

Signaling in Bayesian Posted Price Auctions

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

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

Giulia Romano
Politecnico di Milano
Milano, Italy
giulia.romano@polimi.it

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

ABSTRACT

We study single-item single-unit *Bayesian posted price auctions*, where buyers arrive sequentially and their valuations for the item being sold depend on a random, unknown state of nature. The seller has complete knowledge of the actual state and can send signals to the buyers so as to disclose information about it. For instance, the state of nature may reflect the condition and/or some particular features of the item, which are known to the seller only. The problem faced by the seller is about *how to partially disclose information about the state so as to maximize revenue*. Unlike classical signaling problems, in this setting, the seller must also correlate the signals being sent to the buyers with some price proposals for them. This introduces additional challenges compared to standard settings. We consider two cases: the one where the seller can only send signals *publicly* visible to all buyers, and the case in which the seller can *privately* send a different signal to each buyer. As a first step, we prove that, in both settings, the problem of maximizing the seller’s revenue does *not* admit an FPTAS unless $P = NP$, even for basic instances with a single buyer. As a result, in the rest of the paper, we focus on designing PTASs. In order to do so, we first introduce a unifying framework encompassing both public and private signaling, whose core result is a decomposition lemma that allows focusing on a finite set of possible buyers’ posteriors. This forms the basis on which our PTASs are developed. In particular, in the public signaling setting, our PTAS employs some *ad hoc* techniques based on linear programming, while our PTAS for the private setting relies on the ellipsoid method to solve an exponentially-sized LP in polynomial time. In the latter case, we need a custom approximate separation oracle, which we implement with a dynamic programming approach.

KEYWORDS

Posted Price Auctions; Bayesian Persuasion; Uncertain Agents’ Types

1 INTRODUCTION

In *posted price auctions*, the seller tries to sell an item by proposing *take-it-or-leave-it* prices to buyers arriving sequentially. Each buyer has to choose between declining the offer—without having the possibility of coming back—or accepting it, thus ending the auction.

Appears at the 1st Workshop on Learning with Strategic Agents (LSA 2022). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), N. Bishop, M. Han, L. Tran-Thanh, H. Xu, H. Zhang (chairs), May 9–10, 2022, Online. 2022.

Nowadays, posted pricing is the most used selling format in e-commerce [15], whose sales reach over \$4 trillion in 2020 [16]. Posted price auctions are ubiquitous in settings such as, for example, online travel agencies (e.g., *Expedia*), accommodation websites (e.g., *Booking.com*), and retail platforms (e.g., *Amazon* and *eBay*). As a result, growing attention has been devoted to their analysis, both in economics [22] and in computer science [1–3, 12, 14], within AI and machine learning in particular [4, 8, 11, 20, 21, 23].

We study *Bayesian* posted price auctions, where the buyers’ valuations for the item depend on a random state of nature, which is known to the seller only. By applying the *Bayesian persuasion* framework [18], we consider the case in which the seller (sender) can send signals to the buyers (receivers) so as to disclose information about the state. Thus, in a Bayesian auction, the seller does *not* only have to decide price proposals for the buyers, but also *how to partially disclose information about the state so as to maximize revenue*. Our model finds application in several real-world scenarios. For instance, in an e-commerce platform, the state of nature may reflect the condition (or quality) of the item being sold and/or some of its features. These are known to the seller only since the buyers cannot see the item given that the auction is carried out on the web.

Original Contributions. We study the problem of maximizing seller’s revenue in single-item single-unit Bayesian posted price auctions, focusing on two different settings: *public signaling*, where the signals are publicly visible to all buyers, and *private signaling*, in which the seller can send a different signal to each buyer through private communication channels. As a first negative result, we prove that, in both settings, the problem does *not* admit an FPTAS unless $P = NP$, even for basic instances with a single buyer. Then, we provide tight positive results by designing a PTAS for each setting. In order to do so, we first introduce a unifying framework encompassing both public and private signaling. Its core result is a *decomposition lemma* that allows us to focus on a finite set of buyers’ posterior beliefs over states of nature—called *q*-uniform posteriors—, rather than reasoning about signaling schemes with a (potentially) infinite number of signals. Compared to previous works on signaling, our framework has to deal with some additional challenges. The main one is that, in our model, the seller (sender) is *not* only required to choose how to send signals, but they also have to take some actions in the form of price proposals. This requires significant extensions to standard approaches based on decomposition lemmas [9, 13, 25]. The framework forms the basis on which we design our PTASs. In the public setting, it establishes a

connection between signaling schemes and probability distributions over q -uniform posteriors. This allows us to formulate the seller’s revenue-maximizing problem as an LP of polynomial size, whose objective coefficients are *not* readily available. However, they can be approximately computed in polynomial time by an algorithm for finding approximately-optimal prices in (non-Bayesian) posted price auctions, which may also be of independent interest.¹ Solving the LP with approximate coefficients then gives the desired PTAS. As for the private setting, our framework provides a connection between marginal signaling schemes of each buyer and probability distributions over q -uniform posteriors, which, to the best of our knowledge, is the first of its kind, since previous works are limited to public settings [6, 13].² Such connection allows us to formulate an LP correlating marginal signaling schemes together and with price proposals. Although the LP has an exponential number of variables, we show that it can still be approximately solved in polynomial time by means of the ellipsoid method. This requires the implementation of a problem-specific *approximate separation oracle* that can be implemented in polynomial time by means of a dynamic programming algorithm.

Related Works. The works most related to ours are those addressing signaling in second-price auctions. Emek et al. [17] provide an LP to compute an optimal public signaling scheme in the known-valuation setting, and they show that the problem is NP-hard in the Bayesian setting. Cheng et al. [13] provide a PTAS for this latter case. Finally, Badanidiyuru et al. [5] focus on the design of algorithms whose running time is independent from the number of states of nature. They initiate the study of private signaling, showing that, differently from our Bayesian posted price auctions, in second-price auctions it may introduce non-trivial equilibrium selection issues.

2 PRELIMINARIES

2.1 Bayesian Posted Price Auctions and Signaling

In a *posted price auction*, the seller tries to sell an item to a finite set $\mathcal{N} := \{1, \dots, n\}$ of buyers arriving sequentially according to a fixed ordering. W.l.o.g., we let buyer $i \in \mathcal{N}$ be the i -th buyer according to such ordering. The seller chooses a price proposal $p_i \in [0, 1]$ for each buyer $i \in \mathcal{N}$. Then, each buyer in turn has to decide whether to buy the item for the proposed price or not. Buyer $i \in \mathcal{N}$ buys only if their item valuation is at least the proposed price p_i .³ In that case, the auction ends and the seller gets revenue p_i for selling the item, otherwise the auction continues with the next buyer.

