\mathcal{S})$ if and only if I contains *exactly one* pair for each receiver $r \in \mathcal{R}$. Intuitively, a basis $I \in \mathcal{B}(\mathcal{M}_\mathcal{S})$ defines a direct signal profile $s \in \mathcal{S}$ in which, for each receiver $r \in \mathcal{R}$, all the receiver's types in $s \in \mathcal{S}_r$ such that $(r, s) \in I$ are recommended to play action a_1 , while the others are told to play a_0 .

The following Theorem 6 provides a polynomial-time approximation oracle $\mathcal{O}_{1-\frac{1}{e}}$ for instances in which f_θ is submodular for each state of nature $\theta \in \Theta$. The core idea of its proof is that $\sum_{\mathbf{k} \in K} \lambda_{\mathbf{k}} f_\theta(R_{\mathbf{s}}^{\mathbf{k}})$ (see Equation (3)) can be seen as a submodular function defined for the ground set $\mathcal{G}_\mathcal{S}$ and optimizing over direct signal profiles $s \in \mathcal{S}$ is equivalent to doing that over the bases $\mathcal{B}(\mathcal{M}_\mathcal{S})$ of the matroid $\mathcal{M}_\mathcal{S}$. Then, the result is readily proved by exploiting some results concerning the optimization over matroids.⁹

Theorem 6. *If the sender's utility is such that function f_θ is submodular for each $\theta \in \Theta$, then there exists a polynomial-time separation oracle $\mathcal{O}_{1-\frac{1}{e}}$.*

In conclusion, by letting $\mathcal{K}^T \subseteq \mathcal{K}$ be the set of receivers' type profiles observed by the sender up to iteration T , the following Theorem 7 provides our polynomial-time no- $(1 - \frac{1}{e})$ -regret algorithm working with submodular sender's utilities.

Theorem 7. *If the sender's utility is such that function f_θ is submodular for each $\theta \in \Theta$, then there exists a no- $(1 - \frac{1}{e})$ -regret algorithm having $(1 - \frac{1}{e})$ -regret*

$$R_{1-\frac{1}{e}}^T \leq O\left(\sqrt{T} |\mathcal{K}^T|\right),$$

with a per-iteration running time $\text{poly}(T, n, d)$.

Proof. We can run Algorithm 1 on an instance of our multi-receiver online Bayesian persuasion problem. By Theorem 3, if we set $\eta = \frac{1}{\sqrt{T}}$, $\epsilon = \frac{1}{T}$, and $\alpha = 1 - \frac{1}{e}$, we get the desired regret bound (notice that the set of observed feedbacks is $E^t = \mathcal{K}^t$ in our setting). Algorithm 1 employs an approximate projection oracle $\varphi_{1-\frac{1}{e}}$ that we can implement in polynomial time by using the algorithm provided in Theorem 5. This requires access to a polynomial-time approximate separation oracle $\mathcal{O}_{1-\frac{1}{e}}$, which can be implemented by using Theorem 6, under the assumption that the sender's utility is such that functions f_θ are submodular. \square

⁹The separation oracle provided in Theorem 6 guarantees the desired approximation factor with arbitrary high probability. It is easy to see that, since the algorithm fails with arbitrary small probability, this does not modify our regret bound except for an (arbitrary small) negligible term.

Notice that the regret bound only depends on the number $|\mathcal{K}^T|$ of receivers' type profiles observed up to iteration T , while it is independent of the overall number of possible type profiles $|\mathcal{K}| = m^n$, which is exponential in the number of receivers. Thus, the $(1 - \frac{1}{e})$ -regret is polynomial in the size of the problem instance provided that the type profiles received as feedbacks by the sender are polynomially many (though the sender does not have to know which are these type profiles in advance). This is reasonable in many practical applications, where not all the type profiles can occur, since, e.g., receivers' types are highly correlated. On the other hand, let us remark that, as T goes to infinity, we have $|\mathcal{K}^T| \leq m^n$, so that the regret is sublinear in T .

