POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Public Signaling in Vickrey-Clarke-Groves Ad Auctions

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Francesco Bacchiocchi

Advisor: Prof. Nicola Gatti

Co-advisor: Matteo Castiglioni, Alberto Marchesi, Giulia Romano

Academic year: 2021-2022

---

1. Introduction

Nowadays, worldwide spending in digital advertising is skyrocketing, and this growth is primarily driven by ad auctions. These account for almost all market share, since they are at the core of popular advertising platforms, such as, *e.g.*, those by Google, Amazon, and Facebook. In a standard ad auction, the advertisers (also called bidders) compete for displaying their ads on a limited number of slots, and each bidder has their own private valuation representing how much they value a click on their ad. In this work, we study Bayesian ad auctions, which are characterized by the fact that bidders' valuations depend on a random, unknown state of nature. The auction mechanism has complete knowledge of the actual state of nature, and it can send signals to bidders so as to disclose information about the state and increase revenue. In particular, the auction mechanism commits to a signaling scheme, which is defined as a randomized mapping from states of nature to signals being sent to the bidders. Our model fits many real-world applications that are not captured by classical ad auctions. For instance, a state of nature may collectively encode some features of the user visualizing the ads—such as, *e.g.*, age, gender, or geographical region—that are known to the mechanism only, since the latter has access to data sources inaccessible to the bidders.

1.1. Original Contributions

We start with a negative result, showing that, in general, the problem does not admit a PTAS unless $\mathsf{P} = \mathsf{NP}$, even when bidders' valuations are known to the mechanism. The rest of the thesis is devoted to settings in which such negative result can be circumvented. First, we prove that, with known valuations, the problem can indeed be solved in polynomial time when either the number of states d or the number of slots m is fixed. Moreover, in the same setting, we provide an FPTAS for the case in which bidders are single minded, but d and m can be arbitrary. Then, we switch to the random valuations setting, in which these are randomly drawn according to some probability distribution. In this case, we show that the problem admits an FPTAS, a PTAS, and a QPTAS, when, respectively, d is fixed, m is fixed, and bidders' valuations are bounded away from zero.

1.2. Related Works

The algorithmic study of signaling in auctions is mainly focused on second-price auction, which can be seen as a special ad auction with a single slot. Emek et al. [2014] study second-price auctions in the known valuations setting. They provide an LP to compute an optimal public signaling scheme. Moreover, they show that it is NP-hard to compute an optimal signaling scheme in the random valuations setting. Cheng et al. [2015] complement the hardness result of Emek et al. [2014] by providing a PTAS for the random valuations setting. Finally, Badanidiyuru et al. [2018] study algorithms whose running time does not depend on the number of states of nature and initiate the study of private signaling schemes.

