Revenue Guarantees in Sponsored Search Auctions*

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou

Computer Technology Institute and Press "Diophantus" &
Department of Computer Engineering and Informatics,
University of Patras, 26504 Rio, Greece

Abstract. Sponsored search auctions are the main source of revenue for search engines. In such an auction, a set of utility-maximizing advertisers compete for a set of ad slots. The assignment of advertisers to slots depends on bids they submit; these bids may be different than the true valuations of the advertisers for the slots. Variants of the celebrated VCG auction mechanism guarantee that advertisers act truthfully and, under mild assumptions, lead to revenue or social welfare maximization. Still, the sponsored search industry mostly uses generalized second price (GSP) auctions; these auctions are known to be non-truthful and sub-optimal in terms of social welfare and revenue. In an attempt to explain this tradition, we study a Bayesian setting where the valuations of advertisers are drawn independently from a regular probability distribution. In this setting, it is well known by the work of Myerson (1981) that the optimal revenue is obtained by the VCG mechanism with a particular reserve price that depends on the probability distribution. We show that by appropriately setting the reserve price, the revenue over any Bayes-Nash equilibrium of the game induced by the GSP auction is at most a small constant fraction of the optimal revenue, improving recent results of Lucier, Paes Leme, and Tardos (2012). Our analysis is based on the Bayes-Nash equilibrium conditions and on the properties of regular probability distributions.

1 Introduction

The sale of advertising space is the main source of income for information providers on the Internet. For example, a query to a search engine creates advertising space that is sold to potential advertisers through auctions that are known as *sponsored search auctions* (or ad auctions). In their influential papers, Edelman et al. [6] and Varian [18] have proposed a (now standard) model for this process. According to this model, a set of utility-maximizing advertisers compete for a set of ad slots with non-increasing click-through rates. The auctioneer collects bids from the advertisers and assigns them to slots (usually, in

* This work is co-financed by the European Social Fund and Greek national funds through the research funding program Thales on "Algorithmic Game Theory".

non-increasing order of their bids). In addition, it assigns a payment per click to each advertiser. Depending on the way the payments are computed, different auctions can be defined. Typical examples are the Vickrey-Clark-Groves (VCG), the generalized second price (GSP), and the generalized first price (GFP) auction. Naturally, the advertisers are engaged as players in a strategic game defined by the auction; the bid submitted by each player is such that it maximizes her utility (i.e., the total difference of her valuation minus her payment over all clicks) given the bids of the other players. This behavior leads to equilibria, i.e., states of the induced game from which no player has an incentive to unilaterally deviate.

Traditionally, truthfulness has been recognized as an important desideratum in the Economics literature on auctions [11]. In truthful auctions, truth-telling is an equilibrium according to specific equilibrium notions (e.g., dominant strategy, Nash, or Bayes-Nash equilibrium). Such a mechanism guarantees that the social welfare (i.e., the total value of the players) is maximized. VCG is a typical example of a truthful auction [5, 8, 19]. In contrast, GSP auctions are not truthful [6, 18]; still, they are the main auction mechanisms used in the sponsored search industry adopted by leaders such as Google and Yahoo!

In an attempt to explain this prevalence, several papers have provided bounds on the social welfare of GSP auctions [2, 12, 13, 17] over different classes of equilibria (pure Nash, coarse-correlated, Bayes-Nash). The main message from these studies is that the social welfare is always a constant fraction of the optimal one. However, one would expect that revenue (as opposed to social welfare) maximization is the major concern from the point of view of the sponsored search industry. In this paper, following the recent paper by Lucier et al. [14], we aim to provide a theoretical justification for the wide adoption of GSP by focusing on the revenue generated by these auctions.

In order to model the inherent uncertainty in advertisers' beliefs, we consider a Bayesian setting where the advertisers have random valuations drawn independently from a common probability distribution. This is the classical setting that has been studied extensively since the seminal work of Myerson [15] for single-item auctions (which is a special case of ad auctions). The results of [15] carry over to our model as follows. Under mild assumptions, the revenue generated by a player in a Bayes-Nash equilibrium depends only on the distribution of the click-through rate of the ad slot the player is assigned to for her different valuations. Hence, two Bayes-Nash equilibria that correspond to the same allocation yield the same revenue even if they are induced by different auction mechanisms; this statement is known as *revenue equivalence*. The allocation that optimizes the expected revenue is one in which low-bidding advertisers are excluded and the remaining ones are assigned to ad slots in non-increasing order of their valuations. Such an allocation is a Bayes-Nash equilibrium of the variation of the VCG mechanism where an appropriate reserve price (the Myerson reserve) is set in order to exclude the low-bidding advertisers.

