

Revenue Maximization for Finitely Repeated Ad Auctions

Jiang Rong,¹ Tao Qin,² Bo An,³ Tie-Yan Liu²

¹The Key Lab of Intelligent Information Processing, ICT, CAS
University of Chinese Academy of Sciences, Beijing 100190, China
rongjiang13@mails.ucas.ac.cn

²Microsoft Research, Beijing 100080, China
{taoqin, tyliu}@microsoft.com

³School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798
boan@ntu.edu.sg

Abstract

Reserve price is an effective tool for revenue maximization in ad auctions. The optimal reserve price depends on bidders' value distributions, which, however, are generally unknown to auctioneers. A common practice for auctioneers is to first collect information about the value distributions by a *sampling procedure* and then apply the reserve price estimated with the sampled bids to the following auctions. In order to maximize the total revenue over finite auctions, it is important for the auctioneer to find a proper sample size to trade off between the cost of the sampling procedure and the optimality of the estimated reserve price. We investigate the sample size optimization problem for Generalized Second Price auctions, which is the most widely-used mechanism in ad auctions, and make three main contributions along this line. First, we bound the revenue losses in the form of competitive ratio during and after sampling. Second, we formulate the problem of finding the optimal sample size as a non-convex mixed integer optimization problem. Then we characterize the properties of the problem and prove the uniqueness of the optimal sample size. Third, we relax the integer optimization problem to a continuous form and develop an efficient algorithm based on the properties to solve it. Experimental results show that our approach can significantly improve the revenue for the auctioneer in finitely repeated ad auctions.

1 Introduction

Ad auctions have become a major monetization channel for Internet applications, including sponsored search auctions (Jansen and Mullen 2008; Milgrom 2010; Qin, Chen, and Liu 2015) and realtime bidding (RTB) (Chakraborty et al. 2010; Chen et al. 2011; Yuan, Wang, and Zhao 2013). In sponsored search, when a user issues a query to a search engine, in addition to a list of relevant webpages, a selective set of ads related to the query will also be shown to the user. A position auction is used to determine which ads to show and how much to charge the corresponding advertisers. In RTB for display advertising, when a user visits a publisher's website, an ad impression with related information will be sent to the advertisers (or ad networks) through an ad exchange. Then the bids from the advertisers are collected and an auction is used to determine which ad to show and how

much to charge the advertiser. In these applications, Generalized Second Price¹ (GSP) is the most popularly used auction mechanism and has attracted a lot of research attention in recent years (Caragiannis et al. 2011; Chen et al. 2014; Ma et al. 2014; Rong et al. 2016).

In ad auctions, revenue maximization is a very important goal of the auctioneer and many research works have been conducted on this topic (Xiao, Yang, and Li 2009; Radovanovic and Heavlin 2012; Caragiannis et al. 2014; Yuan et al. 2014). The famous Myerson's theory (Myerson 1981; Nisan et al. 2007; Hartline and Roughgarden 2009) reveals how to achieve the optimal revenue in the setting of a single item for sale and independent and identically distributed (i.i.d.) values of the bidders: one just needs to adopt a second-price auction with an optimal reserve price that is the zero point of the *virtual value* function. This result has been extended to GSP. Edelman and Schwarz (2010) proved that GSP with the optimal reserve price is an optimal mechanism under the *regularity* assumption. Thompson and Leyton-Brown (2013) investigated how to optimally set the reserve price when squashing is considered. Sun, Zhou, and Deng (2014) showed that the optimal reserve price of weighted GSP depends on bidders' ad qualities.

All the above works on revenue maximization rely on the knowledge about the value distribution of bidders. However, bidders' valuations are private information and are invisible to the auctioneer. In practice, auctioneers usually need to collect information about the value distribution using the auction mechanism with a heuristically set reserve price. We use *sampling of bids* to refer to this process. With the estimated value distribution, the auctioneer can infer the optimal reserve price and use it in the future auctions. Ostrovsky and Schwarz (2011) conducted a field experiment on sampling of bids with Yahoo's datasets. They assumed the bidders' values to be drawn from a log normal distribution and inferred the parameters by simulation. Sun et al. (2012) estimated the parameters of the value distribution by maximum likelihood estimation. Cole and Roughgarden (2014) and Dhangwatnotai, Roughgarden, and Yan (2015) directly estimated the

¹This paper is based on unweighted GSP (Edelman, Ostrovsky, and Schwarz 2007; Edelman and Schwarz 2010; Ostrovsky and Schwarz 2011) where bidders are ranked and charged by bids. We discuss how to extend our results to weighted GSP (Varian 2007; Sun, Zhou, and Deng 2014) in Section 7.

optimal reserve price with sampled data and analyzed the revenue guarantee with the estimated reserve price in a very simple mechanism where all the winners pay the same price. Mohri and Medina (2015) proposed a discriminative algorithm for the estimation based on the assumption of symmetric Bayes-Nash equilibrium and derived an additive regret. While these works are good attempts on optimal reserve price estimation from sampled bids, they have neglected the revenue loss during the sampling period. Obviously, more rounds of sampling will lead to a more accurate estimation of the optimal reserve price. However, the sampling process usually cannot achieve the optimal revenue by itself, i.e., there is a “cost” of sampling. Hence, the auctioneer needs to determine how many rounds should be used for sampling in order to maximize the overall revenue. Bulow and Klemperer (1994) provided some basic ideas for the cost of sampling in the same environment as (Cole and Roughgarden 2014; Dhangwatnotai, Roughgarden, and Yan 2015) which cannot be applied to the GSP auctions directly.

We consider the trade-off between the cost of sampling and the optimality of the estimated reserve price for revenue maximization, in the context of finitely repeated GSP auctions. Bandit theory is a general framework for the exploration-exploitation problem, which has been used in post-price auctions (Kleinberg and Leighton 2003). However, the application of bandit algorithms to GSP is limited because 1) they need to update the reserve price in each round, which is impractical in sponsored search auctions since a query may be searched for tens of thousands of times one day and frequently updating the reserve price is very expensive, and 2) they all rely on specific assumptions on the function to optimize (Cesa-Bianchi, Gentile, and Mansour 2013), e.g., smoothness and strong concavity (Cope 2009), unimodality (Jia and Mannor 2011) and local Lipschitz condition (Bubeck et al. 2011). In GSP, it is difficult to satisfy these conditions because the revenue is not a continuous function with respect to the reserve price.