2. Model Formulation

In a standard ad auction, there is a set $\mathcal{N} := [n]$ of advertisers (or bidders) who compete for displaying their ads on a set $\mathcal{M} := [m]$ of slots, with $m \leq n$. Each bidder $i \in \mathcal{N}$ is characterized by a private valuation $v_i \in [0, 1]$, which represents how much they value a click on their ad. Moreover, each slot $j \in \mathcal{M}$ is associated with a click through rate parameter $\lambda_j \in [0, 1]$, which is the probability with which the slot is clicked by a user. W.l.o.g., we assume that the slots are ordered so that $\lambda_1 \geq \ldots \geq \lambda_m$. The auction goes on as follows: first, each bidder $i \in \mathcal{N}$ separately reports a bid $b_i \in [0, 1]$ to the auction mechanism; then, based on the bids, the latter allocates an ad to each slot and defines how much each bidder has to pay the mechanism for a click on their ad. We focus on truthful mechanisms, and the VCG mechanism in particular. In truthful mechanisms, allocation and payments are defined so that it is a dominant strategy for each bidder to report their true valuation to the mechanism, namely $b_i = v_i$ for every $i \in \mathcal{N}$. Moreover, the allocation implemented by the VCG mechanism orderly assigns the first m bidders in decreasing value of b_i to the first m slots (those with the highest click through rates). At the same time, assuming w.l.o.g. that bidder i is assigned to slot i, the mechanism defines an expected payment $p_i := \sum_{j=i+1}^{m+1} b_j(\lambda_{j-1} - \lambda_j)$ for each bidder $i \in [m]$, where, for the ease of notation, we let $\lambda_{m+1} = 0$. The payment is zero for all the other bidders. We study Bayesian ad auctions, which are characterized by a set $\Theta := \{\theta_1, \ldots, \theta_d\}$ of d states of nature. Each bidder $i \in \mathcal{N}$ has a valuation vector $v_i \in [0, 1]^d$, with $v_i(\theta)$ being bidder i's valuation in state $\theta \in \Theta$, and all such vectors are arranged in a matrix of bidders' valuations $V \in [0, 1]^{n \times d}$, whose entries are defined as $V(i, \theta) := v_i(\theta)$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$. We model signaling by means of the Bayesian persuasion framework. We consider the case in which the auction mechanism knows the state of nature and acts as a sender by issuing signals to the bidders (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 $\mu \in \Delta_\Theta$, with μ_θ denoting the probability of state $\theta \in \Theta$.[1] The mechanism publicly commits to a signaling scheme ϕ, which is a randomized mapping from states of nature to signals for the bidders. We focus on the case of public signaling in which all the bidders receive the same signal from the auction mechanism. Formally, a signaling scheme is a function $\phi : \Theta \to \Delta_\mathcal{S}$, where $\mathcal{S}$ is a set of available signals. For the ease of notation, we let $\phi_\theta(s)$ be the probability of sending signal $s \in \mathcal{S}$ when the state is $\theta \in \Theta$. A Bayesian ad auction goes on as follows: (i) the auction mechanism commits to a signaling scheme ϕ, and the bidders observe it; (ii) the mechanism gets to know the state of nature $\theta \sim \mu$ and draws signal $s \sim \phi(\theta)$; and (iv) the bidders observe the signal s and rationally update their prior belief over states according to Bayes rule. After observing signal $s \in \mathcal{S}$, all the bidders infer a posterior distribution $\xi_s \in \Delta_\Theta$ over states (also called posterior for short) such that the posterior probability of state $\theta \in \Theta$ is

$$\xi_s(\theta) := \frac{\mu_\theta \phi_\theta(s)}{\sum_{\theta' \in \Theta} \mu_{\theta'} \phi_{\theta'}(s)}. \tag{1}$$

Finally, each bidder $i \in \mathcal{N}$ truthfully reports to the mechanism their expected valuation given the posterior $\xi_s \in \Delta_\Theta$, namely $\xi_s^\top v_i = \sum_{\theta \in \Theta} v_i(\theta)\, \xi_s(\theta)$, and the mechanism allocates slots and defines payments as in a standard ad auction. Moreover, it is oftentimes useful to represent signaling schemes as convex combinations of the posteriors they can induce. Formally, a signaling scheme $\phi : \Theta \to \Delta_\mathcal{S}$ induces a probability distribution γ over posteriors in Δ_Θ, with $\gamma(\xi)$ denoting the probability of posterior

[1] Given a finite set X, we denote with Δ_X the simplex defined over the elements of X.

$\xi \in \Delta_\Theta$, defined as follow:

$$\gamma(\xi) \coloneqq \sum_{s \in \mathcal{S}: \xi_s = \xi} \sum_{\theta \in \Theta} \mu_\theta \phi_\theta(s).$$

Indeed, we can directly reason about distributions γ over Δ_Θ rather than about signaling schemes, provided that they are consistent with the prior as follows:

$$\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)\, \xi(\theta) = \mu_\theta \quad \forall \theta \in \Theta. \tag{2}$$

