Non-parametric Revenue Optimization for Generalized Second Price Auctions

Mehryar Mohri
Courant Institute and Google Research,
251 Mercer Street, New York, NY

Andrés Muñoz Medina
Courant Institute of Mathematical Sciences,
251 Mercer Street, New York, NY

Abstract

We present an extensive analysis of the key problem of learning optimal reserve prices for generalized second price auctions. We describe two algorithms for this task: one based on density estimation, and a novel algorithm benefiting from solid theoretical guarantees and with a very favorable running-time complexity of $O(nS\log(nS))$, where n is the sample size and S the number of slots. Our theoretical guarantees are more favorable than those previously presented in the literature. Additionally, we show that even if bidders do not play at an equilibrium, our second algorithm is still well defined and minimizes a quantity of interest. To our knowledge, this is the first attempt to apply learning algorithms to the problem of reserve price optimization in GSP auctions. Finally, we present the first convergence analysis of empirical equilibrium bidding functions to the unique symmetric Bayesian-Nash equilibrium of a GSP.

1 INTRODUCTION

The Generalized Second-Price (GSP) auction is currently the standard mechanism used for selling sponsored search advertisement. As suggested by the name, this mechanism generalizes the standard second-price auction of Vickrey (1961) to multiple items. In the case of sponsored search advertisement, these items correspond to ad slots which have been ranked by their position. Given this ranking, the GSP auction works as follows: first, each advertiser places a bid; next, the seller, based on the bids placed, assigns a score to each bidder. The highest scored advertiser is assigned to the slot in the best position, that is, the one with the highest likelihood of being clicked on. The second highest score obtains the second best item and so on, until all slots have been allocated or all advertisers have been assigned to a slot. As with second-price auctions, the bidder's payment is independent of his bid. Instead, it depends solely on the bid of the advertiser assigned to the position below.

In spite of its similarity with second-price auctions, the GSP auction is not an incentive-compatible mechanism, that is, bidders have an incentive to lie about their valuations. This is in stark contrast with second-price auctions where truth revealing is in fact a dominant strategy. It is for this reason that predicting the behavior of bidders in a GSP auction is challenging. This is further worsened by the fact that these auctions are repeated multiple times a day. The study of all possible equilibria of this repeated game is at the very least difficult. While incentive compatible generalizations of the second-price auction exist, namely the Vickrey-Clark-Gloves (VCG) mechanism, the simplicity of the payment rule for GSP auctions as well as the large revenue generated by them has made the adoption of VCG mechanisms unlikely.

Since its introduction by Google, GSP auctions have generated billions of dollars across different online advertisement companies. It is therefore not surprising that it has become a topic of great interest for diverse fields such as Economics, Algorithmic Game Theory and more recently Machine Learning.

The first analysis of GSP auctions was carried out independently by Edelman et al. (2005) and Varian (2007). Both publications considered a full information scenario, that is one where the advertisers' valuations are publicly known. This assumption is weakly supported by the fact that repeated interactions allow advertisers to infer their adversaries' valuations. Varian (2007) studied the so-called Symmetric Nash Equilibria (SNE) which is a subset of the Nash equilibria with several favorable properties. In particular, Varian showed that any SNE induces an efficient allocation, that is an allocation where the highest positions are assigned to advertisers with high values. Furthermore, the revenue earned by the seller when advertisers play an SNE is always at least as much as the one obtained by VCG. The authors also presented some empirical results showing that some bidders indeed play by using an SNE. However, no theoretical justification can be given for the choice of this subset of equilibria (Börgers et al., 2013; Edelman and Schwarz, 2010). A finer analysis of the full information scenario was given by Lucier et al. (2012). The authors proved that, excluding the payment of the highest bidder, the revenue achieved at any Nash equilibrium is at least one half that of the VCG auction.

Since the assumption of full information can be unrealistic, a more modern line of research has instead considered a Bayesian scenario for this auction. In a Bayesian setting, it is assumed that advertisers' valuations are i.i.d. samples drawn from a common distribution. Gomes and Sweeney (2014) characterized all symmetric Bayes-Nash equilibria and showed that any symmetric equilibrium must be efficient. This work was later extended by Sun et al. (2014) to account for the quality score of each advertiser. The main contribution of this work was the design of an algorithm for the crucial problem of revenue optimization for the GSP auction. Lahaie and Pennock (2007) studied different *squashing* ranking rules for advertisers commonly used in practice and showed that none of these rules are necessarily optimal in equilibrium. This work is complemented by the simulation analysis of Vorobeychik (2009) who quantified the distance from equilibrium of bidding truthfully. Lucier et al. (2012) showed that the GSP auction with an optimal reserve price achieves at least 1/6 of the optimal revenue (of any auction) in a Bayesian equilibrium. More recently, Thompson and Leyton-Brown (2013) compared different allocation rules and showed that an *anchoring* allocation rule is optimal when valuations are sampled i.i.d. from a uniform distribution. With the exception of Sun et al. (2014), none of these authors have proposed an algorithm for revenue optimization using historical data.

Zhu et al. (2009) introduced a ranking algorithm to learn an optimal allocation rule. The proposed ranking is a convex combination of a quality score based on the features of the advertisement as well as a revenue score which depends on the value of the bids. This work was later extended in (He et al., 2014) where, in addition to the ranking function, a behavioral model of the advertisers is learned by the authors.

The rest of this paper is organized as follows. In Section 2, we give a learning formulation of the problem of selecting reserve prices in a GSP auction. In Section 3, we discuss previous work related to this problem. Next, we present and analyze two learning algorithms for this problem in Section 4, one based on density estimation extending to this setting an algorithm of Guerre et al. (2000), and a novel discriminative algorithm taking into account the loss function and benefiting from favorable learning guarantees. Section 5 provides a convergence analysis of the empirical equilibrium bidding function to the true equilibrium bidding function in a GSP. On its own, this result is of great interest as it justifies the common assumption of buyers playing a symmetric Bayes-Nash equilibrium. Finally, in Section 6, we report the results of experiments comparing our algorithms and demonstrating in particular the benefits of the second algorithm.

2 MODEL

For the most part, we will use the model defined by Sun et al. (2014) for GSP auctions with incomplete information. We consider N bidders competing for S slots with $N \geq S$. Let $v_i \in [0,1]$ and $b_i \in [0,1]$ denote the per-click valuation of bidder i and his bid respectively. Let the position factor $c_s \in [0,1]$ represent the probability of a user noticing an ad in position s and let $e_i \in [0,1]$ denote the expected click-through rate of advertiser i. That is e_i is the probability of ad i being clicked on given that it was noticed by the user. We will adopt the common assumption that $c_s > c_{s+1}$ (Gomes and Sweeney, 2014; Lahaie and Pennock, 2007; Sun et al., 2014; Thompson and Leyton-Brown, 2013). Define the score of bidder i to be $s_i = e_i v_i$. Following Sun et al. (2014), we assume that s_i is an i.i.d. realization of a random variable with distribution F and density function f. Finally, we assume that advertisers bid in an efficient symmetric Bayes-Nash equilibrium. This is motivated by the fact that even though advertisers may not infer what the valuation of their adversaries is from repeated interactions, they can certainly estimate the distribution F.