Acknowledgments

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

References

- Alonso, R. and Câmara, O. Persuading voters. *American Economic Review*, 106(11):3590–3605, 2016.
- Antioch, G. et al. Persuasion is now 30 per cent of US GDP: Revisiting McCloskey and Klamer after a quarter of a century. *Economic Round-up*, 2013.
- Arieli, I. and Babichenko, Y. Private bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.
- Arora, S., Lund, C., Motwani, R., Sudan, M., and Szegedy, M. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998.
- Babichenko, Y. and Barman, S. Computational aspects of private bayesian persuasion. *arXiv preprint arXiv:1603.01444*, 2016.
- Babichenko, Y. and Barman, S. Algorithmic aspects of private Bayesian persuasion. In *Innovations in Theoretical Computer Science Conference*, 2017.
- Babichenko, Y., Talgam-Cohen, I., Xu, H., and Zabarnyi, K. Regret-minimizing bayesian persuasion. *arXiv preprint arXiv:2105.13870*, 2021.
- Badanidiyuru, A., Bhawalkar, K., and Xu, H. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2545–2563, 2018.
- Bhaskar, U., Cheng, Y., Ko, Y. K., and Swamy, C. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 479–496, 2016.
- Bro Miltersen, P. and Sheffet, O. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 234–247, 2012.
- Candogan, O. Persuasion in networks: Public signals and k-cores. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 133–134, 2019.
- Castiglioni, M. and Gatti, N. Persuading voters in district-based elections. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021.
- Castiglioni, M., Celli, A., and Gatti, N. Persuading voters: It's easy to whisper, it's hard to speak loud. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pp. 1870–1877, 2020a.
- Castiglioni, M., Celli, A., Marchesi, A., and Gatti, N. Online bayesian persuasion. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Castiglioni, M., Celli, A., Marchesi, A., and Gatti, N. Signaling in bayesian network congestion games: the subtle power of symmetry. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021.
- Cheng, Y., Cheung, H. Y., Dughmi, S., Emamjomeh-Zadeh, E., Han, L., and Teng, S.-H. Mixture selection, mechanism design, and signaling. In *56th Annual Symposium on Foundations of Computer Science*, pp. 1426–1445, 2015.
- Dughmi, S. and Xu, H. Algorithmic bayesian persuasion. In *ACM STOC*, pp. 412–425, 2016.
- Dughmi, S. and Xu, H. Algorithmic persuasion with no externalities. In *ACM EC*, pp. 351–368, 2017.
- Emek, Y., Feldman, M., Gamzu, I., PaesLeme, R., and Tenenholtz, M. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation*, 2(2):1–19, 2014.
- Garber, D. Efficient online linear optimization with approximation algorithms. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 627–635. Curran Associates, Inc., 2017.
- Grötschel, M., Lovász, L., and Schrijver, A. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

- Hazan, E., Hu, W., Li, Y., and Li, Z. Online improper learning with an approximation oracle. In *Advances in Neural Information Processing Systems*, volume 31, pp. 5652–5660, 2018.
- Kakade, S. M., Kalai, A. T., and Ligett, K. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009. doi: 10.1137/070701704.
- Kamenica, E. and Gentzkow, M. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Khachiyan, L. G. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics*, 20(1):53–72, 1980.
- Mansour, Y., Slivkins, A., Syrgkanis, V., and Wu, Z. S. Bayesian exploration: Incentivizing exploration in bayesian games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 661–661, 2016.
- McCloskey, D. and Klamler, A. One quarter of GDP is persuasion. *The American Economic Review*, 85(2):191–195, 1995.
- Rabinovich, Z., Jiang, A. X., Jain, M., and Xu, H. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 645–653, 2015.
- Raz, R. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.
- Roughgarden, T. and Wang, J. R. Minimizing regret with multiple reserves. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–18, 2019.
- Schrijver, A. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.
- Sviridenko, M., Vondrák, J., and Ward, J. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4):1197–1218, 2017.
- Vasserman, S., Feldman, M., and Hassidim, A. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, pp. 660–666, 2015.
- Xu, H. On the tractability of public persuasion with no externalities. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2708–2727, 2020.
- Xu, H., Freeman, R., Conitzer, V., Dughmi, S., and Tambe, M. Signaling in bayesian stackelberg games. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, pp. 150–158, 2016.