From now on we will use the term signaling scheme to refer to a consistent distribution γ over Δ_Θ. We focus on the problem of computing an optimal signaling scheme, *i.e.*, one maximizing the revenue of the mechanism. We study two settings: the **known valuations** (KV) setting in which the matrix of bidders' valuations V is known to the mechanism and the **random valuations** (RV) setting in which the matrix of bidders' valuations V is unknown, but randomly drawn according to a probability distribution $\mathcal{V}$. We denote by $\textsc{Rev}(V, \xi)$ the expected revenue of the mechanism when the bidders' valuations are given by V and the posterior induced by the mechanism is $\xi \in \Delta_\Theta$. Formally, given that bidders truthfully report their expected valuations and assuming w.l.o.g. that bidder i is assigned by the mechanism to slot i, we can write $\textsc{Rev}(V, \xi) \coloneqq \sum_{j=1}^m j\, \xi^\top v_{j+1}(\lambda_j - \lambda_{j+1})$. Then, given a signaling scheme γ, the expected revenue of the mechanism is $\mathbb{E}_{\xi \sim \gamma}[\textsc{Rev}(V, \xi)]$. When the valuations are unknown, we let $\textsc{Rev}(\mathcal{V}, \xi) \coloneqq \mathbb{E}_{V \sim \mathcal{V}}[\mathbb{E}_{\xi \sim \gamma}[\textsc{Rev}(V, \xi)]]$ and define the expected revenue analogously. Notice that, given a distribution of valuations $\mathcal{V}$ (or, in the KV setting, a matrix of bidders' valuations V) and a finite set $\Xi \subseteq \Delta_\Theta$ of posteriors, it is possible to formulate the problem of computing an optimal signaling scheme as an LP, as follows:

$$\max_{\gamma \in \Delta_\Xi} \sum_{\xi \in \Xi} \gamma(\xi)\, \textsc{Rev}(\mathcal{V}, \xi) \quad \text{s.t.} \tag{3a}$$

$$\sum_{\xi \in \Xi} \gamma(\xi)\, \xi(\theta) = \mu_\theta \qquad \forall \theta \in \Theta. \tag{3b}$$

Note that LP 3 is written for the RV setting, its analogous for the KV setting can be obtained by substituting $\textsc{Rev}(\mathcal{V}, \xi)$ with $\textsc{Rev}(V, \xi)$. In the following, we let OPT_Ξ be the optimal value of LP 3, while we denote with OPT the optimal expected revenue of the mechanism over all the possible signaling schemes γ.

3. A General Inapproximability Result

We start our analysis with the following negative result:

Theorem 3.1. *The problem of computing an optimal signaling scheme does* not *admit a* PTAS *unless* $\mathsf{P} = \mathsf{NP}$, *even when it is restricted to the KV setting.*

Theorem 3.1 is proved by a reduction from the VERTEX COVER problem in cubic graphs.

4. KV Setting: Parametrized Complexity

In this section, we study the parametrized complexity of the problem of computing an optimal revenue-maximizing signaling scheme, showing that it admits a polynomial-time algorithm when either the number of slots m or the number of states of nature d is fixed. In the following, we let $\Pi_l \subseteq 2^{\mathcal{N}}$ be the set of all the the possible permutations of $l \leq n$ bidders taken from $\mathcal{N}$, with $\pi = (i_1, ..., i_l) \in \Pi_l$ denoting a tuple made by bidders $i_1, \ldots, i_l \in \mathcal{N}$, in that order. We also let $\Xi_\pi \subseteq \Delta_\Theta$ be the (possibly empty) polytope of posteriors in which the expected valuations of bidders in $\pi \in \Pi_l$ are ordered (from the highest to the lowest) according to π; formally, it holds $\Xi_\pi \coloneqq \left\{ \xi \in \Delta_\Theta \mid \xi^\top v_{i_1} \geq \xi^\top v_{i_2} \geq \ldots \geq \xi^\top v_{i_l} \right\}$. Notice that, for any fixed $\pi \in \Pi_l$ with $l \geq m+1$, the term $\textsc{Rev}(V, \xi)$ is linear in ξ over Ξ_π.