To solve the trade-off problem for GSP, we propose an algorithm containing two procedures in this work: the information collection (sampling) phase with zero reserve price and the deployment phase with estimated optimal reserve price. Our destination is to find an optimal sample size which is independent of the value distribution and robust for the auctioneer. We make three key contributions along this line. First, we give the competitive ratio of the revenue during the sampling period. We show that this ratio is distribution-independent and is only related to the number of bidders and ad slots. We assume that the (unknown) value distribution is regular but do not require any further information about its form (e.g., log normal, exponent, or normal). Besides, we use a method inspired by (Dhangwatnotai, Roughgarden, and Yan 2015) to estimate the optimal reserve price and prove its revenue guarantee in GSP. Second, we formulate the trade-off problem as a non-convex mixed integer optimization problem, which aims to maximize the auctioneer’s worst-case total revenue, and derive many properties for it, including the unimodal feature and uniqueness of optimal solution. Third, we propose an efficient algorithm (with exponential convergence rate) based on these properties. We evaluate

the performance of our proposed approach using numerical simulations. The experimental results show significant improvements in revenue for the auctioneer as compared with baseline strategies.

2 Preliminaries

We introduce the GSP mechanism and review basic ideas for optimal reserve price when the value distribution is known to the auctioneer in this section.

2.1 GSP Mechanism

There are N bidders competing for K ad slots and usually we assume that $K < N$. Let v_i denote the value of bidder i ’s ad for a click and b_i express the bid submitted by i to participate in the auction. We use the vector $v = (v_1, v_2, \dots, v_N) = (v_i, v_{-i})$ to represent the value profile. In the real world, the advertisers in each auction may be different (Cesa-Bianchi, Gentile, and Mansour 2013) and the value of the same ad to different users with various experience and identities is not fixed (Abrams and Schwarz 2007). These uncertainties can be captured by assuming that bidders’ values are i.i.d. from a distribution (Ostrovsky and Schwarz 2011; Lucier, Paes Leme, and Tardos 2012; Thompson and Leyton-Brown 2013). Furthermore, as assumed in most works (Sun et al. 2012; Cole and Roughgarden 2014; Dhangwatnotai, Roughgarden, and Yan 2015) that consider real applications, the distribution is unknown to the auctioneer. Each ad slot j has a corresponding click-through-rate (CTR) θ_j . We define $\theta_j \equiv 0$ for $j > K$ and then use a vector $\theta = (\theta_1, \theta_2, \dots, \theta_N)^T$ to denote the CTR profile, where T represents the matrix transposition. Generally, we assume $\theta_1 \geq \theta_2 \geq \dots \geq \theta_K > 0$.

The GSP sorts the bidders in descending order of their bids. The allocation rule, leaving out the reserve price, can be represented as a mapping $x : \mathbb{R}^N \mapsto \theta^N$. Specifically, given the bid profile $b = (b_1, b_2, \dots, b_N)$, $x(b) = (x_1(b), x_2(b), \dots, x_N(b))$ is a vector consisting of CTRs each bidder will receive, where $x_i(b) = \theta_k$ if and only if b_i is the k -th highest bid in b (ties are broken randomly). When considering the reserve price r , the allocation rule $x^r : \mathbb{R}^N \mapsto \theta^N$ is defined as $x_i^r(b) = x_i(b)$ if $b_i \geq r$; $x_i^r(b) = 0$ otherwise. Suppose bidders are labeled such that $b_i \geq b_{i+1}$ and define $b_{N+1} \equiv 0$, then bidder i should pay $\max\{b_{i+1}, r\}$ once his/her ad is clicked and the expected utility is $u_i(v_i, b) = v_i x_i^r(b) - p_i(b)$, where the second part is the expected payment and is written as $p_i(b) = x_i^r(b) \cdot \max\{b_{i+1}, r\}$. The expected revenue of the auctioneer is $\sum_{i=1}^N p_i(b)$.

2.2 Optimal Reserve Price with Known Value Distribution

We learn from the last subsection that the auctioneer’s expected revenue depends on the bid profile and reserve price. Here we present how to compute the optimal reserve price with known bidders’ value distribution. In the theory of revenue maximization, equilibrium solution concept is widely used to model bidders’ behavior (Edelman and Schwarz 2010; Ostrovsky and Schwarz 2011). Sun, Zhou, and Deng (2014)

proved that if all bidders adopt the same equilibrium strategy $b_i = \beta(v_i)$, then the expected revenue of the auctioneer with reserve price r , $R(r)$, in each auction can be written as

$$R(r) = E_v \left\{ \sum_{i=1}^N \psi(v_i) x_i^r(v) \right\}, \quad (1)$$

where $\psi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$ is called the virtual value function (Myerson 1981) with $f(\cdot)$ and $F(\cdot)$ denoting the probability density function and cumulative distribution function of bidders' value distribution respectively. Following the common practice like (Myerson 1981; Hartline 2006; Lucier, Paes Leme, and Tardos 2012), we assume that the virtual value function is regular, i.e., $\psi(\cdot)$ is strictly monotonically increasing on $(0, +\infty)$. Given the value distribution $F(\cdot)$, the revenue $R(r)$ can be maximized when $\sum_{i=1}^N \psi(v_i) x_i^r(v)$ is maximized pointwise for every v , which happens when (i) slots are only allocated to bidders with positive virtual values and (ii) bidders with higher virtual values are assigned higher CTRs. Since both $\psi(v_i)$ and $x_i^r(v)$ are non-decreasing functions of v_i , the second condition is satisfied. To meet the first condition, the auctioneer just needs to set an optimal reserve price r^* such that $\psi(r^*) = 0$, because $x_i^{r^*}(v) = 0$ when $v_i < r^*$, $\forall i \in \{1, 2, \dots, N\}$. We call the GSP mechanism with the optimal reserve price the *optimal GSP*. Since bidders' values are i.i.d., it follows that, $\forall i' \in \{1, 2, \dots, N\}$,

$$E_v \left\{ \sum_{i=1}^N \psi(v_i) x_i^r(v) \right\} = N \cdot E_v \left\{ \psi(v_{i'}) x_{i'}^r(v) \right\}. \quad (2)$$