Define $\pi \colon s \mapsto \pi(s)$ as the function mapping slots to advertisers, i.e. $\pi(s) = i$ if advertiser i is allocated to position s. For a vector $\mathbf{x} = (x_1, \ldots, x_N) \in \mathbb{R}^N$, we use the notation $x^{(s)} := x_{\pi(s)}$. Finally, denote by r_i the reserve price for advertiser i. An advertiser may participate in the auction only if $b_i \geq r_i$. In this paper we present an analysis of the two most common ranking rules (Qin et al., 2014):

1. Rank-by-bid. Advertisers who bid above their reserve price are ranked in descending order of their bids and the payment of advertiser $\pi(s)$ is equal to $\max(r^{(s)}, b^{(s+1)})$.
2. Rank-by-revenue. Each advertiser is assigned a quality score $q_i := q_i(b_i) = e_i b_i \mathbb{1}_{b_i \geq r_i}$ and the ranking is done by sorting these scores in descending order. The payment of advertiser $\pi(s)$ is given by $\max\left(r^{(s)}, \frac{q^{(s+1)}}{e^{(s)}}\right)$.

In both setups, only advertisers bidding above their reserve price are considered. Notice that rank-by-bid is a particular case of rank-by-revenue where all quality scores are equal to 1. Given a vector of reserve prices $\mathbf{r}$ and a bid vector $\mathbf{b}$, we define the revenue function to be

$$\mathrm{Rev}(\mathbf{r}, \mathbf{b}) = \sum_{s=1}^{S} c_s \left(\frac{q^{(s+1)}}{e^{(s)}} \mathbb{1}_{q^{(s+1)} \geq e^{(s)} r^{(s)}} + r^{(s)} \mathbb{1}_{q^{(s+1)} < e^{(s)} r^{(s)} \leq q^{(s)}} \right)$$

Using the notation of Mohri and Medina (2014), we define the loss function

$$L(\mathbf{r}, \mathbf{b}) = -\mathrm{Rev}(\mathbf{r}, \mathbf{b}).$$

Given an i.i.d. sample $\mathcal{S} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of realizations of an auction, our objective will be to find a reserve price vector $\mathbf{r}^*$ that maximizes the expected revenue. Equivalently, $\mathbf{r}^*$ should be a solution of the following optimization problem:

$$\min_{\mathbf{r} \in [0,1]^N} \mathbb{E}_{\mathbf{b}}[L(\mathbf{r}, \mathbf{b})]. \tag{1}$$

3 PREVIOUS WORK

It has been shown, both theoretically and empirically, that reserve prices can increase the revenue of an auction (Myerson, 1981; Ostrovsky and Schwarz, 2011). The choice of an appropriate reserve price therefore becomes crucial. If it is chosen too low, the seller might lose some revenue. On the other hand, if it is set too high, then the advertisers may not wish to bid above that value and the seller will not obtain any revenue from the auction.

Mohri and Medina (2014), Pardoe et al. (2005), and Cesa-Bianchi et al. (2013) have given learning algorithms that estimate the optimal reserve price for a second-price auction in different information scenarios. The scenario we consider is most closely related to that of Mohri and Medina (2014). An extension of this work to the GSP auction, however, is not straightforward. Indeed, as we will show later, the optimal reserve price vector depends on the distribution of the advertisers' valuation. In a second-price auction, these valuations are observed since the corresponding mechanism is an incentive-compatible. This does not hold for GSP auctions. Moreover, for second-price auctions, only one reserve price had to be estimated. In contrast, our model requires the estimation of up to N parameters with intricate dependencies between them.

The problem of estimating valuations from observed bids in a non-incentive compatible mechanism has been previously analyzed. Guerre et al. (2000) described a way of estimating valuations from observed bids in a first-price auction. We will show that this method can be extended to the GSP auction. The rate of convergence of this algorithm, however, in general will be worse than the standard learning rate of $O\big(\frac{1}{\sqrt{n}}\big)$.

Sun et al. (2014) showed that, for advertisers playing an efficient equilibrium, the optimal reserve price is given by $r_i = \frac{\overline{r}}{e_i}$ where $\overline{r}$ satisfies

$$\overline{r} = \frac{1 - F(\overline{r})}{f(\overline{r})}.$$

The authors suggest learning $\overline{r}$ via a maximum likelihood technique over some parametric family to estimate f and F, and to use these estimates in the above expression. There are two main drawbacks for this algorithm. The first is a standard problem of parametric statistics: there are no guarantees on the convergence of their estimation procedure when the density function f is not part of the parametric family considered. While this problem can be addressed by the use of a non-parametric estimation algorithm such as kernel density estimation, the fact remains that the function f is the density for the unobservable scores s_i and therefore cannot be properly estimated. The solution proposed by the authors assumes that the bids in fact form a perfect SNE and so advertisers' valuations can be recovered using the process described by Varian (2007). There is however no justification for this assumption and, in fact, we show in Section 6 that bids played in a Bayes-Nash equilibrium do not in general form a SNE.

4 LEARNING ALGORITHMS

Here, we present and analyze two algorithms for learning the optimal reserve price for a GSP auction when advertisers play a symmetric equilibrium.

4.1 DENSITY ESTIMATION ALGORITHM

First, we derive an extension of the algorithm of Guerre et al. (2000) to GSP auctions. To do so, we first derive a formula for the bidding strategy at equilibrium. Let $z_s(v)$ denote the probability of winning position s given that the advertiser's valuation is v. It is not hard to verify that

$$z_s(v) = \binom{N-1}{s-1}(1 - F(v))^{s-1} F^p(v),$$

where $p = N - s$. Indeed, in an efficient equilibrium, the bidder with the s-th highest valuation must be assigned to the s-th highest position. Therefore an advertiser with valuation v is assigned to position s if and only if $s-1$ bidders have a higher valuation and p have a lower valuation.

For a rank-by-bid auction, Gomes and Sweeney (2014) showed the following results.

Theorem 1 (Gomes and Sweeney (2014))**.** *A GSP auction has a unique efficient symmetric Bayes-Nash equilibrium with bidding strategy β if and only if β is strictly increasing and satisfies the following integral equation:*

$$\sum_{s=1}^{S} c_s \int_0^v \frac{dz_s(t)}{dt} t dt \tag{2}$$
$$= \sum_{s=1}^{S} c_s \binom{N-1}{s-1} (1 - F(v))^{s-1} \int_0^v \beta(t) p F^{p-1}(t) f(t) dt.$$

Furthermore, the optimal reserve price r^* satisfies

$$r^* = \frac{1 - F(r^*)}{f(r^*)}. \tag{3}$$

The authors show that, if the click probabilities c_s are sufficiently diverse, then, β is guaranteed to be strictly increasing. When ranking is done by revenue, Sun et al. (2014) gave the following theorem.