We study *Bayesian* posted price auctions, characterized by a finite set of d states of nature, namely $\Theta := \{\theta_1, \dots, \theta_d\}$. Each buyer $i \in \mathcal{N}$ has a valuation vector $v_i \in [0, 1]^d$, with $v_i(\theta)$ representing buyer i ’s valuation when the state is $\theta \in \Theta$. Each valuation v_i is independently drawn from a probability distribution \mathcal{V}_i supported

¹Notice that the objective of the LP for the public setting is *not* readily available since, as it is common in the literature, we assume that the distributions of buyers’ valuations are only accessible through an oracle providing random samples drawn i.i.d. from them.

²A notable exception is [9], which studies a specific case in between private and public signaling schemes.

³As customary in the literature, we assume that buyers always buy when they are offered a price that is equal to their valuation.

on $[0, 1]^d$. For the ease of presentation, we let $V \in [0, 1]^{n \times d}$ be the matrix of buyers’ valuations, whose entries are $V(i, \theta) := v_i(\theta)$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$.⁴ Moreover, by letting $\mathcal{V} := \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ be the collection of all distributions of buyers’ valuations, we write $V \sim \mathcal{V}$ to denote that V is built by drawing each v_i independently from \mathcal{V}_i .

We model signaling with the *Bayesian persuasion* framework by Kamenica and Gentzkow [18]. We consider the case in which the seller—having knowledge of the state of nature—acts as a *sender* by issuing signals to the buyers (the *receivers*), so as to partially disclose information about the state and increase revenue. As customary in the literature, we assume that the state is drawn from a common prior distribution $\mu \in \Delta_\Theta$, explicitly known to both the seller and the buyers.⁵ We denote by μ_θ the probability of state $\theta \in \Theta$. The seller *commits to a signaling scheme* ϕ , which is a randomized mapping from states of nature to signals for the receivers. Letting \mathcal{S}_i be the set of signals for buyer $i \in \mathcal{N}$, a signaling scheme is a function $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$, where $\mathcal{S} := \times_{i \in \mathcal{N}} \mathcal{S}_i$. An element $s \in \mathcal{S}$ —called *signal profile*—is a tuple specifying a signal for each buyer. We use s_i to refer to the i -th component of any $s \in \mathcal{S}$ (i.e., the signal for buyer i), so that $s = (s_1, \dots, s_n)$. We let $\phi_\theta(s)$ be the probability of drawing signal profile $s \in \mathcal{S}$ when the state is $\theta \in \Theta$. Furthermore, we let $\phi_i : \Theta \rightarrow \Delta_{\mathcal{S}_i}$ be the marginal signaling scheme of buyer $i \in \mathcal{N}$, with $\phi_i(\theta)$ being the marginalization of $\phi(\theta)$ with respect to buyer i ’s signals. As for general signaling schemes, $\phi_{i,\theta}(s_i) := \sum_{s' \in \mathcal{S} : s'_i = s_i} \phi_\theta(s')$ denotes the probability of drawing signal $s_i \in \mathcal{S}_i$ when the state is $\theta \in \Theta$.

Price proposals may depend on the signals being sent to the buyers. Formally, the seller *commits to a price function* $f : \mathcal{S} \rightarrow [0, 1]^n$, with $f(s) \in [0, 1]^n$ being the price vector when the signal profile is $s \in \mathcal{S}$. We assume that prices proposed to buyer i only depend on the signals sent to them, and *not* on the signals sent to other buyers. Thus, w.l.o.g., we can work with functions $f_i : \mathcal{S}_i \rightarrow [0, 1]$ defining prices for each buyer $i \in \mathcal{N}$ independently, with $f_i(s_i)$ denoting the i -th component of $f(s)$ for all $s \in \mathcal{S}$ and $i \in \mathcal{N}$.⁶




Figure 1: Interaction between the seller and the buyers.

The interaction involving the seller and the buyers goes on as follows (Figure 1): (i) the seller commits to a signaling scheme $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$ and a price function $f : \mathcal{S} \rightarrow [0, 1]^n$, and the buyers observe such commitments; (ii) the seller observes the state of nature $\theta \sim \mu$; (iii) the seller draws a signal profile $s \sim \phi(\theta)$; and (iv) the buyers arrive sequentially, with each buyer $i \in \mathcal{N}$ observing

⁴Sometimes, we also write $V_i := v_i^\top$ to denote the i -th row of matrix V , which is the valuation of buyer $i \in \mathcal{N}$.

⁵In this work, given a finite set X , we denote with Δ_X the $(|X| - 1)$ -dimensional simplex defined over the elements of X .

⁶Let us remark that our assumption on the seller’s price function ensures that a buyer does *not* get additional information about the state of nature by observing the proposed price, since the latter only depends on the signal which is revealed to them anyway.

their signal s_i and being proposed price $f_i(s_i)$. Then, each buyer rationally updates their prior belief over states according to Bayes rule, and buys the item only if their expected valuation for the item is greater than or equal to the offered price. The interaction terminates whenever a buyer decides to buy the item or there no more buyers arriving. The following paragraph formally defines the elements involved in step (iv).

Buyers' Posteriors. In step (iv), a buyer $i \in \mathcal{N}$ receiving a signal $s_i \in \mathcal{S}_i$ infers a posterior belief over states (also called *posterior*), which we denote by $\xi_{i,s_i} \in \Delta_\Theta$, with $\xi_{i,s_i}(\theta)$ being the posterior probability of state $\theta \in \Theta$. Formally,

$$\xi_{i,s_i}(\theta) := \frac{\mu_\theta \phi_{i,\theta}(s_i)}{\sum_{\theta' \in \Theta} \mu_{\theta'} \phi_{i,\theta'}(s_i)}. \quad (1)$$

Thus, after receiving signal $s_i \in \mathcal{S}_i$, buyer i 's expected valuation for the item is $\sum_{\theta \in \Theta} v_i(\theta) \xi_{i,s_i}(\theta)$, and the buyer buys it only if such value is at least as large as the price $f_i(s_i)$. In the following, given a signal profile $s \in \mathcal{S}$, we denote by ξ_s a tuple defining all buyers' posteriors resulting from observing signals in s ; formally, $\xi_s := (\xi_{1,s_1}, \dots, \xi_{n,s_n})$.

Distributions on Posteriors. In single-receiver Bayesian persuasion models, it is oftentimes useful to represent signaling schemes as convex combinations of the posteriors they can induce. In our setting, a marginal signaling scheme $\phi_i : \Theta \rightarrow \Delta_{\mathcal{S}_i}$ of buyer $i \in \mathcal{N}$ induces a probability distribution γ_i over posteriors in Δ_Θ , with $\gamma_i(\xi_i)$ denoting the probability of posterior $\xi_i \in \Delta_\Theta$. Formally, it holds that

$$\gamma_i(\xi_i) := \sum_{s_i \in \mathcal{S}_i: \xi_{i,s_i} = \xi_i} \sum_{\theta \in \Theta} \mu_\theta \phi_{i,\theta}(s_i).$$