3 Problem Formulation

Now we realistically consider that bidders' value distribution is unknown to the auctioneer who aims to maximize his/her overall revenue for M rounds. The first τ rounds are used for sampling and the reserve price is set as zero in this period in order to observe complete (i.e., non-truncated) value distribution. Note that the auctioneer can only observe bids of the advertisers but not their values. For ease of analysis, we assume that bidders were playing the widely-used lowest-revenue Symmetric Nash Equilibrium (SNE) (Varian 2007; Edelman and Schwarz 2010; Ostrovsky and Schwarz 2011; Sun, Zhou, and Deng 2014), which can easily recover values from bids ($v_i = \beta^{-1}(b_i)$) (Varian 2007). Then we make an estimation of the optimal reserve price, represented as \bar{r} , with the sampled data and set it for the remaining $M - \tau$ rounds. The total revenue the auctioneer will get is $\tau \cdot R(0) + (M - \tau) \cdot R(\bar{r})$. The revenue loss is thus $M \cdot R(r^*) - (\tau \cdot R(0) + (M - \tau) \cdot R(\bar{r}))$, which is equal to $M \cdot R(r^*) (1 - (\frac{\tau}{M} \cdot \frac{R(0)}{R(r^*)} + (1 - \frac{\tau}{M}) \cdot \frac{R(\bar{r})}{R(r^*)}))$. Then minimizing the total loss means that

$$\max_{\tau \in \{1, 2, \dots, M\}} \frac{\tau}{M} \frac{R(0)}{R(r^*)} + (1 - \frac{\tau}{M}) \frac{R(\bar{r})}{R(r^*)}. \quad (3)$$

Further analysis shows that the competitive ratios $\frac{R(0)}{R(r^*)}$ and $\frac{R(\bar{r})}{R(r^*)}$ are related to the density function $f(\cdot)$. As a result, the solution of the problem defined in Eq. (3) is a function of $f(\cdot)$. Our destination is to find an optimal sample size which

is independent of the priori $f(\cdot)$. Specifically, we prove in Section 4 that both $\frac{R(0)}{R(r^*)}$ and $\frac{R(\bar{r})}{R(r^*)}$ have lower bounds that do not rely on $f(\cdot)$. Based on this result, in Section 5, we reformulate the optimization problem with lower bounds of $\frac{R(0)}{R(r^*)}$ and $\frac{R(\bar{r})}{R(r^*)}$ and derive its optimal sample size maximizing the worst-case total revenue, which is thus robust to any distribution and safe for the auctioneer.

4 Competitive Ratio Derivation

To solve the trade-off problem defined in Eq. (3), we first derive the competitive ratios $\frac{R(0)}{R(r^*)}$ and $\frac{R(\bar{r})}{R(r^*)}$ in this section. Theorem 1 gives the ratio for the sampling period.

Theorem 1. $\frac{R(0)}{R(r^*)} \geq \frac{N-K}{N}$ for GSP with N bidders and K slots.

The theorem can be proved by combining the following two lemmas.

Lemma 2. *The expected revenue of GSP with N bidders, K slots and zero reserve price is at least that of an optimal GSP with $N - K$ bidders and K slots.*

Proof. The expected revenue of the optimal GSP with $N - K$ bidders is equal to $E_{v^{N-K}} \left\{ \sum_{i=1}^{N-K} \psi(v_i) x_i^{r^*}(v^{N-K}) \right\}$, where v^{N-K} is the value profile of the $N - K$ bidders. We learn from Section 2.2 that the first $k \leq \min\{K, N - K\}$ bidders whose values are greater than r^* will be located in descending order of values at the first k slots. We use the function $\max_K : \mathbb{R}^N \mapsto \mathbb{R}^K$ to compute a vector consisting of the first K highest elements, sorted in descending order, from a vector with N elements. Then the optimal expected revenue can be represented as $E_{v^{N-K}} \left\{ \max_K(\psi(v_1), \dots, \psi(v_{N-K}), 0^K) \cdot \theta \right\}$, where 0^K is a vector of K zeros.

When another K bidders, labeled by $N - K + 1, N - K + 2, \dots, N$, are added to the auction with value profile v^K , we have that $E_{v^K} \left\{ \psi(v_i) \right\} = \psi(v_i)$ for $i = 1, 2, \dots, N - K$ and $E_{v^K} \left\{ \psi(v_j) \right\} = E_{v_{N-K+1}^K} \left\{ E_{v_j^K} \left\{ \psi(v_j) \right\} \right\} = E_{v_{N-K+1}^K} \left\{ 0 \right\} = 0$ for $j = N - K + 1, N - K + 2, \dots, N$. Thus,

$$\begin{aligned} & E_{v^{N-K}} \left\{ \max_K(\psi(v_1), \dots, \psi(v_{N-K}), 0^K) \cdot \theta \right\} \\ &= E_{v^{N-K}} \left\{ \max_K(E_{v^K} \left\{ \psi(v_1) \right\}, \dots, E_{v^K} \left\{ \psi(v_N) \right\}) \cdot \theta \right\} \\ &\leq E_{v^{N-K}} \left\{ E_{v^K} \left\{ \max_K(\psi(v_1), \dots, \psi(v_N)) \cdot \theta \right\} \right\} \\ &= E_v \left\{ \max_K(\psi(v_1), \psi(v_2), \dots, \psi(v_N)) \cdot \theta \right\} \\ &= E_v \left\{ \sum_{i=1}^N \psi(v_i) x_i(v) \right\}, \end{aligned} \quad (4)$$

where the last equation is the expected revenue with N bidders and zero reserve price. \square

Lemma 3. *The expected revenue of the optimal GSP with $N - K$ bidders and K slots is at least $\frac{N-K}{N}$ that of the optimal GSP with N bidders and K slots.*

Proof. For any advertiser i in the optimal GSP with $N - K$ bidders, new bidders will not rise i 's rank, i.e., $x_i^{r^*}(v^{N-K}) \geq$

$$\begin{aligned}
& x_i^{r^*}(v^{N-K}, v^K) \text{ for any } v^{N-K} \text{ and } v^K. \text{ It thus follows that} \\
& E_{v^{N-K}}\{\psi(v_i)x_i^{r^*}(v^{N-K})\} = E_{v^K}\{E_{v^{N-K}}\{\psi(v_i)x_i^{r^*}(v^{N-K})\}\} \\
& \geq E_{v^K}\{E_{v^{N-K}}\{\psi(v_i)x_i^{r^*}(v^{N-K}, v^K)\}\} \\
& = E_v\{\psi(v_i)x_i^{r^*}(v)\}. \tag{5}
\end{aligned}$$

The proof is completed by multiplying the first and last equations by $N - K$ and $\frac{N-K}{N} \cdot N$, respectively. \square

Next, we first present the method to compute \bar{r} with sampled data, based on which we derive the ratio $\frac{R(\bar{r})}{R(r^*)}$ for GSP in Theorem 5.