4.1. Fixing the Number of Slots

In this case, the problem can be solved in polynomial time by formulating it as an LP, thanks to the following lemma:

Lemma 4.1. *There always exists an optimal signaling scheme γ such that $|\Xi_\pi \cap \text{supp}(\gamma)| \leq 1$ for every $\pi \in \Pi_{m+1}$.*

By Lemma 4.1, we can re-write the problem of computing a revenue maximizing signaling scheme as $\max \sum_{\pi \in \Pi_{m+1}} \gamma(\xi_\pi) \textsc{Rev}(V, \xi_\pi)$ subject to constraints ensuring that each ξ_π belongs to Ξ_π and that γ is a consistent probability distribution over such posteriors (see Equation 3b). This problem can be formulated as an LP by introducing a variable for each $\pi \in \Pi_{m+1}$ and $\theta \in \Theta$, encoding the products $\gamma(\xi_\pi)\xi_\pi(\theta)$ that define the expected revenue. Overall, the resulting LP has a number of variables and constraints that is $O(n^m)$, which, after fixing m, is polynomial in the size of the input. Thus, we conclude that:

Theorem 4.1. *In the KV setting, if the number of slots m is fixed, then an optimal signaling scheme can be computed in polynomial time.*

4.2. Fixing the Number of States

Our polynomial-time algorithm exploits the fact that an optimal signaling scheme can be computed by restricting the attention to distributions supported on a finite set of posteriors whose cardinality is polynomial in all the parameters, except from d. In particular, it is sufficient to focus on the set $\Xi^* := \bigcup_{\pi \in \Pi_n} V(\Xi_\pi)$, where $V(\cdot)$ denotes the set of vertices of the polytope given as input. Formally:

Lemma 4.2. *It holds that* $\text{OPT}_{\Xi^*} = \text{OPT}$.

Moreover, since it is possible to show that $|\Xi^*| = O((n^2+d)^{d-1})$, an optimal signaling scheme can be computed by means of LP 3 instantiated for the set Ξ^*, which has a number of variables and constraints that is polynomial once d is fixed. This proves the following:

Theorem 4.2. *In the KV setting, if the number of states d is fixed, then an optimal signaling scheme can be computed in polynomial time.*

5. KV Setting: Single-Minded Bidders

In this section, we focus on particular Bayesian ad auctions where the bidders are single minded. Intuitively, in our setting, by single mindedness we mean that each bidder is interested in displaying their ad only when the realized state of nature is a specific (single) state, and that all the bidders interested in the same state value a click on their ad for the same amount. We introduce the following formal definition:

Definition 5.1 (Single-minded bidders)**.** *In a Bayesian ad auction, we say that bidders are* single minded *if there exist $\mathcal{N}_\theta \subseteq \mathcal{N}$ and $\delta_\theta \in [0,1]$ for all $\theta \in \Theta$ such that:*

(i) $\mathcal{N} = \bigcup_{\theta \in \Theta} \mathcal{N}_\theta$ and $\mathcal{N}_\theta \cap \mathcal{N}_{\theta'} = \emptyset$ for all $\theta \neq \theta' \in \Theta$;

(ii) for every $\theta \in \Theta$ and $i \in \mathcal{N}_\theta$, it holds $v_i(\theta) = \delta_\theta$ and $v_i(\theta') = 0$ for all $\theta' \in \Theta : \theta' \neq \theta$.

Note that, since bidders truthfully report their expected valuations, the mechanism will always receive at most d different bids, one per set $\mathcal{N}_\theta$. In the following, we let $\Pi \subseteq 2^\Theta$ be the set of all the permutations of the sates of nature $\Theta = \{\theta_i\}_{i=1}^d$, while we let $\pi = (\theta_{k_1}, \dots, \theta_{k_d}) \in \Pi$ be an ordered tuple made by states $\theta_{k_1}, \dots, \theta_{k_d} \in \Theta$, where $k_1, \dots, k_d \in [d]$. Moreover we define: $\Xi_\pi := \left\{ \xi \in \Delta_\Theta \mid \delta_{\theta_{k_1}} \xi(\theta_{k_1}) \geq \ldots \geq \delta_{\theta_{k_d}} \xi(\theta_{k_d}) \right\}$ as the polytope of posteriors in which the expected valuations are ordered according to π. The first preliminary result that we need in order to derive our approximation algorithm is a characterization of the vertices of the sets Ξ_π for $\pi \in \Pi$, as follows.