Intuitively, $\gamma_i(\xi_i)$ denotes the probability that buyer i has posterior ξ_i . Indeed, it is possible to directly reason about distributions γ_i rather than marginal signaling schemes, provided that such distributions are *consistent* with the prior. Formally, by letting $\text{supp}(\gamma_i) := \{\xi_i \in \Delta_\Theta \mid \gamma_i(\xi_i) > 0\}$ be the support of γ_i , it must be required that

$$\sum_{\xi_i \in \text{supp}(\gamma_i)} \gamma_i(\xi_i) \xi_i(\theta) = \mu_\theta \quad \forall \theta \in \Theta. \quad (2)$$

2.2 Computational Problems

We focus on the problem of computing a signaling scheme $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$ and a price function $f : \mathcal{S} \rightarrow [0, 1]^n$ that maximize the seller's expected revenue, considering both public and private signaling settings.⁷

We denote by $\text{REV}(\mathcal{V}, p, \xi)$ the expected revenue of the seller when the distributions of buyers' valuations are given by $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, the proposed prices are defined by the vector $p \in [0, 1]^n$, and the buyers' posteriors are those specified by the tuple $\xi =$

⁷Formally, a signaling scheme $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$ is *public* if: (i) $\mathcal{S}_i = \mathcal{S}_j$ for all $i, j \in \mathcal{N}$; and (ii) for every $\theta \in \Theta$, $\phi_\theta(s) > 0$ only for signal profiles $s \in \mathcal{S}$ such that $s_i = s_j$ for $i, j \in \mathcal{N}$. Since, given a signal profile $s \in \mathcal{S}$, under a public signaling scheme all the buyers always share the same posterior (i.e., $\xi_{i,s_i} = \xi_{j,s_j}$ for all $i, j \in \mathcal{N}$), we overload notation and sometimes use $\xi_s \in \Delta_\Theta$ to denote the unique posterior appearing in $\xi_s = (\xi_{1,s_1}, \dots, \xi_{n,s_n})$. Similarly, in the public setting, given a posterior $\xi \in \Delta_\Theta$ we sometimes write ξ in place of a tuple of n copies of ξ .

(ξ_1, \dots, ξ_n) containing a posterior $\xi_i \in \Delta_\Theta$ for each buyer $i \in \mathcal{N}$. Then, the seller's expected revenue is:

$$\sum_{\theta \in \Theta} \mu_\theta \sum_{s \in \mathcal{S}} \phi_\theta(s) \text{REV}(\mathcal{V}, f(s), \xi_s).$$

In the following, we denote by OPT the value of the seller's expected revenue for a revenue-maximizing (ϕ, f) pair.

In this work, we assume that algorithms have access to a black-box oracle to sample buyers' valuations according to the probability distributions specified by \mathcal{V} (rather than actually knowing such distributions). Thus, we look for algorithms that output pairs (ϕ, f) such that

$$\mathbb{E} \left[\sum_{\theta \in \Theta} \mu_\theta \sum_{s \in \mathcal{S}} \phi_\theta(s) \text{REV}(\mathcal{V}, f(s), \xi_s) \right] \geq OPT - \lambda,$$

where $\lambda \geq 0$ is an additive error. Notice that the expectation above is with respect to the randomness of the algorithm, which originates from using the black-box sampling oracle.

3 HARDNESS OF SIGNALING WITH A SINGLE BUYER

We start with a negative result: there is no FPTAS for the problem of computing a revenue-maximizing (ϕ, f) pair unless $P = NP$, in both public and private signaling settings. Our result holds even in the basic case with only one buyer, where public and private signaling are equivalent. Notice that, in the reduction that we use to prove our result, we assume that the support of the distribution of valuations of the (single) buyer is finite and that such distribution is perfectly known to the seller. This represents an even simpler setting than that in which the seller has only access to a black-box oracle returning samples drawn from the buyer's distribution of valuations. The result formally reads as follows:⁸

Theorem 1. *There is no additive FPTAS for the problem of computing a revenue-maximizing (ϕ, f) pair unless $P = NP$, even when there is a single buyer.*

4 UNIFYING PUBLIC AND PRIVATE SIGNALING

In this section, we introduce a general mathematical framework related to buyers' posteriors and distributions over them, proving some results that will be crucial in the rest of this work, both in public and private signaling scenarios.

One of the main difficulties in computing sender-optimal signaling schemes is that they might need a (potentially) infinite number of signals, resulting in infinitely-many receiver's posteriors. The trick commonly used to circumvent this issue in settings with a finite number of valuations is to use direct signals, which explicitly specify action recommendations for each receiver's valuation [7, 10]. However, in our auction setting, this solution is *not* viable, since a direct signal for a buyer $i \in \mathcal{N}$ should represent a recommendation for every possible $v_i \in [0, 1]^d$, and these are infinitely many. An alternative technique, which can be employed in our setting, is to restrict the number of possible posteriors.

⁸The proofs of all the results can be found in the Appendix.

Our core idea is to focus on a small set of posteriors, which are those encoded as particular q -uniform probability distributions, as formally stated in the following definition.⁹

Definition 1 (q -uniform posterior). *A posterior $\xi \in \Delta_\Theta$ is q -uniform if it can be obtained by averaging the elements of a multiset defined by $q \in \mathbb{N}_{>0}$ canonical basis vectors of \mathbb{R}^d .*

We denote the set of all q -uniform posteriors as $\Xi^q \subset \Delta_\Theta$. Notice that the set Ξ^q has size $|\Xi^q| = O(d^q)$.

The existence of an approximately-optimal signaling scheme that only uses q -uniform posteriors is usually proved by means of so-called *decomposition lemmas* (see [9, 13, 25]). The goal of these lemmas is to show that, given some signaling scheme encoded as a distribution over posteriors, it is possible to obtain a new signaling scheme whose corresponding distribution is supported only on q -uniform posteriors, and such that the sender's utility only decreases by a small amount. At the same time, these lemmas must also ensure that the distribution over posteriors corresponding to the new signaling scheme is still consistent (according to Equation (2)).

The main result of our framework (Theorem 2) is a decomposition lemma that is suitable for our setting. Before stating the result, we need to introduce some preliminary definitions.

Definition 2 ((α, ϵ) -decreasing distribution). *Let $\alpha, \epsilon > 0$. A probability distribution γ over Δ_Θ is (α, ϵ) -decreasing around a given posterior $\xi \in \Delta_\Theta$ if the following condition holds for every matrix $V \in [0, 1]^{n \times d}$ of buyers' valuations:*

$$\Pr_{\xi \sim \gamma} \left\{ V_i \tilde{\xi} \geq V_i \xi - \epsilon \right\} \geq 1 - \alpha \quad \forall i \in \mathcal{N}.$$

Intuitively, a probability distribution γ as in Definition 2 can be interpreted as a perturbation of the given posterior ξ such that, with high probability, buyers' expected valuations in γ are at most ϵ less than those in posterior ξ .¹⁰

The second definition we need is about functions mapping vectors in $[0, 1]^n$ —defining a valuation for each buyer—to seller's revenues. For instance, one such function could be the seller's revenue given price vector $p \in [0, 1]^n$. In particular, we define the stability of a function g compared to another function h . Intuitively, g is stable compared to h if the value of g , in expectation over buyers' valuations and posteriors drawn from a probability distribution γ that is (α, ϵ) -decreasing around ξ , is “close” to the value of h given ξ , in expectation over buyers' valuations.¹¹ Formally:

Definition 3 ($(\delta, \alpha, \epsilon)$ -stability). *Let $\alpha, \epsilon, \delta > 0$. Given a posterior $\xi \in \Delta_\Theta$, some distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, and two functions $g, h : [0, 1]^n \rightarrow [0, 1]$, g is $(\delta, \alpha, \epsilon)$ -stable compared to h for (ξ, \mathcal{V}) if, for every probability distribution γ over Δ_Θ that is (α, ϵ) -decreasing*

⁹In all the definitions and results of this section (Section 4), we denote by $\xi \in \Delta_\Theta$ a generic posterior common to all the buyers and with γ a probability distribution over Δ_Θ (i.e. over posteriors).