Theorem 2 (Sun et al. (2014))**.** *Let β be defined by the previous theorem. If advertisers bid in a Bayes-Nash equilibrium then $b_i = \frac{\beta(v_i)}{e_i}$. Moreover, the optimal reserve price vector $\mathbf{r}^*$ is given by $r_i^* = \frac{\overline{r}}{e_i}$ where $\overline{r}$ satisfies equation (3).*

We are now able to present the foundation of our first algorithm. Instead of assuming that the bids constitute an SNE as in (Sun et al., 2014), we follow the ideas of Guerre et al. (2000) and infer the scores s_i only from observables b_i. Our result is presented for the rank-by-bid GSP auction but an extension to the rank-by-revenue mechanism is trivial.

Lemma 1. *Let $v_1, \ldots, v_n$ be an i.i.d. sample of valuations from distribution F and let $b_i = \beta(v_i)$ be the bid played at equilibrium. Then the random variables b_i are i.i.d. with distribution $G(b) = F(\beta^{-1}(b))$ and density $g(b) = \frac{f(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}$. Furthermore,*

$$\begin{aligned} v_i &= \beta^{-1}(b_i) \qquad (4)\\ &= \frac{\sum_{s=1}^S c_s \binom{N-1}{s-1}(1-G(b_i))^{s-1} b_i p G(b_i)^{p-1} g(b_i)}{\sum_{s=1}^S c_s \binom{N-1}{s-1} \frac{d\overline{z}}{db}(b_i)} \\ &- \frac{\sum_{s=1}^S c_s (s-1)(1-G(b_i))^{s-2} g(b_i) \int_0^{b_i} pG(u)^{p-1} u g(u) du}{\sum_{s=1}^S c_s \binom{N-1}{s-1} \frac{d\overline{z}}{db}(b_i)}, \end{aligned}$$

where $\overline{z}_s(b) := z_s(\beta^{-1}(b))$ and is given by $\binom{N-1}{s-1}(1 - G(b))^{s-1} G(b)^{p-1}$.

Proof. By definition, $b_i = \beta(v_i)$ is a function of only v_i. Since β does not depend on the other samples either, it follows that $(b_i)_{i=1}^N$ must be an i.i.d. sample. Using the fact that β is a strictly increasing function we also have $G(b) = P(b_i \leq b) = P(v_i \leq \beta^{-1}(b)) = F(\beta^{-1}(b))$ and a simple application of the chain rule gives us the expression for the density $g(b)$. To prove the second statement observe that by the change of variable $v = \beta^{-1}(b)$, the right-hand side of (2) is equal to

$$\begin{aligned} &\sum_{s=1}^S \binom{N-1}{s-1} (1-G(b))^{s-1} \int_0^{\beta^{-1}(b)} p\beta(t) F^{p-1}(t) f(t) dt \\ =&\sum_{s=1}^S \binom{N-1}{s-1} (1-G(b))^{s-1} \int_0^{b} puG(u)^{p-1}(u) g(u) du. \end{aligned}$$

The last equality follows by the change of variable $t = \beta(u)$ and from the fact that $g(b) = \frac{f(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}$. The same change of variables applied to the left-hand side of (2) yields the following integral equation:

$$\begin{aligned} &\sum_{s=1}^S \binom{N-1}{s-1} \int_0^b \beta^{-1}(u) \frac{d\overline{z}}{du}(u) du \\ &= \sum_{s=1}^S \binom{N-1}{s-1} (1-G(b))^{s-1} \int_0^b upG(u)^{p-1}(u) g(u) du. \end{aligned}$$

Taking the derivative with respect to b of both sides of this equation and rearranging terms lead to the desired expression. □

The previous Lemma shows that we can recover the valuation of an advertiser from its bid. We therefore propose the following algorithm for estimating the value of $\overline{r}$.

1. Use the sample $\mathcal{S}$ to estimate G and g.
2. Plug this estimates in (4) to obtain approximate samples from the distribution F.
3. Use the approximate samples to find estimates $\widehat{f}$ and $\widehat{F}$ of the valuations density and cumulative distribution functions respectively.
4. Use $\widehat{F}$ and $\widehat{f}$ to estimate $\overline{r}$.

In order to avoid the use of parametric methods, a kernel density estimation algorithm can be used to estimate g and f. While this algorithm addresses both drawbacks of the algorithm proposed by Sun et al. (2014), it can be shown (Guerre et al., 2000)[Theorem 2] that if f is R times continuously differentiable, then, after seeing n samples, $\|f - \widehat{f}\|_\infty$ is in $\Omega\big(\frac{1}{n^{R/(2R+3)}}\big)$ independently of the algorithm used to estimate f. In particular, note that for $R = 1$ the rate is in $\Omega\big(\frac{1}{n^{1/4}}\big)$. This unfavorable rate of convergence can be attributed to the fact that a two-step estimation algorithm is being used (estimation of g and f). But, even with access to bidder valuations, the rate can only be improved to $\Omega\big(\frac{1}{n^{R/(2R+1)}}\big)$ (Guerre et al., 2000). Furthermore, a small error in the estimation of f affects the denominator of the equation defining $\overline{r}$ and can result in a large error on the estimate of $\overline{r}$.

4.2 DISCRIMINATIVE ALGORITHM

In view of the problems associated with density estimation, we propose to use empirical risk minimization to find an approximation to the optimal reserve price. In particular, we are interested in solving the following optimization problem:

$$\min_{\mathbf{r} \in [0,1]^N} \sum_{i=1}^n L(\mathbf{r}, \mathbf{b}_i). \qquad (5)$$

We first show that, when bidders play in equilibrium, the optimization problem (1) can be considerably simplified.

Proposition 1. *If advertisers play a symmetric Bayes-Nash equilibrium then*

$$\min_{\mathbf{r} \in [0,1]^N} \mathbb{E}_{\mathbf{b}}[L(\mathbf{r}, \mathbf{b})] = \min_{r \in [0,1]} \mathbb{E}_{\mathbf{b}}[\widetilde{L}(r, \mathbf{b})],$$

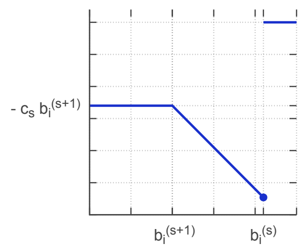

Figure 1: Plot of the loss function $L_{i,s}$. Notice that the loss in fact resembles a broken "V" .

where $\widetilde{q}_i := \widetilde{q}_i(b_i) = e_i b_i$ *and*

$$\widetilde{L}(r, \mathbf{b}) = -\sum_{s=1}^{S} \frac{c_s}{e^{(s)}} \Big(\widetilde{q}^{(s+1)} \mathbb{1}_{\widetilde{q}^{(s+1)} \geq r} + r \mathbb{1}_{\widetilde{q}^{(s+1)} < r \leq \widetilde{q}^{(s)}} \Big).$$