The optimal reserve price r^* satisfies $\psi(r^*) = 0$, which is equivalent to $-\psi(r^*)f(r^*) = 0$. Readers can easily verify that the primitive function of $-\psi(r)f(r)$ is $P(r) = r(1 - F(r))$. Regularity of $F(\cdot)$ thus implies that $P(r)$ increases with respect to r on $(0, r^*)$ and decreases on $(r^*, +\infty)$ and hence reaches its maximum at r^* . Given m samples from F , renamed so that $v_{(1)} \geq v_{(2)} \geq \dots \geq v_{(m)}$, a straightforward idea for finding the optimal reserve price is to view $\frac{i}{m}$ as an approximation to $1 - F(v_{(i)})$, $i = 1, 2, \dots, m$, and then solve the following problem: $\arg \max_{v_{(i)}} \frac{i}{m} v_{(i)}$, s.t. $1 \leq i \leq m$. This naive method is not feasible and may give overly large estimated reserve price with a heavy-tailed distribution $F(\cdot)$. A common practice is to forbid the largest samples from acting as reserve prices, leading to the *guarded empirical reserve* (Cole and Roughgarden 2014; Dhangwatnotai, Roughgarden, and Yan 2015):

$$\bar{r} = \arg \max_{v_{(i)}} \frac{i}{m} \cdot v_{(i)}, \text{ s.t. } \delta m \leq i \leq m. \tag{6}$$

where $0 < \delta < 1$ is an accuracy parameter. The following lemma demonstrates the sample complexity for \bar{r} to be $(1 - \epsilon)$ -optimal in the point of view of $P(r^*)$, which serves as a key component of finding the bound for $\frac{R(\bar{r})}{R(r^*)}$.

Lemma 4 ((Dhangwatnotai, Roughgarden, and Yan 2015), Lemma 4.1²). *For every regular distribution $F(\cdot)$ and $0 < \epsilon < 1$, the following statement holds: with probability at least $1 - \epsilon$, the guarded empirical reserve \bar{r} of Eq. (6) with $m = c\epsilon^{-3} \ln \epsilon^{-1}$ samples from $F(\cdot)$ is a $(1 - \epsilon)$ -optimal reserve price, meaning that $P(\bar{r}) \geq (1 - \epsilon)P(r^*)$ where c is a positive constant number.*

Based on Lemma 4, we can prove the following theorem for the lower bound of $\frac{R(\bar{r})}{R(r^*)}$.

Theorem 5. *For GSP mechanism and \bar{r} obtained from Eq. (6), we have that $\frac{R(\bar{r})}{R(r^*)} \geq (1 - \epsilon)^2$ for any i.i.d. and regular value distributions with $m = c\epsilon^{-3} \ln \epsilon^{-1}$ samples.*

Proof. Fix bidder i and v_{-i} , and let t be the minimal bid for i to be a winner without regard to the reserve price. The expected revenue generated from bidder i with r^* and \bar{r} are $\int_{\max\{r^*, t\}}^{+\infty} \psi(v_i)f(v_i)x_i(v)dv_i$ and

$\int_{\max\{\bar{r}, t\}}^{+\infty} \psi(v_i)f(v_i)x_i(v)dv_i$, respectively. There are three situations depending on the maximum value of r^* , \bar{r} and t . Consider the first case, $\max\{r^*, \bar{r}, t\} = t$, which leads to the same revenue for the auctioneer using r^* or \bar{r} . Now consider the second case, $\max\{r^*, \bar{r}, t\} = r^*$. We assume that there are n bidders, excluding bidder i , whose values are greater than r^* in this case. Then the auctioneer's expected revenue from i with optimal reserve price, denoted by R^* , is $R^* = \theta_{n+1} \int_{r^*}^{v_{(n)}} \phi(v_i) + \theta_n \int_{v_{(n)}}^{v_{(n-1)}} \phi(v_i) + \dots + \theta_1 \int_{v_{(1)}}^{+\infty} \phi(v_i) = \theta_{n+1}(P(r^*) - P(v_{(n)})) + \theta_n(P(v_{(n)}) - P(v_{(n-1)})) + \dots + \theta_1(P(v_{(1)}) - 0) = \theta_{n+1}P(r^*) + \sum_{j=1}^n (\theta_j - \theta_{j+1})P(v_{(j)})$, where $\phi(v_i) = \psi(v_i)f(v_i)dv_i$ for simplicity. Suppose that without consideration of bidder i , there are k bidders with values less than r^* but greater than $\max\{\bar{r}, t\}$. When $k = 0$, the revenues from i with \bar{r} and r^* are the same. When $k > 0$, the auctioneer's expected revenue contributed by i with \bar{r} , represented with \bar{R} , is $\bar{R} = \theta_{n+k+1} \int_{\max\{\bar{r}, t\}}^{v_{(n+k)}} \phi(v_i) + \theta_{n+k} \int_{v_{(n+k)}}^{v_{(n+k-1)}} \phi(v_i) + \dots + \theta_1 \int_{v_{(1)}}^{+\infty} \phi(v_i) = \theta_{n+k+1}P(\max\{\bar{r}, t\}) + \sum_{j=1}^{n+k} (\theta_j - \theta_{j+1})P(v_{(j)})$. Since $P(r)$ is increasing on $(0, r^*)$, and $r^* > v_{(n+j)} > \max\{\bar{r}, t\} \geq \bar{r}$ for $j = 1, 2, \dots, k$ in the second case, and $P(\bar{r}) \geq (1 - \epsilon)P(r^*)$, the equation $P(v_{(n+j)}) > P(\max\{\bar{r}, t\}) \geq P(\bar{r}) \geq (1 - \epsilon)P(r^*)$ holds for $j = 1, 2, \dots, k$. Hence, we have $\bar{R} \geq \theta_{n+k+1}(1 - \epsilon)P(r^*) + \sum_{j=n+1}^{n+k} (\theta_j - \theta_{j+1})(1 - \epsilon)P(r^*) + \sum_{j=1}^n (\theta_j - \theta_{j+1})P(v_{(j)}) = (1 - \epsilon)\theta_{n+1}P(r^*) + \sum_{j=1}^n (\theta_j - \theta_{j+1})P(v_{(j)}) \geq (1 - \epsilon)\theta_{n+1}P(r^*) + (1 - \epsilon) \sum_{j=1}^n (\theta_j - \theta_{j+1})P(v_{(j)}) = (1 - \epsilon)R^*$.