¹⁰Definition 2 is similar to analogous ones in the literature [9, 25], where the distance is usually measured in both directions, as $|V_i \tilde{\xi} - V_i \xi| \leq \epsilon$. We look only at the direction of decreasing values, since in our setting, if a buyer's valuation increases, then the seller's revenue also increases.

¹¹The notion of compared stability has been already used [9, 13]. However, previous works consider the case in which g is a relaxation of h . Instead, our definition is conceptually different, as g and h represent two different functions corresponding to different price vectors of the seller.

around ξ , it holds:

$$\mathbb{E}_{\xi \sim \gamma, V \sim \mathcal{V}} \left[g(V \tilde{\xi}) \right] \geq (1 - \alpha) \mathbb{E}_{V \sim \mathcal{V}} \left[h(V \xi) \right] - \delta \epsilon.$$

Now, we are ready to state our main result. We show that, for any buyer's posterior $\xi \in \Delta_\Theta$, if a function g is stable compare to h , then there exists a suitable probability distribution over q -uniform posteriors such that the expected value of g given such distribution is “close” to that of h given ξ .

Theorem 2. *Let $\alpha, \epsilon, \delta > 0$, and set $q := \frac{32}{\epsilon^2} \log \frac{4}{\alpha}$. Given a posterior $\xi \in \Delta_\Theta$, some distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, and two functions $g, h : [0, 1]^n \rightarrow [0, 1]$, if g is $(\delta, \alpha, \epsilon)$ -stable compared to h for (ξ, \mathcal{V}) , then there exists $\gamma \in \Delta_{\Xi^q}$ such that, for every $\theta \in \Theta$, $\sum_{\tilde{\xi} \in \text{supp}(\gamma)} \gamma(\tilde{\xi}) \tilde{\xi}(\theta) = \xi(\theta)$ and*

$$\mathbb{E}_{\xi \sim \gamma, V \sim \mathcal{V}} \left[\tilde{\xi}(\theta) g(V \tilde{\xi}) \right] \geq \xi(\theta) \left[(1 - \alpha) \mathbb{E}_{V \sim \mathcal{V}} \left[h(V \xi) \right] - \delta \epsilon \right]. \quad (3)$$

The crucial feature of Theorem 2 is that Equation (3) holds for every state. This is fundamental for proving our results in the private signaling scenario. On the other hand, with public signaling, we will make use of the following (weaker) corollary, obtained by summing Equation (3) over all $\theta \in \Theta$.

Corollary 1. *Let $\alpha, \epsilon, \delta > 0$, and set $q := \frac{32}{\epsilon^2} \log \frac{4}{\alpha}$. Given a posterior $\xi \in \Delta_\Theta$, some distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, and two functions $g, h : [0, 1]^n \rightarrow [0, 1]$, if g is $(\delta, \alpha, \epsilon)$ -stable compared to h for (ξ, \mathcal{V}) , then there exists $\gamma \in \Delta_{\Xi^q}$ such that, for every $\theta \in \Theta$, $\sum_{\tilde{\xi} \in \text{supp}(\gamma)} \gamma(\tilde{\xi}) \tilde{\xi}(\theta) = \xi(\theta)$ and*

$$\mathbb{E}_{\xi \sim \gamma, V \sim \mathcal{V}} \left[g(V \tilde{\xi}) \right] \geq (1 - \alpha) \mathbb{E}_{V \sim \mathcal{V}} \left[h(V \xi) \right] - \delta \epsilon. \quad (4)$$

5 WARMING UP: NON-BAYESIAN AUCTIONS

In this section, we focus on non-Bayesian posted price auctions, proving some results that will be useful in the rest of the paper.¹² In particular, we study what happens to the seller's expected revenue when buyers' valuations are “slightly decreased”, proving that the revenue also decreases, but only by a small amount. This result will be crucial when dealing with public signaling, and it also allows to design a poly-time algorithm for finding approximately-optimal price vectors in non-Bayesian auctions, as we show at the end of this section.

In the following, we extensively use distributions of buyers' valuations as specified in the definition below.

Definition 4. *Given $\epsilon > 0$, we denote by $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ and $\mathcal{V}^\epsilon = \{\mathcal{V}_i^\epsilon\}_{i \in \mathcal{N}}$ two collections of distributions of buyers' valuations such that, for every price vector $p \in [0, 1]^n$,*

$$\Pr_{v_i \sim \mathcal{V}_i^\epsilon} \{v_i \geq p_i - \epsilon\} \geq \Pr_{v_i \sim \mathcal{V}_i} \{v_i \geq p_i\} \quad \forall i \in \mathcal{N}.$$

¹²When we study non-Bayesian posted price auctions, we stick to our notation, with the following differences: valuations are scalars rather than vectors, namely $v_i \in [0, 1]$; distributions \mathcal{V}_i are supported on $[0, 1]$ rather than $[0, 1]^d$; the matrix V is indeed a column vector whose components are buyers' valuations; and the price function f is replaced by a single price vector $p \in [0, 1]^n$, with its i -th component p_i being the price for buyer $i \in \mathcal{N}$. Moreover, we continue to use the notation Rev to denote seller's revenues, dropping the dependence on the tuple of posteriors. Thus, in a non-Bayesian auction in which the distributions of buyers' valuations are $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, the notation $\text{Rev}(\mathcal{V}, p)$ simply denotes the seller's expected revenue by selecting a price vector $p \in [0, 1]^n$.