GSP auctions may lead to different Bayes-Nash equilibria [7] in which a player with a higher valuation is assigned with positive probability to a slot with lower click-through rate than another player with lower valuation. This implies that the revenue is suboptimal. Our purpose is to quantify the loss of revenue over all Bayes-Nash equilibria of GSP auctions by proving worst-case *revenue guarantees*. A revenue guarantee of ρ for an auction mechanism implies that, at any Bayes-Nash equilibrium, the revenue generated is at most ρ times smaller than the optimal one. Note that, it is not even clear whether Myerson reserve is the choice that minimizes the revenue guarantee in GSP auctions. This issue is the subject of existing experimental work (see [16]).

Recently, Lucier et al. [14] proved theoretical revenue guarantees for GSP auctions. Among other results for full information settings, they consider two different Bayesian models. When the advertisers' valuations are drawn independently from a common probability distribution with monotone hazard rate (MHR), GSP auctions with Myerson reserve have a revenue guarantee of 6. This bound is obtained by comparing the utility of players at the Bayes-Nash equilibrium with the utility they would have by deviating to a single alternative bid (and by exploiting the special properties of MHR distributions). The class of MHR distributions is wide enough and includes many common distributions (such as uniform, normal, and exponential). In the more general case where the valuations are regular, the same bound is obtained using a different reserve price. This reserve is computed using a prophet inequality [10]. Prophet inequalities have been proved useful in several Bayesian auction settings in the past [4, 9].

In this work, we consider the same Bayesian settings with [14] and improve their results. We show that when the players have i.i.d. valuations drawn from a regular distribution, there is a reserve price so that the revenue guarantee is at most 4.72. For MHR valuations, we present a bound of 3.46. In both cases, the reserve price is either Myerson's or another one that maximizes the revenue obtained by the player allocated to the first slot. The latter is computed by developing new prophet-like inequalities that exploit the particular characteristics of the valuations. Furthermore, we show that the revenue guarantee of GSP auctions with Myerson reserve is at most 3.90 for MHR valuations. In order to analyze GSP auctions with Myerson reserve, we extend the techniques recently developed in [2, 13] (see also [3]). The Bayes-Nash equilibrium condition implies that the utility of each player does not improve when she deviates to any other bid. This yields a series of inequalities which we take into account with different weights. These weights are given by families of functions that are defined in such a way that a relation between the revenue at a Bayes-Nash equilibrium and the optimal revenue is revealed; we refer to them as *deviation weight function families*.

The rest of the paper is structured as follows. We begin with preliminary definitions in Section 2. Our prophet-type bounds are presented in Section 3. The role of deviation weight function families in the analysis is explored in Section 4. Then, Section 5 is devoted to the proofs of our main statements for GSP auctions. We conclude with open problems in Section 6. Due to lack of space several proofs have been omitted.

2 Preliminaries

We consider a Bayesian setting with n players and n slots[1] where slot $j \in [n]$ has a click-through rate α_j that corresponds to the frequency of clicking an ad in slot j. We add an artificial $(n+1)$-th slot with click-through rate 0 and index the slots so that $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n \geq \alpha_{n+1} = 0$. Each player's valuation (per click) is non-negative and is drawn from a publicly known probability distribution.

The auction mechanisms we consider use a reserve price t and assign slots to players according to the bids they submit. Player i submits a bid $b_i(v_i)$ that depends on her valuation v_i; the bidding function b_i is the strategy of player i. Given a realization of valuations, let $\mathbf{b} = (b_1, \dots, b_n)$ denote a bid vector and define the random permutation π so that $\pi(j)$ is the player with the j-th highest bid (breaking ties arbitrarily). The mechanism assigns slot j to player $\pi(j)$ whenever $b_{\pi(j)} \geq t$; if $b_{\pi(j)} < t$, the player is not allocated any slot. In such an allocation, let $\sigma(i)$ denote the slot that is allocated to player i. This is well-defined when player i is assigned a slot; if this is not the case, we follow the convention that $\sigma(i) = n+1$. Given $\mathbf{b}$, the mechanism also defines a payment $p_i \geq t$ for each player i that is allocated a slot. Then, the *utility* of player i is $u_i(\mathbf{b}) = \alpha_{\sigma(i)}(v_i - p_i)$. A set of players' strategies is a Bayes-Nash equilibrium if no player has an incentive to deviate from her strategy in order to increase her expected utility. This means that for every player i and every possible valuation x, $\mathbb{E}[u_i(\mathbf{b})|v_i = x] \geq \mathbb{E}[u_i(b'_i, \mathbf{b}_{-i})|v_i = x]$ for every alternative bid b'_i. Note that the expectation is taken over the randomness of the valuations of the other players and the notation $(b'_i, \mathbf{b}_{-i})$ is used for the bid vector where player i has deviated to b'_i and the remaining players bid as in $\mathbf{b}$. The *social welfare* at a Bayes-Nash equilibrium $\mathbf{b}$ is $\mathcal{W}_t(\mathbf{b}) = \mathbb{E}[\sum_i \alpha_{\sigma(i)} v_i]$, while the revenue generated by the mechanism is $\mathcal{R}_t(\mathbf{b}) = \mathbb{E}[\sum_i \alpha_{\sigma(i)} p_i]$.