Proof. Since advertisers play a symmetric Bayes-Nash equilibrium, the optimal reserve price vector $\mathbf{r}^*$ is of the form $r_i^* = \frac{\overline{r}}{e_i}$. Therefore, letting $D = \{\mathbf{r} | r_i = \frac{r}{e_i},\ r \in [0,1]\}$ we have $\min_{\mathbf{r} \in [0,1]^N} \mathbb{E}_{\mathbf{b}}[L(\mathbf{r}, \mathbf{b})] = \min_{\mathbf{r} \in D} \mathbb{E}_{\mathbf{b}}[L(\mathbf{r}, \mathbf{b})]$. Furthermore, when restricted to D, the objective function L is given by

$$-\sum_{s=1}^{S} \frac{c_s}{e^{(s)}} \Big(q^{(s+1)} \mathbb{1}_{q^{(s+1)} \geq r} + r \mathbb{1}_{q^{(s+1)} < r \leq q^{(s)}} \Big).$$

Thus, we are left with showing that replacing $q^{(s)}$ with $\widetilde{q}^{(s)}$ in this expression does not affect its value. Let $r \geq 0$, since $q_i = \widetilde{q}_i \mathbb{1}_{\widetilde{q}_i \geq r}$, in general the equality $q^{(s)} = \widetilde{q}^{(s)}$ does not hold. Nevertheless, if s_0 denotes the largest index less than or equal to S satisfying $q^{(s_0)} > 0$, then $\widetilde{q}^{(s)} \geq r$ for all $s \leq s_0$ and $q^{(s)} = \widetilde{q}^{(s)}$. On the other hand, for $S \geq s > s_0$, $\mathbb{1}_{q^{(s)} \geq r} = \mathbb{1}_{\widetilde{q}^{(s)} \geq r} = 0$. Thus,

$$\begin{aligned}
&\sum_{s=1}^{S} \frac{c_s}{e^{(s)}} \Big(q^{(s+1)} \mathbb{1}_{q^{(s+1)} \geq r} + r \mathbb{1}_{q^{(s+1)} < r \leq q^{(s)}} \Big) \\
&= \sum_{s=1}^{s_0} \frac{c_s}{e^{(s)}} \Big(q^{(s+1)} \mathbb{1}_{q^{(s+1)} \geq r} + r \mathbb{1}_{q^{(s+1)} < r \leq q^{(s)}} \Big) \\
&= \sum_{s=1}^{s_0} \frac{c_s}{e^{(s)}} \Big(\widetilde{q}^{(s+1)} \mathbb{1}_{\widetilde{q}^{(s+1)} \geq r} + r \mathbb{1}_{\widetilde{q}^{(s+1)} < r \leq \widetilde{q}^{(s)}} \Big) \\
&= -\widetilde{L}(r, \mathbf{b}),
\end{aligned}$$

which completes the proof. □

In view of this proposition, we can replace the challenging problem of solving an optimization problem in $\mathbb{R}^N$ with solving the following simpler empirical risk minimization problem

$$\min_{r \in [0,1]} \sum_{i=1}^{n} \widetilde{L}(r, \mathbf{b}_i) = \min_{r \in [0,1]} \sum_{i=1}^{n} \sum_{s=1}^{S} L_{s,i}(r, \widetilde{q}^{(s)}, \widetilde{q}^{(s+1)}), \tag{6}$$

Algorithm 1 Minimization algorithm

Require: Scores $(\widetilde{q}_i^{(s)})$, $1 \leq n$, $1 \leq s \leq S$.

Define $(p_{is}^{(1)}, p_{is}^{(2)}) = (\widetilde{q}_i^{(s)}, \widetilde{q}_i^{(s+1)})$; $m = nS$;
$\mathcal{N} := \bigcup_{i=1}^{n} \bigcup_{s=1}^{S} \{p_{is}^{(1)}, p_{is}^{(2)}\}$;
$(n_1, ..., n_{2m}) = \textbf{Sort}(\mathcal{N})$;
Set $\mathbf{d}_i := (d_1, d_2) = \mathbf{0}$
Set $d_1 = -\sum_{i=1}^{n} \sum_{s=1}^{S} \frac{c_s}{e_i} p_{is}^{(2)}$;
Set $r^* = -1$ and $L^* = \infty$
for $j = 2, \ldots, 2m$ **do**
  if $n_{j-1} = p_{is}^{(2)}$ **then**
    $d_1 = d_1 + \frac{c_s}{e_i} p_{is}^{(2)}$; $d_2 = d_2 - \frac{c_s}{e_i}$;
  else if $n_{j-1} = p_{is}^{(1)}$ **then**
    $d_2 = d_2 + \frac{c_s}{e_s}$
  end if
  $L = d_1 - n_j d_2$;
  if $L < L^*$ **then**
    $L^* = L$; $r^* = n_j$;
  end if
end for
return r^*;

where $L_{s,i}(r, \widetilde{q}^{(s)}), \widetilde{q}^{(s+1)}) := -\frac{c_s}{e^{(s)}} (\widetilde{q}_i^{(s+1)} \mathbb{1}_{\widetilde{q}_i^{(s+1)} \geq r} - r \mathbb{1}_{\widetilde{q}_i^{(s+1)} < r \leq \widetilde{q}_i^{(s)}})$. In order to efficiently minimize this highly non-convex function, we draw upon the work of Mohri and Medina (2014) on minimization of sums of v-functions.

Definition 1. *A function* $V : \mathbb{R}^3 \to \mathbb{R}$ *is a* v-function *if it admits the following form:*

$$\begin{aligned}
&V(r, q_1, q_2) \\
&= -a^{(1)} \mathbb{1}_{r \leq q_2} - a^{(2)} r \mathbb{1}_{q_2 < r \leq q_1} + \Big[\frac{r}{\eta} - a^{(3)}\Big] \mathbb{1}_{q_1 < r < (1+\eta) q_1},
\end{aligned}$$

with $0 \leq a^{(1)}, a^{(2)}, a^{(3)}, \eta \leq \infty$ *constants satisfying* $a^{(1)} = a^{(2)} q_2$, $-a^{(2)} q_1 \mathbb{1}_{\eta > 0} = \big(\frac{1}{\eta} q_1 - a^{(3)}\big) \mathbb{1}_{\eta > 0}$. *Under the convention that* $0 \cdot \infty = 0$.

As suggested by their name, these functions admit a characteristic "V shape". It is clear from Figure 1 that $L_{s,i}$ is a v-function with $a^{(1)} = \frac{c_s}{e^{(s)}} \widetilde{q}_i^{(s+1)}$, $a^{(2)} = \frac{c_s}{e^{(s)}}$ and $\eta = 0$. Thus, we can apply the optimization algorithm given by Mohri and Medina (2014) to minimize (6) in $O(nS \log nS)$ time. Algorithm 1 gives the pseudocode of that the adaptation of this general algorithm to our problem. A proof of the correctness of this algorithm can be found in (Mohri and Medina, 2014).

We conclude this section by presenting learning guarantees for our algorithm. Our bounds are given in terms of the Rademacher complexity and the VC-dimension.