Intuitively, valuations drawn from \mathcal{V}^ϵ are “slightly decreased” with respect to those drawn from \mathcal{V} , since the probability with which any buyer $i \in \mathcal{N}$ buys the item at the (reduced) price $[p_i - \epsilon]_+$ when their valuation is drawn from \mathcal{V}_i^ϵ is at least as large as the probability of buying at price p_i when their valuation is drawn from \mathcal{V}_i .¹³

Our main contribution in this section (Lemma 2) is to show that $\max_{p \in [0,1]^n} \text{REV}(\mathcal{V}^\epsilon, p) \geq \max_{p \in [0,1]^n} \text{REV}(\mathcal{V}, p) - \epsilon$. By letting $p^* \in \arg \max_{p \in [0,1]^n} \text{REV}(\mathcal{V}, p)$ be any revenue-maximizing price vector under distributions \mathcal{V} , one may naïvely think that, since under distributions \mathcal{V}^ϵ and price vector $[p^* - \epsilon]_+$ each buyer would buy the item at least with the same probability as with distributions \mathcal{V} and price vector p^* , while paying a price that is only ϵ less, then $\text{REV}(\mathcal{V}^\epsilon, [p^* - \epsilon]_+) \geq \text{REV}(\mathcal{V}, p^*) - \epsilon$, proving the result. However, this line of reasoning does *not* work, as shown by Example 1 in Appendix A. The crucial feature of Example 1 is that there exists a p^* in which one buyer is offered a price that is too low, and, thus, the seller prefers *not* to sell the item to them, but rather to a following buyer. This prevents a direct application of the line of reasoning outlined above, as it shows that incrementing the probability with which a buyer buys is *not* always beneficial. One could circumvent this issue by considering a p^* such that the seller is never upset if some buyer buys. In other words, it must be such that each buyer is proposed a price that is at least as large as the seller’s expected revenue in the posted price auction restricted to the following buyers. Next, we show that there always exists a p^* with such desirable property.

Letting $\text{REV}_{>i}(\mathcal{V}, p)$ be the seller’s revenue for price vector $p \in [0, 1]^n$ and distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ in the auction restricted to buyers $j \in \mathcal{N} : j > i$, we prove the following:

Lemma 1. *For any $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, there exists a revenue-maximizing price vector $p^* \in \arg \max_{p \in [0,1]^n} \text{REV}(\mathcal{V}, p)$ such that $p_i^* \geq \text{REV}_{>i}(\mathcal{V}, p^*)$ for every buyer $i \in \mathcal{N}$.*

The proof of Lemma 2 builds upon the existence of a revenue-maximizing price vector $p^* \in [0, 1]^n$ as in Lemma 1 and the fact that, under distributions \mathcal{V}^ϵ , the probability with which each buyer buys the item given price vector $[p^* - \epsilon]_+$ is greater than that with which they would buy given p^* . Since the seller’s expected revenue is larger when a buyer buys compared to when they do *not* buy (as $p_i^* \geq \text{REV}_{>i}(\mathcal{V}, p^*)$), the seller’s expected revenue decreases by at most ϵ .

Lemma 2. *Given $\epsilon > 0$, let $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ and $\mathcal{V}^\epsilon = \{\mathcal{V}_i^\epsilon\}_{i \in \mathcal{N}}$ satisfying the conditions of Definition 4. Then, $\max_{p \in [0,1]^n} \text{REV}(\mathcal{V}^\epsilon, p) \geq \max_{p \in [0,1]^n} \text{REV}(\mathcal{V}, p) - \epsilon$.*

Lemma 2 will be useful to prove Lemma 3 and to show the compared stability of a suitably-defined function that is used to design a PTAS in the public signaling scenario.

Finding Approximately-Optimal Prices. Algorithm 1 computes (in polynomial time) an approximately-optimal price vector for any non-Bayesian posted price auction. It samples $K \in \mathbb{N}_{>0}$ matrices of buyers’ valuations, each one drawn according to the distributions \mathcal{V} . Then, it finds an optimal price vector p in the discretized set

¹³In this work, given $x \in \mathbb{R}$, we let $[x]_+ := \max\{x, 0\}$. We extend the $[\cdot]_+$ operator to vectors by applying it component-wise.

Algorithm 1 FIND-APX-PRICES

Inputs: # of samples $K \in \mathbb{N}_{>0}$; # of discretization steps $b \in \mathbb{N}_{>0}$

- 1: **for** $i \in \mathcal{N}$ **do**
- 2: **for** $k = 1, \dots, K$ **do**
- 3: $v_i^k \leftarrow$ Sample buyer i ’s valuation using oracle for \mathcal{V}_i
- 4: $\mathcal{V}_i^K \leftarrow$ Empirical distribution of the K i.i.d. samples v_i^k
- 5: $\mathcal{V}^K \leftarrow \{\mathcal{V}_i^K\}_{i \in \mathcal{N}}$; $p \leftarrow \mathbf{0}_n$; $r \leftarrow 0$
- 6: **for** $i = n, \dots, 1$ (in reversed order) **do**
- 7: $p_i \leftarrow \arg \max_{p'_i \in \mathcal{P}^b} \Pr_{v_i \sim \mathcal{V}_i^K} \{v_i \geq p'_i\} + (1 - \Pr_{v_i \sim \mathcal{V}_i^K} \{v_i \geq p'_i\})r$
- 8: $r \leftarrow p_i \Pr_{v_i \sim \mathcal{V}_i^K} \{v_i \geq p_i\} + (1 - \Pr_{v_i \sim \mathcal{V}_i^K} \{v_i \geq p_i\})r$
- 9: **return** (p, r)

\mathcal{P}^b , assuming that buyers’ valuations are drawn according to the empirical distribution resulting from the sampled matrices.¹⁴ This last step can be done by backward induction, as it is well known in the literature (see, e.g., [24]). The following Lemma 3 establishes the correctness of Algorithm 1, also providing a bound on its running time. The key ideas of its proof are: (i) the sampling procedure constructs a good estimation of the actual distributions of buyers’ valuations; and (ii) even if the algorithm only considers discretized prices, the components of the computed price vector are at most $1/b$ less than those of an optimal (unconstrained) price vector. As shown in the proof, this is strictly related to reducing buyer’s valuations by $\frac{1}{b}$. Thus, it follows by Lemma 2 that the seller’s expected revenue is at most $1/b$ less than the optimal one.

Lemma 3. *For any $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ and $\epsilon, \tau > 0$, there exist $K \in \text{poly}\left(n, \frac{1}{\epsilon}, \log \frac{1}{\tau}\right)$ and $b \in \text{poly}\left(\frac{1}{\epsilon}\right)$ such that, with probability at least $1 - \tau$, Algorithm 1 returns (p, r) satisfying $\text{REV}(\mathcal{V}, p) \geq \max_{p' \in [0,1]^n} \text{REV}(\mathcal{V}, p') - \epsilon$ and $r \in [\text{REV}(\mathcal{V}, p) - \epsilon, \text{REV}(\mathcal{V}, p) + \epsilon]$ in time $\text{poly}\left(n, \frac{1}{\epsilon}, \log \frac{1}{\tau}\right)$.*

6 PUBLIC SIGNALING

In the following, we design a PTAS for computing a revenue-maximizing (ϕ, f) pair in the public signaling setting. Notice that this positive result is tight by Theorem 1.