We focus on the case where the valuations of players are drawn independently from a common probability distribution $\mathcal{D}$ with probability density function f and cumulative distribution function F. Given a distribution $\mathcal{D}$ over players' valuations, the *virtual valuation* function is $\phi(x) = x - \frac{1-F(x)}{f(x)}$. We consider *regular* probability distributions where $\phi(x)$ is non-decreasing. The work of Myerson [15] implies that the expected revenue from player i at a Bayes-Nash equilibrium $\mathbf{b}$ of any auction mechanism is $\mathbb{E}[\alpha_{\sigma(i)}\phi(v_i)]$, i.e., it depends only on the allocation of player i and her virtual valuation. Hence, the total expected revenue is maximized when the players with non-negative virtual valuations are assigned to slots in non-increasing order of their virtual valuations and players with negative virtual valuations are not assigned any slot. A mechanism that imposes this allocation as a Bayes-Nash equilibrium (and, hence, is revenue-maximizing) is the celebrated VCG mechanism with reserve price t such that $\phi(t) = 0$. We refer to this as Myerson reserve and denote it by r in the following. We use the notation μ to denote such an allocation. Note that, in μ, players with zero virtual

[1] Our model can simulate cases where the number of slots is smaller than the number of players by adding fictitious slots with zero click-through rate.

valuation can be either allocated slots or not; such players do not contribute to the optimal revenue.

A particular subclass of regular probability distributions are those with *monotone hazard rate* (MHR). A regular distribution $\mathcal{D}$ is MHR if its hazard rate function $h(x) = f(x)/(1 - F(x))$ is non-decreasing. These distributions have some nice properties (see [1]). For example, $F(r) \leq 1 - 1/e$ and $\phi(x) \geq x - r$ for every $x \geq r$.

In this paper, we focus on the GSP mechanism. For each player i that is allocated a slot (i.e., with bid at least t), GSP computes her payment as the maximum between the reserve price t and the next highest bid $b_{\pi(i+1)}$ (assuming that $b_{\pi(n+1)} = 0$). As it has been observed in [7], GSP may not admit the allocation μ as a Bayes-Nash equilibrium. This immediately implies that the revenue over Bayes-Nash equilibria would be suboptimal. In order to capture the revenue loss due to the selfish behavior of the players, we use the notion of *revenue guarantee*.

Definition 1. *The revenue guarantee of an auction game with reserve price t is $\max_{\mathbf{b}} \frac{\mathcal{R}_{OPT}}{\mathcal{R}_t(\mathbf{b})}$, where $\mathbf{b}$ runs over all Bayes-Nash equilibria of the game.*

In our proofs, we use the notation σ to refer to the random allocation that corresponds to a Bayes-Nash equilibrium. Note that, a player with valuation strictly higher than the reserve has always an incentive to bid at least the reserve and be allocated a slot. When her valuation equals the reserve, she is indifferent between bidding the reserve or not participating in the auction. For auctions with Myerson reserve, when comparing a Bayes-Nash equilibrium to the revenue-maximizing allocation μ, we assume that a player with valuation equal to the reserve has the same behavior in both σ and μ (this implies that $\mathbb{E}[\sum_i \alpha_{\sigma(i)}] = \mathbb{E}[\sum_i \alpha_{\mu(i)}]$). This assumption is without loss of generality since such a player contributes zero to the optimal revenue anyway. In our proofs, we also use the random variable $o(j)$ to denote the player with the j-th highest valuation (breaking ties arbitrarily). Hence, $\mu(i) = o^{-1}(i)$ if the virtual valuation of player i is positive and $\mu(i) = n + 1$ if it is negative. When the virtual valuation of player i is zero, it can be either $\mu(i) = o^{-1}(i)$ or $\mu(i) = n + 1$.

When considering GSP auctions, we make the assumption that players are *conservative*: whenever the valuation of player i is v_i, she only selects a bid $b_i(v_i) \in [0, v_i]$ at Bayes-Nash equilibria. This is a rather natural assumption since any bid $b_i(v_i) > v_i$ is weakly dominated by bidding $b_i(v_i) = v_i$ [17].