Definition 2. *Let* $\mathcal{X}$ *be a set and let* $G := \{g : \mathcal{X} \to \mathbb{R}\}$ *be a family of functions. Given a sample* $\mathcal{S} = (x_1, \ldots, x_n) \in$

$\mathcal{X}$, *the empirical Rademacher complexity of* G *is defined by*

$$\widehat{\mathfrak{R}}_S(G) = \frac{1}{n}\,\mathbb{E}_\sigma\Big[\sup_{g\in G}\frac{1}{n}\sum_{i=1}^n \sigma_i g(x_i)\Big],$$

where σ_i*s are independent random variables distributed uniformly over the set* $\{-1,1\}$.

Proposition 2. *Let* $\mathfrak{m} = \min_i e_i > 0$ *and* $\mathfrak{M} = \sum_{s=1}^S c_s$. *Then, for any* $\delta > 0$, *with probability at least* $1-\delta$ *over the draw of a sample* $\mathcal{S}$ *of size* n, *each of the following inequalities holds for all* $r \in [0,1]$*:*

$$\mathbb{E}[\widetilde{L}(r,\mathbf{b})] \le \frac{1}{n}\sum_{i=1}^n \widetilde{L}(r,\mathbf{b}_i) + C(\mathfrak{M},\mathfrak{m},n,\delta) \quad (7)$$

$$\frac{1}{n}\sum_{i=1}^n \widetilde{L}(r,\mathbf{b}_i) \le \mathbb{E}[\widetilde{L}(r,\mathbf{b})] + C(\mathfrak{M},\mathfrak{m},n,\delta), \quad (8)$$

where $C(\mathfrak{M},\mathfrak{m},n,\delta) = \frac{1}{\sqrt{n}} + \sqrt{\frac{\log(en)}{n}} + \sqrt{\frac{\mathfrak{M}\log(1/\delta)}{2\mathfrak{m}n}}$.

Proof. Let $\Psi\colon S \mapsto \sup_{r\in[0,1]}\frac{1}{n}\sum_{i=1}^n \widetilde{L}(r,\mathbf{b}_i) - \mathbb{E}[\widetilde{L}(r,\mathbf{b})]$. Let $\mathcal{S}^i$ be a sample obtained from $\mathcal{S}$ by replacing $\mathbf{b}_i$ with $\mathbf{b}_i'$. It is not hard to verify that $|\Psi(\mathcal{S}) - \Psi(\mathcal{S}^i)| \le \frac{\mathfrak{M}}{n\mathfrak{m}}$. Thus, it follows from a standard learning bound that, with probability at least $1-\delta$,

$$\mathbb{E}[\widetilde{L}(r,\mathbf{b})] \le \frac{1}{n}\sum_{i=1}^n \widetilde{L}(r,\mathbf{b}_i) + \widehat{\mathfrak{R}}_S(\mathcal{R}) + \sqrt{\frac{\mathfrak{M}\log(1/\delta)}{2\mathfrak{m}n}},$$

where $\mathcal{R} = \{\overline{L}_r : \mathbf{b} \mapsto \widetilde{L}(r,\mathbf{b}) | r \in [0,1]\}$. We proceed to bound the empirical Rademacher complexity of the class $\mathcal{R}$. For $q_1 > q_2 \ge 0$ let $\overline{L}(r,q_1,q_2) = q_2 \mathbb{1}_{q_2>r} + r\mathbb{1}_{q_1\ge r\ge q_2}$. By definition of the Rademacher complexity we can write

$$\begin{aligned}
\widehat{\mathfrak{R}}_S(\mathcal{R}) &= \frac{1}{n}\,\mathbb{E}_\sigma\Big[\sup_{r\in[0,1]}\sum_{i=1}^n \sigma_i \overline{L}_r(\mathbf{b}_i)\Big] \\
&= \frac{1}{n}\,\mathbb{E}_\sigma\Big[\sup_{r\in[0,1]}\sum_{i=1}^n \sigma_i \sum_{s=1}^S \frac{c_s}{e_s}\overline{L}(r,\widetilde{q}_i^{(s)},\widetilde{q}_i^{(s+1)})\Big] \\
&\le \frac{1}{n}\,\mathbb{E}_\sigma\Big[\sum_{s=1}^S \sup_{r\in[0,1]}\sum_{i=1}^n \sigma_i \psi_s(\overline{L}(r,\widetilde{q}_i^{(s)},\widetilde{q}_i^{(s+1)}))\Big],
\end{aligned}$$

where ψ_s is the $\frac{c_s}{\mathfrak{m}}$-Lipschitz function mapping $x \mapsto \frac{c_s}{e^{(s)}}x$. Therefore, by Talagrand's contraction lemma (Ledoux and Talagrand, 2011), the last term is bounded by

$$\sum_{s=1}^S \frac{c_s}{n\mathfrak{m}}\,\mathbb{E}_\sigma \sup_{r\in[0,1]}\sum_{i=1}^n \sigma_i \overline{L}(r,\widetilde{q}_i^{(s)},\widetilde{q}_i^{(s+1)}) = \sum_{s=1}^S \frac{c_s}{\mathfrak{m}}\widehat{\mathfrak{R}}_{\mathcal{S}_s}(\widetilde{\mathcal{R}}),$$

where $\mathcal{S}_s = \big((\widetilde{q}_1^{(s)},\widetilde{q}_1^{(s+1)}),\ldots,(\widetilde{q}_n^{(s)},\widetilde{q}_n^{(s+1)})\big)$ and $\widetilde{\mathcal{R}} := \{\overline{L}(r,\cdot,\cdot)|r\in[0,1]\}$. The loss $\overline{L}(r,\widetilde{q}^{(s)},\widetilde{q}^{(s+1)})$ in fact evaluates to the negative revenue of a second-price auction with highest bid $\widetilde{q}^{(s)}$ and second highest bid $\widetilde{q}^{(s+1)}$ (Mohri and Medina, 2014). Therefore, by Propositions 9 and 10 of Mohri and Medina (2014) we can write

$$\begin{aligned}
\widehat{\mathfrak{R}}_{\mathcal{S}_s}(\widetilde{\mathcal{R}}) &\le \frac{1}{n}\,\mathbb{E}_\sigma\Big[\sup_{r\in[0,1]}\sum_{i=1}^n r\sigma_i\Big] + \sqrt{\frac{2\log en}{n}} \\
&\le \Big(\frac{1}{\sqrt{n}} + \sqrt{\frac{2\log en}{n}}\Big),
\end{aligned}$$

which concludes the proof. ☐

Corollary 1. *Under the hypotheses of Proposition 2, let* $\widehat{r}$ *denote the empirical minimizer and* r^* *the minimizer of the expected loss. Then, for any* $\delta > 0$, *with probability at least* $1-\delta$, *the following inequality holds:*

$$\mathbb{E}[\widetilde{L}(\widehat{r},\mathbf{b})] - \mathbb{E}[\widetilde{L}(r^*,\mathbf{b})] \le 2C\Big(\mathfrak{M},\mathfrak{m},n,\frac{\delta}{2}\Big).$$

Proof. By the union bound, (7) and (8) hold simultaneously with probability at least $1-\delta$ if δ is replaced by $\delta/2$ in those expression. Adding both inequalities and using the fact that $\widehat{r}$ is an empirical minimizer yields the result. ☐

It is worth noting that our algorithm is well defined whether or not the buyers bid in equilibrium. Indeed, the algorithm consists of the minimization over r of an observable quantity. While we can guarantee convergence to a solution of (1) only when buyers play a symmetric BNE, our algorithm will still find an approximate solution to

$$\min_{r\in[0,1]} \mathbb{E}_{\mathbf{b}}[L(r,\mathbf{b})],$$

which remains a quantity of interest that can be close to (1) if buyers are close to the equilibrium.