As a first intermediate result, we prove the compared stability of suitably-defined functions, which are intimately related to the seller’s revenue. In particular, for every price vector $p \in [0, 1]^n$, we conveniently let $g_p : [0, 1]^n \rightarrow [0, 1]$ be a function that takes a vector of buyers’ valuations and outputs the seller’s expected revenue achieved by selecting p when the buyers’ valuations are those specified as input. The following Lemma 4 shows that, given some distributions of buyers’ valuations \mathcal{V} and a posterior $\xi \in \Delta_\Theta$, there always exists a price vector $p \in [0, 1]^n$ such that g_p is stable compared with $g_{p'}$ for every other $p' \in [0, 1]^n$. This result crucially allows us to decompose any posterior $\xi \in \Delta_\Theta$ by means of the decomposition lemma in Corollary 1, while guaranteeing a small loss in terms of seller’s expected revenue.

Lemma 4. *Given $\alpha, \epsilon > 0$, a posterior $\xi \in \Delta_\Theta$, and some distributions of buyers’ valuations $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, there exists $p \in [0, 1]^n$ such*

¹⁴In this work, for a discretization step $b \in \mathbb{N}_{>0}$, we let $\mathcal{P}^b \subset [0, 1]$ be the set of prices multiples of $1/b$, while $\mathcal{P}^b := \times_{i \in \mathcal{N}} \mathcal{P}^b$.

that, for every other $p' \in [0, 1]^n$, the function g_p is $(1, \alpha, \epsilon)$ -stable compared with $g_{p'}$ for (ξ, \mathcal{V}) .

Our PTAS leverages the fact that public signaling schemes can be represented as probability distributions over buyers' posteriors (recall that, in the public signaling setting, all the buyers share the same posterior, as they all observe the same signal). In particular, the algorithm returns a pair (γ, f°) , where γ is a probability distribution over Δ_Θ satisfying consistency constraints (see Equation (2)), while $f^\circ : \Delta_\Theta \rightarrow [0, 1]^n$ is a function mapping each posterior to a price vector. In single-receiver settings, it is well known (see Subsection 2.1) that using distributions over posteriors rather than signaling schemes ϕ is without loss of generality. The following lemma shows that the same holds in our case, *i.e.*, given a pair (γ, f°) , it is always possible to obtain a pair (ϕ, f) providing the seller with the same expected revenue.

Lemma 5. *Given a pair (γ, f°) , where γ is a probability distribution over Δ_Θ with $\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \xi(\theta) = \mu_\theta$ for all $\theta \in \Theta$ and $f^\circ : \Delta_\Theta \rightarrow [0, 1]^n$, there is a pair (ϕ, f) s.t.*

$$\sum_{\theta \in \Theta} \mu_\theta \sum_{s \in \mathcal{S}} \phi_\theta(s) \text{REV}(\mathcal{V}, f(s), \xi_s) = \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \text{REV}(\mathcal{V}, f^\circ(\xi), \xi).$$

Next, we show that, in order to find an approximately-optimal pair (γ, f°) , we can restrict the attention to q -uniform posteriors (with q suitably defined). First, we introduce the following LP that computes an optimal probability distribution restricted over q -uniform posteriors.

$$\max_{\gamma \in \Delta_{\Xi^q}} \sum_{\xi \in \Xi^q} \gamma(\xi) \max_{p \in [0, 1]^n} \text{REV}(\mathcal{V}, p, \xi) \text{ s.t.} \quad (5a)$$

$$\sum_{\xi \in \Xi^q} \gamma(\xi) \xi(\theta) = \mu_\theta \quad \forall \theta \in \Theta. \quad (5b)$$

The following Lemma 6 shows the optimal value of LP 5 is “close” to *OPT*. Its proof is based on the following core idea. Given the signaling scheme ϕ in a revenue-maximizing pair (ϕ, f) , letting γ be the distribution over Δ_Θ induced by ϕ , we can decompose each posterior in the support of γ according to Corollary 1. Then, the obtained distributions over q -uniform posteriors are consistent according to Equation (2), and, thus, they satisfy Constraints (5b). Moreover, since such distributions are also decreasing around the decomposed posteriors, by Lemma 4 each time a posterior is decomposed there exists a price vector resulting in a small revenue loss. These observations allow us to conclude that the seller’s expected revenue provided by an optimal solution to LP 5 is within some small additive loss of *OPT*.

Lemma 6. *Given $\eta > 0$ and letting $q = \frac{1}{\eta^2} 128 \log \frac{6}{\eta}$, an optimal solution to LP 5 has value at least $OPT - \eta$.*

Finally, we are ready to provide our PTAS. Its main idea is to solve LP 5 (of polynomial size) for the value of q in Lemma 6. This results in a small revenue loss. The last part missing for the algorithm is computing the terms appearing in the objective of LP 5, *i.e.*, a revenue-maximizing price vector (together with its revenue) for every q -uniform posterior. In order to do so, we can use Algorithm 1 (see also Lemma 3), which allows us to obtain in polynomial time good approximations of such price vectors, with high probability.

Theorem 3. *There exists an additive PTAS for computing a revenue-maximizing (ϕ, f) pair with public signaling.*

7 PRIVATE SIGNALING

With private signaling, computing a (ϕ, f) pair amounts to specifying a pair (ϕ_i, f_i) for each buyer $i \in \mathcal{N}$ —composed by a marginal signaling scheme $\phi_i : \Theta \rightarrow \Delta_{\mathcal{S}_i}$ and a price function $f_i : \mathcal{S}_i \rightarrow [0, 1]$ for buyer i —, and, then, correlating the ϕ_i so as to obtain a (non-marginal) signaling scheme $\phi : \Theta \rightarrow \Delta_{\mathcal{S}}$. We leverage this fact to design our PTAS.

In Subsection 7.1, we first show that it is possible to restrict the set of marginal signaling schemes of a given buyer $i \in \mathcal{N}$ to those encoded as distributions over q -uniform posteriors, as we did with public signaling. Then, we provide an LP formulation for computing an approximately-optimal (ϕ, f) pair, dealing with the challenge of correlating marginal signaling schemes in a non-trivial way. Finally, in Subsection 7.2, we show how to compute a solution to the LP in polynomial time, which requires the application of the ellipsoid method in a non-trivial way, due to the features of the formulation.

7.1 LP for Approximate Signaling Schemes

Before providing the LP, we show that restricting marginal signaling schemes to q -uniform posteriors results in a buyer’s behavior which is similar to the one with arbitrary posteriors. This amounts to showing that suitably-defined functions related to the probability of buying are comparatively stable.

For $i \in \mathcal{N}$ and $p_i \in [0, 1]$, let $g_{i,p_i} : [0, 1]^n \rightarrow \{0, 1\}$ be a function that takes as input a vector of buyers’ valuations and outputs 1 if and only if $v_i \geq p_i$ (otherwise it outputs 0).