In the following, we use the notation x^+ to denote $\max\{x, 0\}$ while the expression $x\mathbb{1}\{E\}$ equals x when the event E is true and 0 otherwise.

3 Achieving Minimum Revenue Guarantees

Our purpose in this section is to show that by appropriately setting the reserve price, we can guarantee a high revenue from the advertiser that occupies the first slot at any Bayes-Nash equilibrium. Even though this approach will not give us a "standalone" result, it will be very useful later when we will combine it with the analysis of GSP auctions with Myerson reserve. These bounds are similar in spirit to prophet inequalities in optimal stopping theory [10].

We begin with a simple lemma.

Lemma 1. *Consider n random valuations v_1, ..., v_n that are drawn i.i.d. from a regular distribution $\mathcal{D}$. Then, for every $t \geq r$, it holds that*

$$\mathbb{E}[\max_i \phi(v_i)^+] \leq \phi(t) + \frac{n(1-F(t))^2}{f(t)}.$$

We can use Lemma 1 in order to bound the revenue in the case of regular valuations.

Lemma 2. *Let $\mathbf{b}$ be a Bayes-Nash equilibrium for a GSP auction game with n players with random valuations v_1, ..., v_n drawn i.i.d. from a regular distribution $\mathcal{D}$. Then, there exists $r' \geq r$ such that $\mathcal{R}_{r'}(\mathbf{b}) \geq (1-1/e)\alpha_1\mathbb{E}[\max_i \phi(v_i)^+]$.*

For MHR valuations, we show an improved bound.

Lemma 3. *Let $\mathbf{b}$ be a Bayes-Nash equilibrium for a GSP auction game with n players with random valuations v_1, ..., v_n drawn i.i.d. from an MHR distribution $\mathcal{D}$. Then, there exists $r' \geq r$ such that $\mathcal{R}_{r'}(\mathbf{b}) \geq (1-e^{-2})\alpha_1\mathbb{E}[\max_i \phi(v_i)^+] - (1-e^{-2})\alpha_1 r(1-F^n(r))$.*

Proof. We will assume that $\mathbb{E}[\max_i \phi(v_i)^+] \geq r(1-F^n(r))$ since the lemma holds trivially otherwise. Let t^* be such that $F(t^*) = 1-\eta/n$ where $\eta = 2-(1-1/e)^n$. We will distinguish between two cases depending on whether $t^* \geq r$ or not.

We first consider the case $t^* \geq r$. We will use the definition of the virtual valuation, the fact that the hazard rate function satisfies $h(t^*) \geq h(r) = 1/r$, the definition of t^*, Lemma 1 (with $t = t^*$), and the fact that $F(r) \leq 1-1/e$ which implies that $1-F^n(r) \geq \eta - 1$. We have

$$\begin{aligned}
&t^*(1-F^n(t^*)) \\
&= \phi(t^*)(1-F^n(t^*)) + \frac{1}{h(t^*)}(1-F^n(t^*)) \\
&= \phi(t^*)(1-F^n(t^*)) + \frac{\eta}{h(t^*)}(1-F^n(t^*)) - \frac{\eta-1}{h(t^*)}(1-F^n(t^*)) \\
&\geq \phi(t^*)(1-F^n(t^*)) + \frac{n(1-F(t^*))^2}{f(t^*)} \cdot \frac{\eta(1-F^n(t^*))}{n(1-F(t^*))} - (\eta-1)r(1-F^n(t^*)) \\
&= (1-F^n(t^*))\left(\phi(t^*) + \frac{n(1-F(t^*))^2}{f(t^*)} - (\eta-1)r\right) \\
&\geq \left(1-\left(1-\frac{2-(1-1/e)^n}{n}\right)^n\right)\left(\mathbb{E}[\max_i \phi(v_i)^+] - r(1-F^n(r))\right).
\end{aligned}$$

Note that the left side of the above equality multiplied with α_1 is a lower bound on the revenue of GSP with reserve t^*. Also, $\left(1-\frac{2-(1-1/e)^n}{n}\right)^n$ is non-decreasing in n and approaches e^{-2} from below as n tends to infinity. Furthermore, the right-hand side of the above inequality in non-negative. Hence,

$$\mathcal{R}_{t^*}(\mathbf{b}) \geq (1-e^{-2})\alpha_1 \mathbb{E}[\max_i \phi(v_i)^+] - (1-e^{-2})\alpha_1 r(1-F^n(r))$$

as desired.