The analysis for the third case where $\max\{r^*, \bar{r}, t\} = \bar{r}$ is similar. Overall, the expected revenue from bidder i with guarded empirical reserve is at least a $(1 - \epsilon)$ fraction of that of the optimal GSP mechanism in each case. Taking expectation over v_{-i} and i , and considering the probability $(1 - \epsilon)$, we have that $\frac{R(\bar{r})}{R(r^*)} \geq (1 - \epsilon)^2$. \square

5 Sample Size Optimization

In this section we formulate the problem of finding the optimal sample size as a constrained optimization problem and prove the uniqueness of the solution.

Considering that $\frac{R(0)}{R(r^*)} \geq \frac{N-K}{N}$ (Theorem 1) and $\frac{R(\bar{r})}{R(r^*)} \geq (1 - \epsilon)^2$ with $\tau N = c\epsilon^{-3} \ln \epsilon^{-1}$ (Theorem 5), we maximize the worst-case total revenue:

$$\max_{\tau} \widehat{\mathcal{R}}(\tau) = \tau \frac{N-K}{NM} + (1 - \frac{\tau}{M})(1 - \epsilon)^2, \tag{7}$$

$$\text{s.t. } \begin{cases} \tau \in \{1, 2, \dots, M\}; \\ \tau = \frac{c\epsilon^{-3} \ln \epsilon^{-1}}{N}. \end{cases} \tag{8}$$

Since the bounds we find are prior-independent, the revenue generated with τ^* is robust for any distribution, where τ^* is the solution of the above optimization problem.

We see from Eqs. (7) and (8) that it is difficult to explicitly express ϵ using τ , so we replace τ with ϵ and get the following

²We restate the lemma where we let $\epsilon = \delta$. Readers can check the proof in (Dhangwatnotai, Roughgarden, and Yan 2015) and verify that Lemma 4 always holds if ϵ is bounded by $(0, 1)$.

equivalent formulation:

$$\max_{\epsilon} \mathcal{R}(\epsilon) = c\epsilon^{-3} \ln \epsilon^{-1} \frac{N-K}{N^2 M} + \left(1 - \frac{c\epsilon^{-3} \ln \epsilon^{-1}}{NM}\right) (1-\epsilon)^2, \quad (9)$$

$$\text{s.t. } 1 \leq \frac{c\epsilon^{-3} \ln \epsilon^{-1}}{N} \leq M. \quad (10)$$

We have the following theorem for the above optimization problem.

Theorem 6 (Uniqueness of ϵ^*). *The solution ϵ^* of the optimization problem defined in Eqs. (9) and (10) is unique.*

The proof Theorem 6 is straightforward based on the unimodal property of $\mathcal{R}(\epsilon)$ demonstrated in Theorem 7.

Theorem 7 (Unimodal property of $\mathcal{R}(\epsilon)$). *$\mathcal{R}(\epsilon)$ is not concave on $(0, 1)$, but there exists an $s \in (0, 1)$ such that $\mathcal{R}'(\epsilon) > 0$ for $\epsilon \in (0, s)$, $\mathcal{R}'(s) = 0$ and $\mathcal{R}'(\epsilon) < 0$ for $\epsilon \in (s, 1)$.*

Proof. We use $0 < \alpha < 1$ to denote $\frac{K}{N}$ in this proof. By taking the first-order derivative of $\mathcal{R}(\epsilon)$, we have that $\mathcal{R}'(\epsilon) = -2(1-\epsilon) - \frac{c \ln \epsilon^{-1}}{NM} (-3\alpha\epsilon^{-4} + 4\epsilon^{-3} - \epsilon^{-2}) + \frac{c}{NM} (\alpha\epsilon^{-4} - 2\epsilon^{-3} + \epsilon^{-2})$. We can check that $\lim_{\epsilon \rightarrow 0^+} \mathcal{R}'(\epsilon) \rightarrow +\infty > 0$ and $\lim_{\epsilon \rightarrow 1^-} \mathcal{R}'(\epsilon) \rightarrow \frac{c(\alpha-1)}{NM} < 0$. If the statement that there is a point $\epsilon' \in (0, 1)$ such that $\mathcal{R}'(\epsilon)$ is decreasing on $(0, \epsilon')$ and increasing on $(\epsilon', 1)$ holds, then combining with the fact that $\mathcal{R}'(0^+) > 0$ and $\mathcal{R}'(1^-) < 0$, Theorem 7 is proved. Next we will demonstrate the validity of the statement, which is equivalent to proving that $\mathcal{R}''(\epsilon)$ is negative on $(0, \epsilon')$ and positive on $(\epsilon', 1)$.

The second-order derivative of $\mathcal{R}(\epsilon)$ is $\mathcal{R}''(\epsilon) = 2 - \frac{c}{NM} (\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon))$, where $\mathcal{R}_1''(\epsilon) = \epsilon^{-5} \ln \epsilon^{-1} (2\epsilon^2 - 12\epsilon + 12\alpha)$ and $\mathcal{R}_2''(\epsilon) = \epsilon^{-5} (3\epsilon^2 - 10\epsilon + 7\alpha)$. It is easy to get that $\lim_{\epsilon \rightarrow 0^+} (\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon)) \rightarrow +\infty$. The zero points of the convex functions $2\epsilon^2 - 12\epsilon + 12\alpha$ and $3\epsilon^2 - 10\epsilon + 7\alpha$ are $3 \pm \sqrt{9 - 6\alpha}$ and $\frac{5 \pm \sqrt{25 - 21\alpha}}{3}$ respectively. Further analysis indicates that $0 < \frac{5 - \sqrt{25 - 21\alpha}}{3} < \eta \leq 1 < \frac{5 + \sqrt{25 - 21\alpha}}{3} < 3 + \sqrt{9 - 6\alpha}$, where $\eta = \min\{3 - \sqrt{9 - 6\alpha}, 1\}$, hence $\mathcal{R}_2''(\eta) < 0$. Since $\mathcal{R}_1''(\eta) = 0$, we have $\mathcal{R}_1''(\eta) + \mathcal{R}_2''(\eta) < 0$. We can further prove that $\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon)$ is convex on $(0, \eta)$ (see Lemma 8). So there exists a point ϵ'' on $(0, \eta)$ such that $\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon)$ decreases monotonically from positive infinity to zero as ϵ rises from 0^+ to ϵ'' , and then becomes negative when $\epsilon \in (\epsilon'', \eta)$. Because ϵ^{-5} and $\ln \epsilon^{-1}$ are always positive on $(0, 1)$, $\mathcal{R}_1''(\epsilon)$ and $\mathcal{R}_2''(\epsilon)$ are negative on $(\eta, 1)$, which implies that $\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon)$ is negative on $(\epsilon'', 1)$.