5 CONVERGENCE OF EMPIRICAL EQUILIBRIA

A crucial assumption in the study of GSP auctions, including this work, is that advertisers bid in a Bayes-Nash equilibrium (Lucier et al., 2012; Sun et al., 2014). This assumption is partially justified by the fact that advertisers can infer the underlying distribution F using as observations the outcomes of the past repeated auctions and can thereby implement an efficient equilibrium.

In this section, we provide a stronger theoretical justification in support of this assumption: we quantify the difference between the bidding function calculated using observed empirical distributions and the true symmetric bidding function in equilibria. For the sake of notation simplicity, we will consider only the rank-by-bid GSP auction.

Let $\mathcal{S}_v = (v_1,\ldots,v_n)$ be an i.i.d. sample of values drawn from a continuous distribution F with density function f.

Assume without loss of generality that $v_1 \leq \ldots \leq v_n$ and let $\mathbf{v}$ denote the vector defined by $\mathbf{v}_i = v_i$. Let $\widehat{F}$ denote the empirical distribution function induced by $\mathcal{S}_v$ and let $\mathbf{F} \in \mathbb{R}^n$ and $\mathbf{G} \in \mathbb{R}^n$ be defined by $\mathbf{F}_i = \widehat{F}(v_i) = i/n$ and $\mathbf{G}_i = 1 - \mathbf{F}_i$.

We consider a *discrete* GSP auction where the advertiser's valuations are i.i.d. samples drawn from a distribution $\widehat{F}$. In the event where two or more advertisers admit the same valuation, ties are broken randomly. Denote by $\widehat{\beta}$ the bidding function for this auction in equilibrium (when it exists). We are interested in characterizing $\widehat{\beta}$ and in providing guarantees on the convergence of $\widehat{\beta}$ to β as the sample size increases.

We first introduce the notation used throughout this section.

Definition 3. *Given a vector $\mathbf{F} \in \mathbb{R}^n$, the backwards difference operator $\Delta : \mathbb{R}^n \to \mathbb{R}^n$ is defined as:*

$$\Delta \mathbf{F}_i = \mathbf{F}_i - \mathbf{F}_{i-1},$$

for $i > 1$ and $\Delta \mathbf{F}_1 = \mathbf{F}_1$.

We will denote $\Delta\Delta\mathbf{F}_i$ by $\Delta^2\mathbf{F}_i$. Given any $k \in \mathbb{N}$ and a vector $\mathbf{F}$, the vector $\mathbf{F}^k$ is defined as $\mathbf{F}^k_i = (\mathbf{F}_i)^k$. Let us now define the discrete analog of the function z_s that quantifies the probability of winning slot s.

Proposition 3. *In a symmetric efficient equilibrium of the discrete GSP, the probability $\widehat{z}_s(v)$ that an advertiser with valuation v is assigned to slot s is given by*

$$\widehat{z}_s(v) = \sum_{j=0}^{N-s} \sum_{k=0}^{s-1} \binom{N-1}{j, k, N-1-j-k} \frac{\mathbf{F}^j_{i-1}\mathbf{G}^k_i}{(N-j-k)n^{N-1-j-k}},$$

if $v = v_i$ and otherwise by

$$\widehat{z}_s(v) = \binom{N-1}{s-1} \lim_{v' \to v^-} \widehat{F}(v')^p (1 - \widehat{F}(v))^{s-1} =: \widehat{z}_s^-(v),$$

where $p = N - s$.

In particular, notice that $\widehat{z}_s^-(v_i)$ admits the simple expression

$$\widehat{z}_s^-(v_i) = \binom{N-1}{s-1} \mathbf{F}^p_{i-1}\mathbf{G}^{s-1}_{i-1},$$

which is the discrete version of the function z_s. On the other hand, even though $\widehat{z}_s(v_i)$ does not admit a closed-form, it is not hard to show that

$$\widehat{z}_s(v_i) = \binom{N-1}{s-1} \mathbf{F}^p_{i-1}\mathbf{G}^{s-1}_i + O\Big(\frac{1}{n}\Big). \tag{9}$$

Which again can be thought of as a discrete version of z_s. The proof of this and all other propositions in this section are deferred to the Appendix. Let us now define the lower triangular matrix $\mathbf{M}(s)$ by:

$$\mathbf{M}_{ij}(s) = -\binom{N-1}{s-1} \frac{n\Delta\mathbf{F}^p_j \Delta\mathbf{G}^s_i}{s},$$

for $i > j$ and

$$\mathbf{M}_{ii}(s) = \sum_{j=0}^{N-s-1} \sum_{k=0}^{s-1} \binom{N-1}{j, k, N-1-j-k} \frac{\mathbf{F}^j_{i-1}\mathbf{G}^k_i}{(N-j-k)n^{N-1-j-k}}.$$

Proposition 4. *If the discrete GSP auction admits a symmetric efficient equilibrium, then its bidding function $\widehat{\beta}$ satisfies $\widehat{\beta}(v_i) = \boldsymbol{\beta}_i$, where $\boldsymbol{\beta}$ is the solution of the following linear equation.*

$$\mathbf{M}\boldsymbol{\beta} = \mathbf{u}, \tag{10}$$

with $\mathbf{M} = \sum_{s=1}^S c_s \mathbf{M}(s)$ and $\mathbf{u}_i = \sum_{s=1}^S \Big(c_s z_s(v_i) v_i - \sum_{j=1}^i \widehat{z}_s^-(v_j) \Delta \mathbf{v}_j \Big)$.