We now consider the case $t^* < r$. We have $1-\eta/n = F(t^*) \leq F(r) \leq 1-1/e$ which implies that $n \leq 5$. Tedious calculations yield

$$\frac{1-F^n(r)}{n(1-F(r))} = \frac{1+F(r)+...+F^{n-1}(r)}{n} \geq \frac{1-e^{-2}}{2-e^{-2}}$$

for $n \in \{2,3,4,5\}$ since $F(r) \geq 1-\eta/n$. Hence,

$$\begin{aligned}
\mathcal{R}_r(\mathbf{b}) &\geq \alpha_1 r(1-F^n(r)) \\
&\geq (1-e^{-2})\alpha_1 nr(1-F(r)) - (1-e^{-2})\alpha_1 r(1-F^n(r)) \\
&\geq (1-e^{-2})\alpha_1 \mathbb{E}[\max_i \phi(v_i)^+] - (1-e^{-2})\alpha_1 r(1-F^n(r)),
\end{aligned}$$

where the last inequality follows by applying Lemma 1 with $t = r$. □

4 Deviation Weight Function Families

The main idea we use for the analysis of Bayes-Nash equilibria of auction games with reserve price t is that the utility of player i with valuation $v_i = x \geq t$ does not increase when this player deviates to any other bid in $[t, x]$. This provides us with infinitely many inequalities on the utility of player i that are expressed in terms of her valuation, the bids of the other players, and the reserve price. Our technique combines these infinite lower bounds by considering their weighted average. The specific weights with which we consider the different inequalities are given by families of functions with particular properties that we call deviation weight function families.

Definition 2. *Let $\beta, \gamma, \delta \geq 0$ and consider the family of functions $\mathcal{G} = \{g_\xi : \xi \in [0,1)\}$ where g_ξ is a non-negative function defined in $[\xi, 1]$. $\mathcal{G}$ is a (β,γ,δ)-DWFF (deviation weight function family) if the following two properties hold for every $\xi \in [0,1)$:*

$$\begin{aligned}
&i) \int_\xi^1 g_\xi(y)\,\mathrm{d}y = 1, \\
&ii) \int_z^1 (1-y) g_\xi(y)\,\mathrm{d}y \geq \beta - \gamma z + \delta\xi, \quad \forall z \in [\xi, 1].
\end{aligned}$$

The next lemma is used in order to prove most of our bounds together with the deviation weight function family presented in Lemma 5.

Lemma 4. *Consider a Bayes-Nash equilibrium* $\mathbf{b}$ *for a GSP auction game with* n *players and reserve price* t*. Then, the following two inequalities hold for every player* i*.*

$$\mathbb{E}[u_i(\mathbf{b})] \geq \sum_{j=c}^{n} \mathbb{E}[\alpha_j(\beta v_i - \gamma b_{\pi(j)} + \delta t)\mathbb{1}\{\mu(i) = j\}], \tag{1}$$

$$\mathbb{E}[\alpha_{\sigma(i)}\phi(v_i)] \geq \sum_{j=c}^{n} \mathbb{E}[\alpha_j(\beta\phi(v_i) - \gamma b_{\pi(j)})\mathbb{1}\{\mu(i) = j\}], \tag{2}$$

where c *is any integer in* $[n]$*,* β*,* γ*, and* δ *are such that a* (β, γ, δ)*-DWFF exists, and* μ *is any revenue-maximizing allocation.*

Lemma 5. *Consider the family of functions* $\mathcal{G}_1$ *consisting of the functions* $g_\xi : [\xi, 1] \to \mathbb{R}_+$ *defined as follows for every* $\xi \in [0, 1)$*:*

$$g_\xi(y) = \begin{cases} \frac{\kappa}{1-y}, & y \in [\xi, \xi + (1-\xi)\lambda), \\ 0, & \text{otherwise}, \end{cases}$$

where $\lambda \in (0, 1)$ *and* $\kappa = -\frac{1}{\ln(1-\lambda)}$*. Then,* $\mathcal{G}_1$ *is a* $(\kappa\lambda, \kappa, \kappa(1-\lambda))$*-DWFF.*

We remark that the bound for GSP auctions with Myerson reserve (and players with MHR valuations) follows by a slightly more involved deviation weight function family. Due to lack of space, we omit it from this extended abstract; it will appear in the final version of the paper.

5 Revenue Guarantees in GSP Auctions

We will now exploit the techniques developed in the previous sections in order to prove our bounds for GSP auctions. Throughout this section, we denote by O_j the event that slot j is occupied in the revenue-maximizing allocation considered. The next lemma provides a lower bound on the revenue of GSP auctions.