Lemma 5.1. *Given $\pi \in \Pi$ and $\xi \in \Xi_\pi$, it holds that $\xi \in V(\Xi_\pi)$ if and only if there exists $\ell \in [d]$ such that:*

(i) $\delta_{\theta_{k_1}} \xi(\theta_{k_1}) = \ldots = \delta_{\theta_{k_\ell}} \xi(\theta_{k_\ell}) > 0$; and

(ii) $\delta_{\theta_{k_{\ell+1}}} \xi(\theta_{k_{\ell+1}}) = \ldots = \delta_{\theta_{k_d}} \xi(\theta_{k_d}) = 0$.

By letting $\Xi^* = \bigcup_{\pi \in \Pi} V(\Xi_\pi)$, since the term $\textsc{Rev}(V, \xi)$ is linear in ξ over Ξ_π for every permutation $\pi \in \Pi$, we can conclude that $\text{OPT}_{\Xi^*} = \text{OPT}$. Thus, Lemma 5.1 allows us to find an optimal signaling scheme by solving LP 3 for the set Ξ^* and the matrix of bidders' valuations V. However, notice that, since the size of Ξ^* is exponential in d, the resulting LP has exponentially-many variables. Nevertheless, since the LP has polynomially-many constraints, we can still solve it in polynomial time by applying the ellipsoid algorithm to its dual, provided that a polynomial-time separation oracle is available. In order to design a polynomial-time separation oracle, we apply the procedure described above to a relaxed version of LP 3, whose optimal value is sufficiently "close" to that of the original LP. In order to do that we have to design a polynomial-time separation oracle which reads as follow:

Definition 5.2 (Separation problem)**.** *Given values for the dual variables $y_\theta \in [-\beta, 0]$ for all $\theta \in \Theta$, compute:*

$$\max_{\xi \in \Xi^*} \textsc{Rev}(V, \xi) - \sum_{\theta \in \Theta} y_\theta \, \xi(\theta). \tag{4}$$

The following Lemma 5.2 shows that Problem 4 can be solved optimally up to any given additive loss $\lambda > 0$, by means of a dynamic programming algorithm that runs in time polynomial in the size of the input, in $\frac{1}{\lambda}$, and in β. Formally:

Lemma 5.2. *Given $\lambda > 0$, there exists an algorithm that finds an additive λ-approximation to*

Problem 4, in time polynomial in the size of the input, in $\frac{1}{\lambda}$, and in β.

Since the algorithm in Lemma 5.2 only returns an approximate solution to Problem 4, we need to carefully apply the ellipsoid algorithm to solve a relaxed version of the dual of LP 3, so that it correctly works even with an approximated oracle. Some non-trivial duality arguments allow us to prove that, indeed, this can be achieved by only incurring in a small additive loss on the quality of the returned solution, and without degrading the running time of the algorithm. Overall, this allows us to conclude that:

Theorem 5.1. *In the KV setting, if the bidders are single minded, then the problem of computing an optimal signaling scheme admits an (additive) FPTAS.*

6. RV Setting

In this setting, as stated in Section 2, we assume that the auction mechanism has access to the distribution of bidders' valuations $\mathcal{V}$ only through a black-box sampling oracle. In the following, given $s \in N_{>0}$ *i.i.d* samples of matrices of bidders' valuations, namely $V_1, \ldots, V_s \in [0,1]^{n\times d}$, we let $\mathcal{V}^s$ be their empirical distribution, which is such that:

$$\Pr_{V\sim\mathcal{V}^s}\left\{V=\hat{V}\right\} := \frac{\sum_{t=1}^{s}\mathbb{1}\{V_t=\hat{V}\}}{s}$$

for all $\hat{V} \in [0,1]^{n\times d}$. In this section, we first study the parametrized complexity of the problem of computing an optimal signaling scheme in general auctions (Section 6.1), and, then, we address special auction settings in which the bidders' valuations are bounded away from zero, namely $v_i(\theta) > \delta$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$, for some threshold $\delta > 0$. In the latter case, we show that the problem admits a QPTAS and the result is tight (Section 6.2). Before stating our main results (Theorems 6.1, 6.2, 6.3, and 6.4), we introduce some preliminary useful lemmas. The first one (Lemma 6.1) works under the true distribution of bidders' valuations $\mathcal{V}$, and it establishes a connection between the optimal expected revenue (OPT) and the optimal value of LP 3 for suitably-defined finite sets $\Xi \subseteq \Delta_\Theta$ of posteriors (OPT_Ξ). In particular, we look at sets $\Xi \subseteq \Delta_\Theta$ for which the function $\textsc{Rev}(\mathcal{V},\cdot)$ is "*stable*" according to the following definition:

Definition 6.1 ((α,ε)-stability)**.** *Given $\alpha,\varepsilon \geq 0$ and a finite set $\Xi \subseteq \Delta_\Theta$, we say that $\textsc{Rev}(\mathcal{V},\cdot)$ is* (α,ε)-stable *for Ξ if, for every $\xi \in \Delta_\Theta$, there exists a distribution $\gamma_\xi \in \Delta_\Xi$ such that:*

$$E_{\gamma_\xi}[\textsc{Rev}(\mathcal{V},\xi')] \geq (1-\alpha)\textsc{Rev}(\mathcal{V},\xi) - \varepsilon. \quad (5)$$

For any finite set $\Xi \subseteq \Delta_\Theta$ such that $\textsc{Rev}(\mathcal{V},\cdot)$ is (α,ε)-*stable* for Ξ, starting from an optimal signaling scheme γ one can recover an optimal solution to LP 3, only incurring in "small" multiplicative and additive losses in the expected revenue, respectively of $1-\alpha$ and ε, formally we have:

Lemma 6.1. *Given $\alpha,\varepsilon \geq 0$ and $\Xi \subseteq \Delta_\Theta$ such that $\textsc{Rev}(\mathcal{V},\cdot)$ is (α,ε)-stable for Ξ, it holds* $\text{OPT}_\Xi \geq (1-\alpha)\text{OPT} - \varepsilon$.

The second lemma (Lemma 6.2) deals with the approximation error introduced by using an empirical distribution of bidders' valuations $\mathcal{V}^s$, rather than the actual distribution $\mathcal{V}$. Given a finite set $\Xi \subseteq \Delta_\Theta$ of posteriors, let $\gamma_{\mathcal{V}^s} \in \Delta_\Xi$ be an optimal solution to LP 3 for distribution $\mathcal{V}^s$ and set Ξ. Moreover, let $\text{OPT}_{\Xi,s} := E\left[\sum_{\xi\in\Xi}\gamma_{\mathcal{V}^s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\right]$ be the average expected revenue of signaling schemes $\gamma_{\mathcal{V}^s}$ under the true distribution of valuations $\mathcal{V}$, where the expectation is with respect to the sampling procedure that determines $\mathcal{V}^s$. Then, a concentration argument proves the following:

Lemma 6.2. *Given $\rho,\tau > 0$, let $\Xi \subseteq \Delta_\Theta$ be finite and $s := \left\lceil \frac{2(\lambda_1 m)^2}{\tau^2}\log\frac{2}{\rho}\right\rceil$, $\text{OPT}_{\Xi,s} \geq (1-\rho|\Xi|)\,\text{OPT}_\Xi - \tau$.*

Finally, the last lemma (Lemma 6.3) exploits Lemma 6.1 to provide two useful bounds on the value of OPT_{Ξ_q}, where $\Xi_q \subseteq \Delta_\Theta$ (for a given $q \in N_{>0}$) is the finite set of all the q-uniform posteriors, according to the following definition:

Definition 6.2 (q-uniform posterior)**.** *Given $q \in N_{>0}$, a posterior $\xi \in \Delta_\Theta$ is* q-uniform *if each $\xi(\theta)$ is a multiple of $\frac{1}{q}$.*

We first observe that the set Ξ_q has size $|\Xi_q| = \binom{q+d-1}{d-1} \leq \min\{d^q, q^d\}$. The two points in the following lemma are readily proved by applying Lemma 6.1, after noticing that the sets Ξ_q in the statement are such that the function $\textsc{Rev}(\mathcal{V},\cdot)$ is (α,ε)-stable for them, with suitable values of $\alpha \geq 0$ and $\varepsilon \geq 0$. Formally:

Lemma 6.3. *Given $q := \left\lceil \frac{1}{2\eta^2}\log\frac{m+1}{\eta}\right\rceil$, $\eta > 0$ it holds:*

(i) $\text{OPT}_{\Xi_q} \geq \text{OPT} - 2\eta m$;

(ii) if, for some $\delta > 0$, it is the case that $v_i(\theta) > \delta$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$, then $\text{OPT}_{\Xi_q} \geq (1 - \frac{\eta}{\delta})^2\, \text{OPT}$.