To gain some insight about the relationship between $\widehat{\beta}$ and β, we compare equations (10) and (2). An integration by parts of the right-hand side of (2) and the change of variable $G(v) = 1 - F(v)$ show that β satisfies

$$\sum_{s=1}^S c_s v z_s(v) - \int_0^v \frac{dz_s(t)}{dt} t dt = \sum_{s=1}^S c_s \binom{N-1}{s-1} G(v)^{s-1} \int_0^v \beta(t) dF^p. \tag{11}$$

On the other hand, equation (10) implies that for all i

$$\mathbf{u}_i = \sum_{s=1}^S c_s \Big[\mathbf{M}_{ii}(s)\boldsymbol{\beta}_i - \binom{N-1}{s-1} \frac{n\Delta\mathbf{G}^s_i}{s} \sum_{j=1}^{i-1} \Delta\mathbf{F}^p_j \boldsymbol{\beta}_j \Big]. \tag{12}$$

Moreover, by Lemma 2 and Proposition 10 in the Appendix, the equalities $-\frac{n\Delta\mathbf{G}^s_i}{s} = \mathbf{G}^{s-1}_i + O\big(\frac{1}{n}\big)$ and

$$\mathbf{M}_{ii}(s) = \frac{1}{2n} \binom{N-1}{s-1} p \mathbf{F}^{p-1}_{i-1} \mathbf{G}^{s-1}_i + O\Big(\frac{1}{n^2}\Big),$$

hold. Thus, equation (12) resembles a numerical scheme for solving (11) where the integral on the right-hand side is approximated by the trapezoidal rule. Equation (11) is in fact a Volterra equation of the first kind with kernel

$$K(t, v) = \sum_{s=1}^S \binom{N-1}{s-1} G(v)^{s-1} p F^{p-1}(t).$$

Therefore, we could benefit from the extensive literature on the convergence analysis of numerical schemes for this type of equations (Baker, 1977; Kress et al., 1989; Linz, 1985). However, equations of the first kind are in general

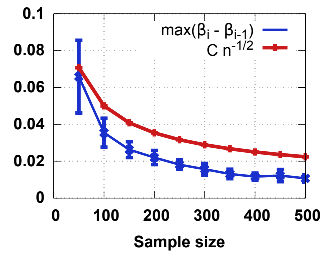

Figure 2: (a) Empirical verification of Assumption 2. The blue line corresponds to the quantity $\max_i \Delta\boldsymbol{\beta}_i$. In red we plot the desired upper bound for $C = 1/2$.

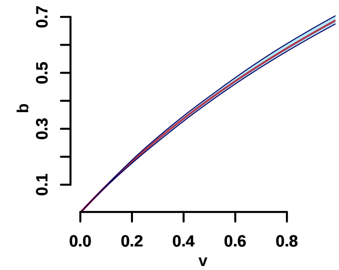

Figure 3: Approximation of the empirical bidding function $\widehat{\beta}$ to the true solution β. The true solution is shown in red and the shaded region represents the confidence interval of $\widehat{\beta}$ when simulating the discrete GSP 10 times with a sample of size 200. Where $N = 3$, $S = 2$, $c_1 = 1$, $c_2 = 0.5$ and bids were sampled uniformly from $[0, 1]$

ill-posed problems (Kress et al., 1989), that is small perturbations on the equation can produce large errors on the solution. When the kernel K satisfies $\min_{t\in[0,1]} K(t,t) > 0$, there exists a standard technique to transform an equation of the first kind to an equation of the second kind, which is a well posed problem. Thus, making the convergence analysis for these types of problems much simpler. The kernel function appearing in (11) does not satisfy this property and therefore these results are not applicable to our scenario. To the best of our knowledge, there exists no quadrature method for solving Volterra equations of the first kind with vanishing kernel.

In addition to dealing with an uncommon integral equation, we need to address the problem that the elements of (10) are not exact evaluations of the functions defining (11) but rather stochastic approximations of these functions. Finally, the grid points used for the numerical approximation are also random.

In order to prove convergence of the function $\widehat{\beta}$ to β we will make the following assumptions

Assumption 1. *There exists a constant $c > 0$ such that $f(x) > c$ for all $x \in [0, 1]$.*

This assumption is needed to ensure that the difference between consecutive samples $v_i - v_{i-1}$ goes to 0 as $n \to \infty$, which is a necessary condition for the convergence of any numerical scheme.

Assumption 2. *The solution $\boldsymbol{\beta}$ of (10) satisfies $v_i, \boldsymbol{\beta}_i \geq 0$ for all i and $\max_{i\in 1,\ldots,n} \Delta\boldsymbol{\beta}_i \leq \frac{C}{\sqrt{n}}$, for some universal constant C.*

Since $\boldsymbol{\beta}_i$ is a bidding strategy in equilibrium, it is reasonable to expect that $v_i \geq \boldsymbol{\beta}_i \geq 0$. On the other hand, the assumption on $\Delta\boldsymbol{\beta}_i$ is related to the smoothness of the solution. If the function β is smooth, we should expect the approximation $\widehat{\beta}$ to be smooth too. Both assumptions can in practice be verified empirically, Figure 2 depicts the quantity $\max_{i\in 1,\ldots,n} \Delta\boldsymbol{\beta}_i$ as a function of the sample size n.

Assumption 3. *The solution β to (2) is twice continuously differentiable.*

This is satisfied if for instance the distribution function F is twice continuously differentiable. We can now present our main result.

Theorem 3. *If Assumptions 1, 2 and 3 are satisfied, then, for any $\delta > 0$, with probability at least $1 - \delta$ over the draw of a sample of size n, the following bound holds for all $i \in [1, n]$:*

$$|\widehat{\beta}(v_i) - \beta(v_i)| \leq e^C \left[\frac{\log(\frac{2}{\delta})^{\frac{N}{2}}}{\sqrt{n}} q\Big(n, \frac{2}{\delta}\Big)^3 + \frac{Cq(n, \frac{2}{\delta})}{n^{3/2}} \right].$$

where $q(n, \delta) = \frac{2}{c} \log(nc/2\delta)$ with c defined in Assumption 1, and where C is a universal constant.

The proof of this theorem is highly technical, thus, we defer it to Appendix F.

6 EXPERIMENTS

Here we present preliminary experiments showing the advantages of our algorithm. We also present empirical evidence showing that the procedure proposed in Sun et al. (2014) to estimate valuations from bids is incorrect. In contrast, our density estimation algorithm correctly recovers valuations from bids in equilibrium.

6.1 SETUP

Let F_1 and F_2 denote the distributions of two truncated log-normal random variables with parameters $\mu_1 = \log(.5)$, $\sigma_1 = .8$ and $\mu_2 = \log(2)$, $\sigma = .1$; the mixture parameter was set to $1/2$. Here, F_1 is truncated to have support in $[0, 1.5]$ and the support of $F_2 = [0, 2.5]$. We consider a GSP with $N = 4$ advertisers with $S = 3$ slots and position factors $c_1 = 1$, $c_2 =, 45$ and $c_3 = 1$. Based on the results of Section 5 we estimate the bidding function β with a sample of 2000 points and we show its plot in Figure 4. We proceed to evaluate the method proposed by Sun et al. (2014) for

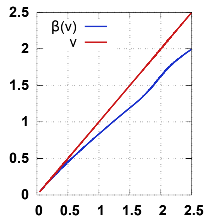

Figure 4: Bidding function for our experiments in blue and identity function in red.

recovering advertisers valuations from bids in equilibrium. The assumption made by the authors is that the advertisers play a SNE in which case valuations can be inferred by solving a simple system of inequalities defining the SNE (Varian, 2007). Since the authors do not specify which SNE the advertisers are playing we select the one that solves the SNE conditions with equality.