Lemma 6. *Consider a Bayes-Nash equilibrium* $\mathbf{b}$ *for a GSP auction game with Myerson reserve price* r *and* n *players. It holds that*

$$\sum_{j \geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] \leq \mathcal{R}_r(\mathbf{b}) - \alpha_1 r \cdot \Pr[\mathrm{O}_1].$$

Proof. Consider a Bayes-Nash equilibrium $\mathbf{b}$ for a GSP auction game with Myerson reserve price r. Define $\Pr[\mathrm{O}_{n+1}] = 0$. Consider some player whose valuation exceeds r and is thus allocated some slot. Note that the player's payment per click is determined by the bid of the player allocated just below her, if there is one, otherwise, the player's (per click) payment is set to r. It holds that

$$\mathcal{R}_r(\mathbf{b}) = \sum_j \alpha_j r(\Pr[\mathrm{O}_j] - \Pr[\mathrm{O}_{j+1}]) + \sum_j \mathbb{E}[\alpha_j b_{\pi(j+1)} \mathbb{1}\{\mathrm{O}_{j+1}\}]$$

$$
\begin{aligned}
&= \sum_{j\geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] + \sum_j \alpha_j r(\Pr[\mathrm{O}_j] - \Pr[\mathrm{O}_{j+1}]) \\
&\quad + \sum_j \mathbb{E}[(\alpha_j - \alpha_{j+1}) b_{\pi(j+1)} \mathbb{1}\{\mathrm{O}_{j+1}\}] \\
&\geq \sum_{j\geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] + \sum_j \alpha_j r(\Pr[\mathrm{O}_j] - \Pr[\mathrm{O}_{j+1}]) \\
&\quad + \sum_j (\alpha_j - \alpha_{j+1}) r \cdot \Pr[\mathrm{O}_{j+1}] \\
&= \sum_{j\geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] + \sum_j \alpha_j r \Pr[\mathrm{O}_j] - \sum_j \alpha_{j+1} r \cdot \Pr[\mathrm{O}_{j+1}] \\
&= \sum_{j\geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] + \alpha_1 r \cdot \Pr[\mathrm{O}_1].
\end{aligned}
$$

The proof follows by rearranging the terms in the last inequality. □

The next statement follows by Lemmas 2 and 4 using the DWFF defined in Lemma 5.

Theorem 1. *Consider a regular distribution $\mathcal{D}$. There exists some r^*, such that the revenue guarantee over Bayes-Nash equilibria of GSP auction games with reserve price r^* is 4.72, when valuations are drawn i.i.d. from $\mathcal{D}$.*

Proof. By Lemma 2, we have that there exists $r' \geq r$ such that the expected revenue over any Bayes-Nash equilibrium $\mathbf{b}'$ of the GSP auction game with reserve price r' satisfies

$$
\mathcal{R}_{r'}(\mathbf{b}') \geq (1 - 1/e)\mathbb{E}[\alpha_1 \phi(v_{o(1)})^+]. \tag{3}
$$

Now, let $\mathbf{b}''$ be any Bayes-Nash equilibrium of the GSP auction game with Myerson reserve and let β, γ, and δ be parameters so that a (β,γ,δ)-DWFF exists. Using inequality (2) from Lemma 4 with $c = 2$ and Lemma 6 we obtain

$$
\begin{aligned}
\mathcal{R}_r(\mathbf{b}'') &= \sum_i \mathbb{E}[\alpha_{\sigma(i)} \phi(v_i)] \\
&\geq \sum_i \sum_{j\geq 2} \mathbb{E}[\alpha_j (\beta\phi(v_i) - \gamma b_{\pi(j)}) \mathbb{1}\{\mu(i) = j\}] \\
&= \beta \sum_{j\geq 2} \mathbb{E}[\alpha_j \phi(v_{o(j)})^+] - \gamma \sum_{j\geq 2} \mathbb{E}[\alpha_j b_{\pi(j)} \mathbb{1}\{\mathrm{O}_j\}] \\
&\geq \beta \sum_{j\geq 2} \mathbb{E}[\alpha_j \phi(v_{o(j)})^+] - \gamma \mathcal{R}_r(\mathbf{b}'').
\end{aligned}
$$

In other words,

$$
(1+\gamma)\mathcal{R}_r(\mathbf{b}'') \geq \beta \sum_{j\geq 2} \mathbb{E}[\alpha_j \phi(v_{o(j)})^+].
$$

Using this last inequality together with inequality (3), we obtain

$$\begin{aligned}\left(1+\gamma+\frac{e\beta}{e-1}\right)\max\{\mathcal{R}_r(\mathbf{b}''),\mathcal{R}_{r'}(\mathbf{b}')\} &\geq (1+\gamma)\mathcal{R}_r(\mathbf{b}'')+\frac{e\beta}{e-1}\mathcal{R}_{r'}(\mathbf{b}') \\ &\geq \beta\sum_j \mathbb{E}[\alpha_j\phi(v_{o(j)})^+] \\ &= \beta\mathcal{R}_{OPT}.\end{aligned}$$