Lemma 7. *Given $\alpha, \epsilon > 0$ and some distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, for every buyer $i \in \mathcal{N}$, posterior $\xi_i \in \Delta_\Theta$, and price $p_i \in [0, 1]$, the function $g_{i, [p_i - \epsilon]_+}$ is $(0, \alpha, \epsilon)$ -stable compared with g_{i, p_i} for (ξ_i, \mathcal{V}) .*

The following remark will be crucial for proving Lemma 9. It shows that, if for every $i \in \mathcal{N}$ we decompose buyer i ’s posterior $\xi_i \in \Delta_\Theta$ by means of a distribution over q -uniform posteriors (α, ϵ)-decreasing around ξ_i , then the probability with which buyer i buys only decreases by a small amount.¹⁵

Remark 1. *Lemma 7 and Theorem 2 imply that, given a tuple of posteriors $\xi = (\xi_1, \dots, \xi_n) \in \times_{i \in \mathcal{N}} \Delta_\Theta$ and some distributions $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$, for every buyer $i \in \mathcal{N}$ and price $p_i \in [0, 1]$, there exists $\gamma_i \in \Delta_{\Xi^q}$ with $q = \frac{32}{\epsilon^2} \log \frac{4}{\alpha}$ s.t.*

$$\mathbb{E}_{\xi_i \sim \gamma_i} \left[\tilde{\xi}_i(\theta) \Pr_{V \sim \mathcal{V}} \left\{ V_i \tilde{\xi}_i \geq [p_i - \epsilon]_+ \right\} \right] \geq \xi_i(\theta) (1 - \alpha) \Pr_{V \sim \mathcal{V}} \{ V_i \xi_i \geq p_i \}$$

and $\sum_{\xi_i \in \Xi^q} \gamma_i(\tilde{\xi}_i) \tilde{\xi}_i(\theta) = \xi_i(\theta)$ for all $\theta \in \Theta$.

Next, we show that an approximately-optimal pair (ϕ, f) can be found by solving LP 6 instantiated with suitably-defined $q \in \mathbb{N}_{>0}$ and $b \in \mathbb{N}_{>0}$. LP 6 employs:

¹⁵In this section, for the ease of presentation, we abuse notation and use Ξ_i^q to denote the (all equal) sets of q -uniform posteriors (Definition 1), one per buyer $i \in \mathcal{N}$, while $\Xi^q := \times_{i \in \mathcal{N}} \Xi_i^q$ is the set of tuples $\xi = (\xi_1, \dots, \xi_n)$ specifying a $\xi_i \in \Xi_i^q$ for each $i \in \mathcal{N}$.

- Variables γ_{i,ξ_i} (for $i \in \mathcal{N}$ and $\xi_i \in \Xi_i^q$), which encode the distributions over posteriors representing the marginal signaling schemes $\phi_i : \Theta \rightarrow \Delta_{\mathcal{S}_i}$ of the buyers.
- Variables t_{i,ξ_i,p_i} (for $i \in \mathcal{N}$, $\xi_i \in \Xi_i^q$, and $p_i \in P^b$), with t_{i,ξ_i,p_i} encoding the probability that the seller offers price p_i to buyer i and buyer i 's posterior is ξ_i .
- Variables $y_{\theta,\xi,p}$ (for $\theta \in \Theta$, $\xi \in \Xi^q$, and $p \in P^b$), with $y_{\theta,\xi,p}$ encoding the probability that the state is θ , the buyers' posteriors are those specified by ξ , and the prices that the seller offers to the buyers are those given by p .

$$\max_{\gamma, t, y \geq 0} \sum_{\theta \in \Theta} \sum_{\xi \in \Xi^q} \sum_{p \in P^b} y_{\theta,\xi,p} \text{REV}(\mathcal{V}, p, \xi) \quad \text{s.t.} \quad (6a)$$

$$\xi_i(\theta) t_{i,\xi_i,p_i} = \sum_{\xi' \in \Xi^q: \xi'_i = \xi_i} \sum_{p' \in P^b: p'_i = p_i} y_{\theta,\xi',p'} \quad \forall \theta \in \Theta, \forall i \in \mathcal{N}, \forall \xi_i \in \Xi_i^q, \forall p_i \in P^b \quad (6b)$$

$$\sum_{p_i \in P^b} t_{i,\xi_i,p_i} = \gamma_{i,\xi_i} \quad \forall i \in \mathcal{N}, \forall \xi_i \in \Xi_i^q \quad (6c)$$

$$\sum_{\xi_i \in \Xi_i^q} \gamma_{i,\xi_i} \xi_i(\theta) = \mu_\theta \quad \forall i \in \mathcal{N}, \forall \theta \in \Theta. \quad (6d)$$

Variables t_{i,ξ_i,p_i} represent marginal signaling schemes, allowing for multiple signals inducing the same posterior. This is needed since signals may correspond to different price proposals.¹⁶ One may think of marginal signaling schemes in LP 6 as if they were using signals defined as pairs $s_i = (\xi_i, p_i)$, with the convention that $f_i(s_i) = p_i$. Variables $y_{\theta,\xi,p}$ and Constraints (6b) ensure that marginal signaling schemes are correctly correlated together, by directly working in the domain of the distributions over posteriors.

To show that an optimal solution to LP 6 provides an approximately-optimal (ϕ, f) pair, we need the following two lemmas. Lemma 8 proves that, given a feasible solution to LP 6, we can recover a pair (ϕ, f) providing the seller with an expected revenue equal to the value of the LP solution. Lemma 9 shows that the optimal value of LP 6 is "close" to *OPT*. These two lemmas imply that an approximately-optimal (ϕ, f) pair can be computed by solving LP 6.

Lemma 8. *Given a feasible solution to LP 6, it is possible to recover a pair (ϕ, f) that provides the seller with an expected revenue equal to the value of the solution.*

Lemma 9. *For every $\eta > 0$, there exist $b(\eta), q(\eta) \in \mathbb{N}_{>0}$ such that LP 6 has optimal value at least $OPT - \eta$.*

7.2 PTAS

We provide an algorithm that approximately solves LP 6 in polynomial time, which completes our PTAS for computing a revenue-maximizing pair (ϕ, f) in the private setting. The core idea of our algorithm is to apply the ellipsoid method on the dual of LP 6.¹⁷

¹⁶Notice that, in a classical setting in which the sender does *not* have to propose a price (or, in general, select some action after sending signals), there always exists a signaling scheme with no pair of signals inducing the same posterior. Indeed, two signals that induce the same posterior can always be joined into a single signal. This is *not* the case in our setting, where we can only join signals that induce the same posterior and correspond to the same price.

¹⁷To be precise, we apply the ellipsoid method to the dual of a relaxed version of LP 6, since we need an over-constrained dual. More details on these technicalities can be found in the Appendix.

In particular, our implementation of the ellipsoid algorithm uses an approximate separation oracle that needs to solve the following optimization problem.