We generated a sample $\mathcal{S}$ consisting of $n = 300$ i.i.d. outcomes of our simulated auction. Since $N = 4$, the effective size of this sample is of 1200 points. We generated the outcome bid vectors $\mathbf{b}_i, \ldots, \mathbf{b}_n$ by using the equilibrium bidding function β. Assuming that the bids constitute a SNE we estimated the valuations and Figure 5 shows an histogram of the original sample as well as the histogram of the estimated valuations. It is clear from this figure that this procedure does not accurately recover the distribution of the valuations. By contrast, the histogram of the estimated valuations using our density estimation algorithm is shown in Figure 5(c). The kernel function used by our algorithm was a triangular kernel given by $K(u) = (1 - |u|)\mathbb{1}_{|u| \leq 1}$. Following the experimental setup of Guerre et al. (2000) the bandwidth h was set to $h = 1.06\widehat{\sigma}n^{1/5}$, where $\widehat{\sigma}$ denotes the standard deviation of the sample of bids.

Finally, we use both our density estimation algorithm and discriminative learning algorithm to infer the optimal value of r. To test our algorithm we generated a test sample of size $n = 500$ with the procedure previously described. The results are shown in Table 1.

Density estimation	Discriminative
1.42 ± 0.02	1.85 ± 0.02

Table 1: Mean revenue for our two algorithms.

7 CONCLUSION

We proposed and analyzed two algorithms for learning optimal reserve prices for generalized second price auctions. Our first algorithm is based on density estimation and therefore suffers from the standard problems associated with this family of algorithms. Furthermore, this algorithm is only well defined when bidders play in equilibrium. Our second algorithm is novel and is based on learning theory guarantees. We show that the algorithm admits an efficient $O(nS\log(nS))$ implementation. Furthermore, our theoretical guarantees are more favorable than those presented for the previous algorithm of Sun et al. (2014). Moreover, even though it is necessary for advertisers to play in equilibrium for our algorithm to converge to optimality, when bidders do not play an equilibrium, our algorithm is still well defined and minimizes a quantity of interest albeit over a smaller set. We also presented preliminary experimental results showing the advantages of our algorithm. To our knowledge, this is the first attempt to apply learning algorithms to the problem of reserve price selection in GSP auctions. We believe that the use of learning algorithms in revenue optimization is crucial and that this work may preface a rich research agenda including extensions of this work to a general learning setup where auctions and advertisers are represented by features. Additionally, in our analysis, we considered two different ranking rules. It would be interesting to combine the algorithm of Zhu et al. (2009) with this work to learn both a ranking rule and an optimal reserve price. Finally, we provided the first analysis of convergence of bidding functions in an empirical equilibrium to the true bidding function. This result on its own is of great importance as it justifies the common assumption of advertisers playing in a Bayes-Nash equilibrium.

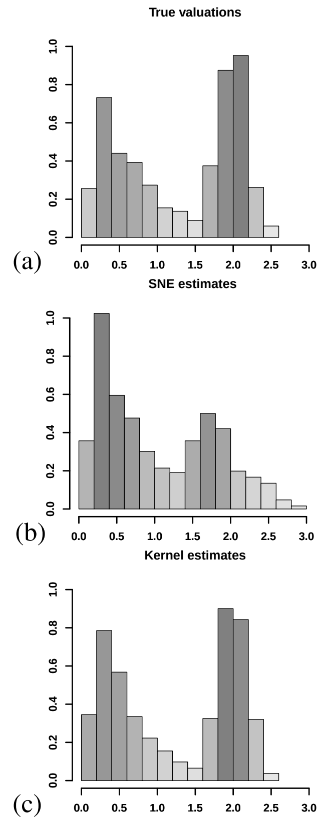

Figure 5: Comparison of methods for estimating valuations from bids. (a) Histogram of true valuations. (b) Valuations estimated under the SNE assumption. (c) Density estimation algorithm.

References

Baker, C. T. (1977). *The numerical treatment of integral equations*. Clarendon press.

Börgers, T., I. Cox, M. Pesendorfer, and V. Petricek (2013). Equilibrium bids in sponsored search auctions: Theory and evidence. *American Economic Journal: Microeconomics 5*(4), 163–87.

Cesa-Bianchi, N., C. Gentile, and Y. Mansour (2013). Regret minimization for reserve prices in second-price auctions. In *Proceedings of SODA 2013*, pp. 1190–1204.

Edelman, B., M. Ostrovsky, and M. Schwarz (2005). Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review 97*.

Edelman, B. and M. Schwarz (2010). Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review 100*(2), 597–602.

Gibbons, R. (1992). *Game theory for applied economists*. Princeton University Press.

Gomes, R. and K. S. Sweeney (2014). Bayes-Nash equilibria of the generalized second-price auction. *Games and Economic Behavior 86*, 421–437.

Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica 68*(3), 525–574.

He, D., W. Chen, L. Wang, and T. Liu (2014). A game-theoretic machine learning approach for revenue maximization in sponsored search. *CoRR abs/1406.0728*.

Kress, R., V. Maz'ya, and V. Kozlov (1989). *Linear integral equations*, Volume 82. Springer.

Lahaie, S. and D. M. Pennock (2007). Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of ACM EC*, pp. 50–56.

Ledoux, M. and M. Talagrand (2011). *Probability in Banach spaces*. Classics in Mathematics. Berlin: Springer-Verlag. Isoperimetry and processes, Reprint of the 1991 edition.

Linz, P. (1985). *Analytical and numerical methods for Volterra equations*, Volume 7. SIAM.

Lucier, B., R. P. Leme, and É. Tardos (2012). On revenue in the generalized second price auction. In *Proceedings of WWW*, pp. 361–370.

Milgrom, P. and I. Segal (2002). Envelope theorems for aribtrary choice sets. *Econometrica 70*(2), 583–601.

Mohri, M. and A. M. Medina (2014). Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of ICML*, pp. 262–270.

Myerson, R. (1981). Optimal auction design. *Mathematics of operations research 6*(1), 58–73.

Ostrovsky, M. and M. Schwarz (2011). Reserve prices in internet advertising auctions: a field experiment. In *Proceedings of ACM EC*, pp. 59–60.

Pardoe, D., P. Stone, M. Saar-Tsechansky, and K. Tomak (2005). Adaptive auctions: Learning to adjust to bidders. In *Proceedings of WITS 2005*.

Qin, T., W. Chen, and T. Liu (2014). Sponsored search auctions: Recent advances and future directions. *ACM TIST 5*(4), 60.

Sun, Y., Y. Zhou, and X. Deng (2014). Optimal reserve prices in weighted GSP auctions. *Electronic Commerce Research and Applications 13*(3), 178–187.

Thompson, D. R. M. and K. Leyton-Brown (2013). Revenue optimization in the generalized second-price auction. In *Proceedings of ACM EC*, pp. 837–852.

Varian, H. R. (2007, December). Position auctions. *International Journal of Industrial Organization 25*(6), 1163–1178.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance 16*(1), 8–37.

Vorobeychik, Y. (2009). Simulation-based analysis of keyword auctions. *SIGecom Exchanges 8*(1).

Zhu, Y., G. Wang, J. Yang, D. Wang, J. Yan, J. Hu, and Z. Chen (2009). Optimizing search engine revenue in sponsored search. In *Proceedings of ACM SIGIR*, pp. 588–595.