We conclude that there exists some reserve price r^* (either r or r') such that for any Bayes-Nash equilibrium $\mathbf{b}$ it holds that

$$\frac{\mathcal{R}_{OPT}}{\mathcal{R}_{r^*}(\mathbf{b})} \leq \frac{1+\gamma}{\beta}+\frac{e}{e-1}.$$

By Lemma 5, the family $\mathcal{G}_1$ is a $(\beta,\gamma,0)$-DWFF with $\beta=\kappa\lambda$ and $\gamma=\kappa$, where $\lambda\in(0,1)$ and $\kappa=-\frac{1}{\ln(1-\lambda)}$. By substituting β and γ with these values and using $\lambda\approx 0.682$, the right-hand side of our last inequality is upper-bounded by 4.72. □

The next statement applies to MHR valuations. It follows by Lemmas 3 and 4 using the DWFF defined in Lemma 5.

Theorem 2. *Consider an MHR distribution $\mathcal{D}$. There exists some r^*, such that the revenue guarantee over Bayes-Nash equilibria of GSP auction games with reserve price r^* is 3.46, when valuations are drawn i.i.d. from $\mathcal{D}$.*

Proof. Let $\mathbf{b}'$ be any Bayes-Nash equilibrium of the GSP auction game with Myerson reserve and let β, γ, and δ be parameters so that a (β,γ,δ)-DWFF exists. Since $\mathcal{D}$ is an MHR probability distribution, we have

$$\mathbb{E}[\alpha_{\sigma(i)}r] \geq \mathbb{E}[\alpha_{\sigma(i)}(v_i-\phi(v_i))] = \mathbb{E}[u_i(\mathbf{b}')]$$

for every player i. By summing over all players and using inequality (1) from Lemma 4 with $c=2$, we obtain

$$\begin{aligned}&\sum_i \mathbb{E}[\alpha_{\sigma(i)}r] \geq \sum_i \mathbb{E}[u_i(\mathbf{b}')] \\ &\geq \sum_i\sum_{j=2}^n \mathbb{E}[\alpha_j(\beta v_i-\gamma b_{\pi(j)}+\delta r)\mathbb{1}\{\mu(i)=j\}] \\ &\geq \sum_{j\geq 2}\mathbb{E}[\alpha_j(\beta\phi(v_{o(j)})^+-\gamma b_{\pi(j)}+\delta r)\mathbb{1}\{\mathrm{O}_j\}] \\ &\geq \beta\sum_{j\geq 2}\mathbb{E}[\alpha_j\phi(v_{o(j)})^+]-\gamma\sum_{j\geq 2}\mathbb{E}[\alpha_j b_{\pi(j)}\mathbb{1}\{\mathrm{O}_j\}]+\delta\sum_{j\geq 2}\mathbb{E}[\alpha_j r\mathbb{1}\{\mathrm{O}_j\}] \\ &\geq \beta\sum_{j\geq 2}\mathbb{E}[\alpha_j\phi(v_{o(j)})^+]-\gamma\mathcal{R}_r(\mathbf{b}')+(\gamma-\delta)\mathbb{E}[\alpha_1 r\mathbb{1}\{O_1\}]+\delta\sum_j\mathbb{E}[\alpha_j r\mathbb{1}\{\mathrm{O}_j\}] \\ &= \beta\sum_{j\geq 2}\mathbb{E}[\alpha_j\phi(v_{o(j)})^+]-\gamma\mathcal{R}_r(\mathbf{b}')+(\gamma-\delta)\mathbb{E}[\alpha_1 r\mathbb{1}\{O_1\}]+\delta\sum_i\mathbb{E}[\alpha_{\mu(i)}r].\end{aligned}$$

The last inequality follows by Lemma 6. Since $\sum_i \mathbb{E}[\alpha_{\mu(i)} r] = \sum_i \mathbb{E}[\alpha_{\sigma(i)} r]$, we obtain that

$$\gamma \mathcal{R}_r(\mathbf{b}') \geq \beta \sum_{j \geq 2} \mathbb{E}[\alpha_j \phi(v_{o(j)})^+] + (\gamma - \delta)\alpha_1 r \cdot \Pr[O_1] + (\delta - 1) \sum_i \mathbb{E}[\alpha_{\sigma(i)} r]. \tag{4}$$

By Lemma 3, we have that there exists $r' \geq r$ such that the expected revenue over any Bayes-Nash equilibrium $\mathbf{b}''$ of the GSP auction game with reserve price r' satisfies

$$\mathcal{R}_{r'}(\mathbf{b}'') \geq (1 - e^{-2})\mathbb{E}[\alpha_1 \phi(v_{o(1)})^+] - (1 - e^{-2})\mathbb{E}[\alpha_1 r \mathbb{1}\{O_1\}].$$