6.1. Parametrized Complexity

First, we study the computational complexity of the problem of computing an optimal signaling scheme when the number of states d is fixed. We provide an (additive) FPTAS that works by performing the following two steps: (i) it collects a suitable number $s \in N_{>0}$ of matrices of bidders' valuations, by invoking the sampling oracle; and (ii) it solves LP 3 for the resulting empirical distribution $\mathcal{V}^s$ and a suitably-defined set of q-uniform posteriors. In particular, given a desired (additive) error $\lambda > 0$, the algorithm works on the set Ξ_q for $q = \lceil \frac{md}{\lambda} \rceil$ and its approximation guarantees rely on the following Lemma 6.4, proved again by means of Lemma 6.1.

Lemma 6.4. *Given $\lambda > 0$ and $q = \lceil \frac{md}{\lambda} \rceil$, then* $\text{OPT}_{\Xi_q} \geq \text{OPT} - \lambda$.

Thanks to Lemmas 6.2 and 6.4 (the former applied for suitable values $\rho, \tau > 0$), we can prove that the procedure described in steps (i) and (ii) above gives a signaling scheme achieving an expected revenue at most a function of λ lower than OPT, provided that the number of samples s is defined as in Lemma 6.4. Moreover, let us notice that, since $|\Xi_q| = O(q^d) = O((\frac{1}{\lambda}md)^d)$, if d is fixed, then the overall procedure runs in time polynomial in the input size and in $\frac{1}{\lambda}$. Thus, we can conclude that:

Theorem 6.1. *In the RV setting, if the number of states d is fixed, then the problem of computing an optimal signaling scheme admits and (additive) FPTAS.*

Next, we switch the attention to the case in which the number of slots m is fixed. We provide an (additive) PTAS that works as the FPTAS in Theorem 6.1, but whose approximation guarantees follow from Lemma 6.2 and point (i) in Lemma 6.3 (rather than Lemma 6.4). Thus, the only difference with respect to the previous case is that the algorithm works on the set Ξ_q of q-uniform posteriors for q defined as in Lemma 6.3. As a result, since $|\Xi_q| = O(d^q)$ and q depends on a parameter $\eta > 0$ that is related to the quality of the obtained approximation, the algorithm is only a PTAS rather than an FPTAS. Formally, we can prove the following:

Theorem 6.2. *In the RV setting, if the number of slots m is fixed, then the problem of computing an optimal signaling scheme admits and (additive) PTAS.*

6.2. Valuations Bounded Away From Zero

We conclude the section by studying the case in which the bidders' valuations are bounded away from zero. This case is dealt with an algorithm identical to the one in Theorem 6.2, but carrying on the approximation analysis by using Lemma 6.2 and point (ii) in Lemma 6.3 (rater than point (i)). Thus, since the value of q in Lemma 6.3 is related to the quality of the approximation thorough a parameter $\eta > 0$ and also depends logarithmically on the number of slots m, we obtain:

Theorem 6.3. *In the RV setting, if $v_i(\theta) \geq \delta$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$ for some $\delta > 0$, then the problem of computing an optimal signaling scheme admits a (multiplicative) QPTAS.*

The following theorem shows that the result is tight.

Theorem 6.4. *Assuming the ETH, there exists a constant $\omega > 0$ such that finding a signaling scheme that provides an expected revenue at least of $(1-\omega)$OPT requires $I^{\tilde{\Omega}(\log I)}$ time, where I is the size of the problem instance. This holds even when $v_i(\theta) > \frac{1}{3}$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$.*[2]

References

[1] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation (TEAC)*, 2014.

[2] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE, 2015.

[3] Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2018.

[2] The $\tilde{\Omega}$ notation hides poly-logarithmic factors.