Since $\mathcal{R}''(\epsilon)$ has the opposite monotonicity of $\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon)$, $\mathcal{R}''(\epsilon)$ increases monotonically from negative infinity to 2 as ϵ increases from 0^+ to ϵ'' , and keeps positive on $(\epsilon'', 1)$. Then a point $\epsilon' \in (0, \epsilon'')$ exists such that the statement holds. \square

Lemma 8 (Concavity of $\mathcal{R}''(\epsilon)$). *The second-order derivative $\mathcal{R}''(\epsilon)$ of $\mathcal{R}(\epsilon)$ is concave on $(0, \eta)$, where $\eta = \min\{3 - \sqrt{9 - 6K/N}, 1\}$.*

Proof. Let $\alpha = \frac{K}{N}$ for simplicity. $\mathcal{R}''(\epsilon) = 2 - \frac{c}{NM} (\mathcal{R}_1''(\epsilon) + \mathcal{R}_2''(\epsilon))$, where $\mathcal{R}_1''(\epsilon) = \epsilon^{-5} \ln \epsilon^{-1} (2\epsilon^2 - 12\epsilon + 12\alpha)$ and

$\mathcal{R}_2''(\epsilon) = \epsilon^{-5} (3\epsilon^2 - 10\epsilon + 7\alpha)$. We just need to prove that both $\mathcal{R}_1''(\epsilon)$ and $\mathcal{R}_2''(\epsilon)$ are convex on $(0, \eta)$.

We first prove the convexity of $\mathcal{R}_1''(\epsilon)$, which is represented as the product of the following two equations: $g_1(\epsilon) = \epsilon^{-5} \ln \epsilon^{-1}$ and $g_2(\epsilon) = 2\epsilon^2 - 12\epsilon + 12\alpha$. The convexity of $\mathcal{R}_1''(\epsilon) = g_1(\epsilon)g_2(\epsilon)$ means that the limitation $\lim_{\Delta \rightarrow 0^+} \frac{g_1(\epsilon+\Delta)g_2(\epsilon+\Delta) - g_1(\epsilon)g_2(\epsilon)}{\Delta}$ is increasing w.r.t. ϵ . We use $\Delta_{g_i(\epsilon)}$ to denote $g_i(\epsilon) - g_i(\epsilon+\Delta)$, $i = 1, 2$. Given $\Delta > 0$, we have that $g_1(\epsilon+\Delta)g_2(\epsilon+\Delta) - g_1(\epsilon)g_2(\epsilon) = (g_1(\epsilon) - \Delta_{g_1(\epsilon)})(g_2(\epsilon) - \Delta_{g_2(\epsilon)}) - g_1(\epsilon)g_2(\epsilon) = -\Delta_{g_2(\epsilon)}(g_1(\epsilon) - \Delta_{g_1(\epsilon)}) - \Delta_{g_1(\epsilon)}g_2(\epsilon) = -\Delta_{g_2(\epsilon)}g_1(\epsilon+\Delta) - \Delta_{g_1(\epsilon)}g_2(\epsilon)$. Suppose ϵ_1 and ϵ_2 satisfy $\epsilon_1 < \epsilon_2$ in $(0, \eta)$. It is easy to verify that $g_1(\epsilon)$ and $g_2(\epsilon)$ are positive decreasing convex functions on $(0, \eta)$, which implies that $g_1(\epsilon_1 + \Delta) > g_1(\epsilon_2 + \Delta) > 0$, $g_2(\epsilon_1) > g_2(\epsilon_2) > 0$ and $\Delta_{g_i(\epsilon_1)} > \Delta_{g_i(\epsilon_2)} > 0$. Hence the equation $-\Delta_{g_2(\epsilon_1)}g_1(\epsilon_1 + \Delta) - \Delta_{g_1(\epsilon_1)}g_2(\epsilon_1) < -\Delta_{g_2(\epsilon_2)}g_1(\epsilon_2 + \Delta) - \Delta_{g_1(\epsilon_2)}g_2(\epsilon_2)$ holds for any $\Delta > 0$ and $\epsilon_1 < \epsilon_2$ in $(0, \eta)$. Then we can conclude that $-\Delta_{g_2(\epsilon)}g_1(\epsilon+\Delta) - \Delta_{g_1(\epsilon)}g_2(\epsilon)$ is increasing w.r.t. $\epsilon \in (0, \eta)$, so is the limitation. Thus $\mathcal{R}_1''(\epsilon) = g_1(\epsilon)g_2(\epsilon)$ is convex on $(0, \eta)$.

Next we focus on $\mathcal{R}_2''(\epsilon)$. The second-order derivative of $\mathcal{R}_2''(\epsilon)$ is $2\epsilon^{-7} (18\epsilon^2 - 100\epsilon + 105\alpha)$, which is positive in the support $(0, \frac{50 - \sqrt{2500 - 1890\alpha}}{18})$. Furthermore, we can verify that $\eta \leq 3 - \sqrt{9 - 6\alpha} < \frac{50 - \sqrt{2500 - 1890\alpha}}{18}$. So $\mathcal{R}_2''(\epsilon)$ is convex (with positive second-order derivative) on $(0, \eta)$. \square

As declared in Theorem 7, $\mathcal{R}(\epsilon)$ is not a concave function on $(0, 1)$. Fortunately, $\mathcal{R}(\epsilon)$ is proved to be unimodal, based on which we propose an efficient algorithm to maximize it. Let $[l, 1) \subset (0, 1)$, $l > 0$, denote the domain of ϵ constrained by Eqs. (7) and (9), and then the optimal ϵ^* of Eqs. (9) and (10) can be computed as

$$\epsilon^* = \begin{cases} s, & \text{if } \mathcal{R}'(l) > 0 \text{ (i.e., } l < s); \\ l, & \text{otherwise.} \end{cases} \quad (11)$$

Next we give the idea on how to find τ^* for the problem defined in Eq. (7) and (8) based on ϵ^* .