Using this last inequality together with inequality (4), we obtain

$$\begin{aligned}
&\left(\gamma + \frac{e^2 \beta}{e^2 - 1}\right) \max\{\mathcal{R}_r(\mathbf{b}'), \mathcal{R}_{r'}(\mathbf{b}'')\} \\
&\geq \gamma \mathcal{R}_r(\mathbf{b}') + \frac{e^2 \beta}{e^2 - 1} \mathcal{R}_{r'}(\mathbf{b}'') \\
&\geq \beta \sum_j \mathbb{E}[\alpha_j \phi(v_{o(j)})^+] + (\gamma - \delta - \beta)\mathbb{E}[\alpha_1 r \mathbb{1}\{O_1\}] + (\delta - 1) \sum_i \mathbb{E}[\alpha_{\sigma(i)} r] \\
&\geq \beta \mathcal{R}_{OPT} + (\gamma - \delta - \beta)\mathbb{E}[\alpha_1 r \mathbb{1}\{O_1\}] + (\delta - 1) \sum_i \mathbb{E}[\alpha_{\sigma(i)} r].
\end{aligned}$$

By Lemma 5, the family $\mathcal{G}_1$ is a (β, γ, δ)-DWFF with $\beta = \gamma - \delta = \kappa\lambda$, $\gamma = \kappa$, and $\delta = \kappa(1 - \lambda)$, where $\lambda \in (0, 1)$ and $\kappa = -\frac{1}{\ln(1-\lambda)}$. By setting $\lambda \approx 0.432$ so that $\delta = \kappa(1 - \lambda) = 1$, the above inequality implies that there exists some reserve price r^* (either r or r') such that for any Bayes-Nash equilibrium $\mathbf{b}$ of the corresponding GSP auction game, it holds that

$$\frac{\mathcal{R}_{OPT}}{\mathcal{R}_{r^*}(\mathbf{b})} \leq \frac{1}{\lambda} + \frac{e^2}{e^2 - 1} \approx 3.46,$$

as desired. □

For GSP auctions with Myerson reserve, our revenue bound follows using a slightly more involved deviation weight function family.

Theorem 3. *Consider an MHR distribution $\mathcal{D}$. The revenue guarantee over Bayes-Nash equilibria of GSP auction games with Myerson reserve price r is* 3.90*, when valuations are drawn i.i.d. from $\mathcal{D}$.*

6 Conclusions

Even though we have significantly improved the results of [14], we conjecture that our revenue guarantees could be further improved. The work of Gomes and Sweeney [7] implies that the revenue guarantee of GSP auctions with Myerson reserve is in general higher than 1; however, no explicit lower bound is known. Due to the difficulty in computing Bayes-Nash equilibria analytically, coming up with a concrete lower bound construction is interesting and would reveal the gap of our revenue guarantees.

References

1. R. Barlow and R. Marshall. Bounds for distributions with monotone hazard rate. *Annals of Mathematicals Statistics*, 35(3): 1234–1257, 1964.
2. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pp. 81–90, 2011.
3. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. *arXiv:1201.6429*, 2012.
4. S. Chawla, J. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 41th ACM Symposium on Theory of Computing (STOC)*, pp. 311–320, 2010.
5. E.H. Clarke. Multipart pricing of public goods. *Public Choice*, 11: 17–33, 1971.
6. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertizing and the generalized second-price auction: selling billions of dollars worth of keywords. *The American Economic Review*, 97(1): 242-259, 2007.
7. R. Gomes and K. Sweeney. Bayes-Nash equilibria of the generalized second price auction. Working paper, 2011. Preliminary version in *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pp. 107–108, 2009.
8. T. Groves. Incentives in teams. *Econometrica*, 41(4): 617–631, 1973.
9. M. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 58–65, 2007.
10. U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4): 745–747, 1977.
11. V. Krishna *Auction Theory.* Academic Press, 2002.
12. S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*, pp. 218–227, 2006.
13. B. Lucier and R. Paes Leme. GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pp. 71–80, 2011.
14. B. Lucier, R. Paes Leme, and É. Tardos. On revenue in generalized second price auctions. In *Proceedings of the 21st World Wide Web Conference (WWW)*, pp. 361–370, 2012.
15. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.
16. M. Ostrovsky and M. Schwarz. Reserve prices in Internet advertising auctions: a field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pp. 59–60, 2011.
17. R. Paes Leme and É. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 735–744, 2010.
18. H. Varian. Position auctions. *International Journal of Industrial Organization*, 25: 1163–1178, 2007.
19. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1): 8–37, 1961.