Definition 5 (MAX-LINREV). *Given some distributions of buyers' valuations $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ such that each \mathcal{V}_i has finite support and a vector $w \in [0, 1]^{n \times |\Xi_i^q| \times |P^b|}$, solve*

$$\max_{\xi \in \Xi^q, p \in P^b} \text{REV}(\mathcal{V}, p, \xi) + \sum_{i \in \mathcal{N}} w_{i,\xi_i,p_i}.$$

As a first step, we provide an FPTAS for MAX-LINREV using a dynamic programming approach. This will be the main building block of our approximate separation oracle.¹⁸

The FPTAS works as follows. Given an error tolerance $\delta > 0$, it first defines a step size $\frac{1}{c}$, with $c = \lceil \frac{n}{\delta} \rceil$, and builds a set $A = \{0, \frac{1}{c}, \frac{2}{c}, \dots, n\}$ of possible discretized values for the linear term appearing in the MAX-LINREV objective. Then, for every buyer $i \in \mathcal{N}$ (in reversed order) and value $a \in A$, the algorithm computes $M(i, a)$, which is an approximation of the largest seller's revenue provided by a pair (ξ, p) when considering buyers i, \dots, n only, and restricted to pairs (ξ, p) such that the inequality $\sum_{j \in \mathcal{N}: j \geq i} w_{j,\xi_j,p_j} \geq a$ is satisfied. By letting $z_i := \Pr_{v_i \sim \mathcal{V}_i} \{v_i^\top \xi_i \geq p_i\}$, the value $M(i, a)$ can be defined by the following recursive formula:¹⁹

$$M(i, a) := \max_{\substack{\xi_i \in \Xi_i^q, p_i \in P^b \\ a' \in A: w_{i,\xi_i,p_i} + a' \geq a}} z_i p_i + (1 - z_i) M(i + 1, a').$$

Finally, the algorithm returns $\max_{a \in A} \{M(1, a) + a\}$. Thus:

Lemma 10. *For any $\delta > 0$, there exists a dynamic programming algorithm that provides a δ -approximation (in the additive sense) to MAX-LINREV. Moreover, the algorithm runs in time polynomial in the size of the input and $\frac{1}{\delta}$.*

Now, we are ready to prove the main result of this section.

Theorem 4. *There exists an additive PTAS for computing a revenue-maximizing (ϕ, f) pair with private signaling.*

¹⁸Notice that, since MAX-LINREV takes as input distributions with a *finite support*, we can safely assume that such distributions can be explicitly represented in memory. In our PTAS, the inputs to the dynamic programming algorithm are obtained by building empirical distributions thorough samples from the actual distributions of buyers' valuations, thus ensuring finiteness of the supports.

¹⁹Notice that, given a pair (ξ, p) with $\xi \in \Xi^q$ and $p \in P^b$, it is possible to compute in polynomial time the probability with which a buyer $i \in \mathcal{N}$ buys the item.

REFERENCES

- [1] Marek Adamczyk, Allan Borodin, Diodato Ferraioli, Bart De Keijzer, and Stefano Leonardi. 2017. Sequential posted-price mechanisms with correlated valuations. *ACM Transactions on Economics and Computation (TEAC)* 5, 4 (2017), 1–39.
- [2] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. 2017. Posting prices with unknown distributions. *ACM Transactions on Economics and Computation (TEAC)* 5, 2 (2017), 1–20.
- [3] Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandr Slivkins. 2015. Dynamic pricing with limited supply. *ACM Transactions on Economics and Computation (TEAC)* 3, 1 (2015), 1–26.
- [4] Francesco Bacchiocchi, Matteo Castiglioni, Alberto Marchesi, Giulia Romano, and Nicola Gatti. 2022. Public Signaling in Bayesian Ad Auctions. *CoRR* abs/2201.09728 (2022). arXiv:2201.09728 <https://arxiv.org/abs/2201.09728>
- [5] 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* (New Orleans, Louisiana) (SODA '18). Society for Industrial and Applied Mathematics, USA, 2545–2563.
- [6] Matteo Castiglioni, Andrea Celli, and Nicola Gatti. 2020. Public Bayesian Persuasion: Being Almost Optimal and Almost Persuasive. *CoRR* abs/2002.05156 (2020). arXiv:2002.05156 <https://arxiv.org/abs/2002.05156>
- [7] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. 2020. Online Bayesian Persuasion. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/ba5451d3c91a0f982f103cde249bc78-Abstract.html>
- [8] Matteo Castiglioni, Diodato Ferraioli, Nicola Gatti, Alberto Marchesi, and Giulia Romano. 2022. Efficiency of Ad Auctions with Price Displaying. *CoRR* abs/2201.12275 (2022). arXiv:2201.12275 <https://arxiv.org/abs/2201.12275>
- [9] Matteo Castiglioni and Nicola Gatti. 2021. Persuading Voters in District-based Elections. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2–9, 2021*. AAAI Press, 5244–5251. <https://ojs.aaai.org/index.php/AAAI/article/view/16662>
- [10] 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 (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 1314–1323. <https://proceedings.mlr.press/v139/castiglioni21a.html>
- [11] Matteo Castiglioni, Alessandro Nuara, Giulia Romano, Giorgio Spadaro, Francesco Trovò, and Nicola Gatti. 2022. Safe Online Bid Optimization with Return-On-Investment and Budget Constraints subject to Uncertainty. *CoRR* abs/2201.07139 (2022). arXiv:2201.07139 <https://arxiv.org/abs/2201.07139>
- [12] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. 311–320.
- [13] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. 2015. Mixture Selection, Mechanism Design, and Signaling. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17–20 October, 2015*, Venkatesan Guruswami (Ed.). IEEE Computer Society, 1426–1445. <https://doi.org/10.1109/FOCS.2015.91>
- [14] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 169–186.
- [15] Liran Einav, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. 2018. Auctions versus posted prices in online markets. *Journal of Political Economy* 126, 1 (2018), 178–215.
- [16] eMarketer. 2021. Worldwide ecommerce will approach \$5 trillion this year. <https://www.emarketer.com/content/worldwide-ecommerce-will-approach-5-trillion-this-year>. Accessed: 2021-09-02.
- [17] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tenenholtz. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation* 2, 2 (2014), 1–19.
- [18] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [19] S. Khot and R. Saket. 2012. Hardness of Finding Independent Sets in Almost q -Colorable Graphs. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, Los Alamitos, CA, USA, 380–389. <https://doi.org/10.1109/FOCS.2012.75>
- [20] Robert Kleinberg and Tom Leighton. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. IEEE, 594–605.
- [21] Giulia Romano, Gianluca Tartaglia, Alberto Marchesi, and Nicola Gatti. 2021. Online Posted Pricing with Unknown Time-Discounted Valuations. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2–9, 2021*. AAAI Press, 5682–5689. <https://ojs.aaai.org/index.php/AAAI/article/view/16713>
- [22] Stefan Seifert. 2006. *Posted price offers in internet auction markets*. Vol. 580. Springer Science & Business Media.
- [23] Virag Shah, Ramesh Johari, and Jose Blanchet. 2019. Semi-Parametric Dynamic Contextual Pricing. In *Advances in Neural Information Processing Systems*. 2363–2373.
- [24] Tao Xiao, Zhengyang Liu, and Wenhan Huang. 2020. On the Complexity of Sequential Posted Pricing. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1521–1529.
- [25] Haifeng Xu. 2020. On the Tractability of Public Persuasion with No Externalities. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5–8, 2020*, Shuchi Chawla (Ed.). SIAM, 2708–2727. <https://doi.org/10.1137/1.9781611975994.165>