We first define the function g as $t = g(\epsilon) = \frac{c\epsilon^{-3} \ln \epsilon^{-1}}{N}$, $\epsilon \in [l, 1)$, where $l = g^{-1}(M)$. Because $g(\epsilon)$ is a decreasing function on $[l, 1)$, so is $g^{-1}(t)$ for $t \in (0, g(l)]$. Thus $\widehat{\mathcal{R}}(t) = \mathcal{R}(g^{-1}(t))$ has the opposite monotonicity of $\mathcal{R}(\epsilon)$ in the corresponding intervals. Specifically, the increasing and decreasing intervals of $\widehat{\mathcal{R}}(t)$ are $(0, g(\epsilon^*)]$ and $[g(\epsilon^*), g(l)]$ respectively, which implies that $\widehat{\mathcal{R}}(t)$ is also a unimodal on $(0, g(l)]$ and the unique maximal point is $t^* = g(\epsilon^*)$. Then τ^* can be represented with respect to t^* .

Corollary 9. *The optimal sample size τ^* satisfies*

$$\tau^* = \begin{cases} \lceil t^* \rceil, & \text{if } \widehat{\mathcal{R}}(\lceil t^* \rceil) > \widehat{\mathcal{R}}(\lfloor t^* \rfloor); \\ \lfloor t^* \rfloor, & \text{otherwise.} \end{cases} \quad (12)$$

Proof. First note that $g(l) = M$ according to Eqs. (7) and (9), and hence $\{1, 2, \dots, M\} \subset (0, g(l)]$. Since $\lceil t^* \rceil \geq t^* \geq \lfloor t^* \rfloor$, according to the monotonicity of $\widehat{\mathcal{R}}(\cdot)$, we have that $\widehat{\mathcal{R}}(\tau) < \widehat{\mathcal{R}}(\lceil t^* \rceil)$, if $\tau > \lceil t^* \rceil$; $\widehat{\mathcal{R}}(\tau) <$

$\widehat{\mathcal{R}}(\lfloor t^* \rfloor)$, if $\tau < \lfloor t^* \rfloor$, which implies that $\widehat{\mathcal{R}}(\tau)$ reaches its maximum at either $\lceil t^* \rceil$ (when $\widehat{\mathcal{R}}(\lceil t^* \rceil) > \widehat{\mathcal{R}}(\lfloor t^* \rfloor)$) or $\lfloor t^* \rfloor$ (when $\widehat{\mathcal{R}}(\lfloor t^* \rfloor) > \widehat{\mathcal{R}}(\lceil t^* \rceil)$). \square

Theorem 10 (Uniqueness of optimal sample size). *The optimal sample size τ^* is unique.*

The proof is straightforward based on Theorem 6 and Corollary 9.

6 Algorithm and Experimental Evaluation

Algorithm 1: Optimal sample size

```

1 lower  $\leftarrow$  0, upper  $\leftarrow$  1, s  $\leftarrow$  1/2;
2 while  $|\mathcal{R}'(s)| \neq 0$  do
3   if  $\mathcal{R}'(s) > 0$  then lower  $\leftarrow$  s;
4   ;
5   else upper  $\leftarrow$  s;
6   ;
7   s  $\leftarrow$  (lower + upper)/2;
8 Compute  $\epsilon^*$  based on Eq. (11),  $t^* \leftarrow g(\epsilon^*)$ ;
9 Use Eq. (12) to compute  $\tau^*$ ;
10 return  $\tau^*$ ;

```

Based on Theorem 7 and Corollary 9, we propose Algorithm 1 to compute τ^* . The first step is to iteratively compute s based on the unimodal property of $\mathcal{R}(\epsilon)$ (Lines 1-5), i.e., if the current point is in the increasing interval of $\mathcal{R}(\cdot)$, search in the right side of $(lower, upper)$ in the next iteration; otherwise, search in the left side of $(lower, upper)$. This procedure is similar to the binary search and hence has exponential convergence rate. Then ϵ^* and τ^* are computed based on s in Lines 6 and 7.

We use the commonly used log normal distribution and exponential distribution to evaluate the performance of Algorithm 1. The parameters for the former distribution are set as $\mu = 0$, $\sigma = 1.5$ and the corresponding optimal reserve price r^* is 4.2755. The parameter for the latter is $\lambda = 3$, with which we have that $r^* = 3.0005$.

We first analyze the robustness of Algorithm 1. M can be estimated by the auctioneer with historical data, but the estimation could be inaccurate. We use $\tau^{(5)}$ to represent the sample sizes computed by Algorithm 1 when the errors of the estimation for M is 5%. The overall revenues with τ^* and $\tau^{(5)}$ samples are presented in Tables 1 and 2, where $\widehat{\mathcal{R}}(\cdot)$ for each setting ($N = 5$) is averaged over 400 instances. The theoretical maximal revenue when the auctioneer knows bidders value distribution is 1. We see that τ^* can gain very high total revenue ($\widehat{\mathcal{R}}(\tau^*) > 0.9$ in all the experiments) and Algorithm 1 is robust against the inaccurate estimation for M given that the performance degradation is very tiny.

Next we compare the revenue generated with τ^* samples with three baseline strategies. The first two, denoted by S_1 and S_2 , are fraction-based strategies, which use $0.1 \cdot M$ and $0.3 \cdot M$ rounds for sampling respectively. The third one (S_3) serves as a bench mark which sets the reserve price as zero for the auctions. $G(S_j)$ measures the relative revenue gain of our method over the baseline S_j , $j = 1, 2, 3$. The larger $G(S_j)$ is,

the better our method is over S_j . We see from Tables 1 and 2 that our proposed algorithm outperforms the other three strategies in all settings. We find that $G(S_1)$ decreases while $G(S_2)$ increases as K grows. That is because a larger K usually leads to more cost for sampling according to Theorem 1, and hence smaller-fraction strategy is more preferred. The comparison with S_3 implies that setting proper reserve prices can significantly improve the auctioneer's revenue. Overall, the strategy we proposed can achieve the highest revenue for the auctioneer among all the strategies.

M	K	$\widehat{\mathcal{R}}(\tau^*)$	$\widehat{\mathcal{R}}(\tau^{(5)})$	$G(S_1)$	$G(S_2)$	$G(S_3)$
5e+6	2	.9199	.9174	13.2%	0.54%	16.2%
	3	.9079	.9047	6.65%	0.61%	45.5%
	4	.9005	.8969	3.14%	4.10%	101%
5e+7	2	.9614	.9601	1.67%	1.15%	20.2%
	3	.9513	.9496	0.28%	5.37%	50.8%
	4	.9449	.9428	0.09%	11.6%	114%

Table 1: Evaluation for log normal distribution

M	K	$\widehat{\mathcal{R}}(\tau^*)$	$\widehat{\mathcal{R}}(\tau^{(5)})$	$G(S_1)$	$G(S_2)$	$G(S_3)$
5e+6	2	.9677	.9542	11.2%	0.48%	15.3%
	3	.9446	.9401	5.64%	0.56%	43.5%
	4	.9186	.9045	3.10%	3.89%	89.1%
5e+7	2	.9842	.9837	1.55%	1.03%	18.7%
	3	.9700	.9652	0.22%	4.97%	47.6%
	4	.9542	.9450	0.08%	11.2%	105%

Table 2: Evaluation for exponential distribution.

7 Conclusion and Discussion

In this paper, we studied the problem of finding the optimal sample size for finitely repeated GSP auctions. We analyzed the competitive ratios of expected revenues during and after sampling for GSP and formulated the problem as a constrained mixed integer non-convex program. We proved the uniqueness of the optimal sample size. The evaluation showed that the solution we provided outperforms baseline strategies.

Our results can be extended to weighted GSP and VCG. We can follow (Sun, Zhou, and Deng 2014) to assume $s_i = v_i e_i$ is drawn from an i.i.d. regular distribution $F(s)$, where e_i is bidder i 's ad quality. The optimal reserve price for bidder i is modified as s^*/e_i , where s^* satisfies $\psi(s^*) = 0$. The proofs in Section 4 still hold if we replace v_i and b_i with $v_i e_i$ and $b_i e_i$ respectively. Thus all results can be extended to weighted GSP. Besides, it is known that bidders' payments are the same for the (dominant) truth-telling equilibrium with VCG mechanism and the SNE with GSP mechanism. Hence, if we assume bidders in VCG auctions to have i.i.d. values and to bid truthfully, all the results can be directly extended to VCG.

In future work, we will try other (non-zero) reserve prices for sampling and investigate their convergence properties. Besides, we will consider other strategic behavior model of advertisers.

8 Acknowledgments

This paper is supported by 2015 Microsoft Research Asia Collaborative Research Program.

References

- Abrams, Z., and Schwarz, M. 2007. Ad auction design and user experience. In *Internet and Network Economics*. 529–534.
- Bubeck, S.; Munos, R.; Stoltz, G.; and Szepesvari, C. 2011. X-armed bandits. *Journal of Machine Learning Research* 12(May):1655–1695.
- Bulow, J., and Klemperer, P. 1994. Auctions vs. negotiations. Technical report, National Bureau of Economic Research.
- Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2011. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*, 81–90.
- Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2014. Revenue guarantees in the generalized second price auction. *ACM Transactions on Internet Technology (TOIT)* 14(2-3):17.
- Cesa-Bianchi, N.; Gentile, C.; and Mansour, Y. 2013. Regret minimization for reserve prices in second-price auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1190–1204.
- Chakraborty, T.; Even-Dar, E.; Guha, S.; Mansour, Y.; and Muthukrishnan, S. 2010. Selective call out and real time bidding. In *Internet and Network Economics*. 145–157.
- Chen, Y.; Berkhin, P.; Anderson, B.; and Devanur, N. R. 2011. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1307–1315.
- Chen, W.; He, D.; Liu, T.-Y.; Qin, T.; Tao, Y.; and Wang, L. 2014. Generalized second price auction with probabilistic broad match. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 39–56.
- Cole, R., and Roughgarden, T. 2014. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 243–252.
- Cope, E. W. 2009. Regret and convergence bounds for a class of continuum-armed bandit problems. *IEEE Transactions on Automatic Control* 54(6):1243–1253.
- Dhangwatnotai, P.; Roughgarden, T.; and Yan, Q. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* 91:318–333.
- Edelman, B., and Schwarz, M. 2010. Optimal auction design and equilibrium selection in sponsored search auctions. *The American Economic Review* 597–602.
- Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American economic review* 97(1):242–259.
- Hartline, J. D., and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234.
- Hartline, J. 2006. Lectures on optimal mechanism design. *Lecture notes*.
- Jansen, B. J., and Mullen, T. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business* 6(2):114–131.
- Jia, Y. Y., and Mannor, S. 2011. Unimodal bandits. In *Proceedings of the 28th International Conference on Machine Learning*, 41–48.
- Kleinberg, R., and Leighton, T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, 594–605.
- Lucier, B.; Paes Leme, R.; and Tardos, E. 2012. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*, 361–370.
- Ma, W.; Wu, T.; Qin, T.; and Liu, T.-Y. 2014. Generalized second price auctions with value externalities. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 1457–1458.
- Milgrom, P. 2010. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior* 70(1):62–70.
- Mohri, M., and Medina, A. M. 2015. Non-parametric revenue optimization for generalized second price auctions. In *Proceedings of the 31st conference on Uncertainty in Artificial Intelligence*.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6(1):58–73.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge.
- Ostrovsky, M., and Schwarz, M. 2011. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*, 59–60.
- Qin, T.; Chen, W.; and Liu, T.-Y. 2015. Sponsored search auctions: Recent advances and future directions. *ACM Transactions on Intelligent Systems and Technology* 5(4):60:1–60:34.
- Radovanovic, A., and Heavlin, W. D. 2012. Risk-aware revenue maximization in display advertising. In *Proceedings of the 21st international conference on World Wide Web*, 91–100.
- Rong, J.; Qin, T.; An, B.; and Liu, T.-Y. 2016. Modeling bounded rationality for sponsored search auctions. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, 515–523.
- Sun, Y.; Zhou, Y.; Yin, M.; and Deng, X. 2012. On the convergence and robustness of reserve pricing in keyword auctions. In *Proceedings of the 14th Annual International Conference on Electronic Commerce*, 113–120.
- Sun, Y.; Zhou, Y.; and Deng, X. 2014. Optimal reserve prices in weighted GSP auctions. *Electronic Commerce Research and Applications* 13(3):178–187.
- Thompson, D. R., and Leyton-Brown, K. 2013. Revenue optimization in the generalized second-price auction. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 837–852.
- Varian, H. R. 2007. Position auctions. *international Journal of industrial Organization* 25(6):1163–1178.
- Xiao, B.; Yang, W.; and Li, J. 2009. Optimal reserve price for the generalized second-price auction in sponsored search advertising. *Journal of Electronic Commerce Research* 10(3):114.
- Yuan, S.; Wang, J.; Chen, B.; Mason, P.; and Seljan, S. 2014. An empirical study of reserve price optimization in real-time bidding. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1897–1906.
- Yuan, S.; Wang, J.; and Zhao, X. 2013. Real-time bidding for online advertising: measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*.