

On Revenue in the Generalized Second Price Auction ^{*}

Brendan Lucier [†]
 Microsoft Research
 New England
 brlucier@microsoft.com

Renato Paes Leme [‡]
 Dept of Computer Science
 Cornell University
 renatopl@cs.cornell.edu

Éva Tardos [§]
 Dept of Computer Science
 Cornell University
 eva@cs.cornell.edu

ABSTRACT

The Generalized Second Price (GSP) auction is the primary auction used for selling sponsored search advertisements. In this paper we consider the revenue of this auction at equilibrium. We prove that if agent values are drawn from identical regular distributions, then the GSP auction paired with an appropriate reserve price generates a constant fraction (1/6th) of the optimal revenue.

In the full-information game, we show that at any Nash equilibrium of the GSP auction obtains at least half of the revenue of the VCG mechanism excluding the payment of a single participant. This bound holds also with any reserve price, and is tight.

Finally, we consider the tradeoff between maximizing revenue and social welfare. We introduce a natural convexity assumption on the click-through rates and show that it implies that the revenue-maximizing equilibrium of GSP in the full information model will necessarily be envy-free. In particular, it is always possible to maximize revenue and social welfare simultaneously when click-through rates are convex. Without this convexity assumption, however, we demonstrate that revenue may be maximized at a non-envy-free equilibrium that generates a socially inefficient allocation.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms

games, economics, algorithms

Keywords

sponsored search auctions, revenue

^{*}a preliminary version of this work was presented in the AdAuctions workshop 2011.

[†]This work was done in part at the University of Toronto.

[‡]Supported in part by NSF grants CCF-0729006.

[§]Supported in part by NSF grants CCF-0910940 and CCF-0729006, ONR grant N00014-98-1-0589, a Yahoo! Research Alliance Grant, and a Google Research Grant.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2012, April 16–20, 2012, Lyon, France.
 ACM 978-1-4503-1229-5/12/04.

1. INTRODUCTION

The sale of sponsored search advertising space is a primary source of income for Internet companies, and responsible for billions of dollars in annual advertising revenue [7]. The Generalized Second Price (GSP) auction is the premier method by which sponsored search advertising space is sold; it is currently employed by Google, Bing, and Yahoo. However, use of the GSP auction is not universal: the classical VCG mechanism was adopted by Facebook for its AdAuction system [12]. In fact, Google also considered switching its advertising platform to a VCG auction some years ago, but eventually decided against it [24]. This apparent tension underlines the importance of understanding how these mechanisms compare. There are many factors in comparing possible mechanisms: The welfare of three distinct user groups (the experience of the searchers, the welfare of advertisers, and the revenue of the auction) are all important considerations, as well as the simplicity of the auction design. In this paper, we take the point of view of the seller and compare the revenue properties of the GSP and VCG auctions.

Previous studies of the revenue of GSP have largely focused on outcomes of the full information game, restricted to the subclass of envy-free equilibria [7, 26, 9]. Here we consider revenue in equilibria of the Bayesian partial information version of the game, as well as equilibria of the full information game that are not envy-free, and are interested to show that GSP generates close to as much revenue as the classical optimal VCG auction. This comparison is natural, as in the Bayesian model the VCG auction is revenue optimal with the appropriate reserve price [21].

Let us first briefly describe the model introduced by Edelman et al [7] and Varian [25]. In sponsored search, a user makes a query for certain keywords in a search engine and is presented with a list of relevant advertisements in addition to organic search results. We assume a “pay-per-click” pricing model, in which the advertiser pays a fee to the search provider whenever a user clicks on an advertisement. There are multiple positions (or “slots”) in which an ad may appear, and the probability that a user clicks an ad depends on its slot. This is modeled as a click-through rate (CTR), a probability α associated with each slot, the probability of getting a click for an advertisement in that position. The search engine must therefore determine which ads to place where, and determine a price per click for each slot. This is done via an auction in which advertisers make bids, which are viewed as the advertiser’s maximum willingness to pay per click. We note that this simplified model as a one-shot

game does not account for advertiser budgets, so models the case when budgets are large. Also, for simplicity of presentation, we will assume that all ads have the same quality (i.e. click-through rate does not depend on the advertiser selected for a slot), though our results for the full information game extend to the version of the model with separable click-through rates.

The VCG and GSP mechanisms differ in the way in which the aforementioned auction is resolved. In both auctions, advertisers are assigned slots in order of their bids, with higher bidders receiving slots with higher click probabilities. The two auctions differ in their payment schemes. In VCG, each agent pays an amount equal to his externality on the other agents: the decrease in the total welfare of all other agents caused by the presence of this advertiser. By contrast, in GSP each advertiser simply pays a price per click equal to the next highest bid. The VCG auction has the strong property of being truthful in dominant strategies. The GSP auction is not truthful, and is therefore prone to strategic bidding behavior. Indeed, strategic manipulation of bids is well-documented in historical GSP bidding data [6].

Since the VCG mechanism is truthful, the revenue of VCG is simply the revenue generated when all bidders declare their values truthfully. If bidders declare their values truthfully in a GSP auction, GSP generates strictly more revenue than VCG. However, rational agents may not declare their values truthfully when participating in a GSP auction. Thus, when studying the revenue of GSP, we consider the revenue generated at a Nash equilibrium; that is, a profile of bidding strategies such that no advertiser can improve his utility (or expected utility in the Bayesian case) by unilaterally deviating. Our goal, then, is a comparison between the revenue of the VCG auction and the revenue of GSP at equilibrium.

Results.

In Section 3 we consider the Bayesian version of this game when valuations are drawn from identical and independent distributions that satisfy the regularity condition. We show that if we allow the auctioneer to include reserve prices the GSP auction always obtains a constant fraction (1/6th) of the optimal VCG revenue, in expectation. The Bayesian setting is motivated by player uncertainty: for many keywords, the ability of a player to exactly predict his opponents types is impaired. Due to complicating factors such as the budgets, quality scores (which depend on many characteristics of each query, such as origin, time, and search history of the user), and the underlying ad allocation algorithm, each auction is different (even those triggered by the same search term) [23]. The resulting uncertainty is captured by equilibria in a Bayesian partial information model, where bidders can only evaluate the expected welfare of the equilibrium, with expectation taken over the valuations of other bidders. It is well-known that, in this setting the revenue-optimal truthful mechanism is the VCG auction with Myerson's reserve price [21].

One might also wish to bound the revenue of GSP with respect to the revenue of VCG without reserve prices, but we show that this is not possible: there are cases in which the VCG revenue is unboundedly greater than the GSP revenue. However, if the slot CTRs satisfy a certain well-separatedness condition - namely that the click through rates

of adjacent slots differ by at least a certain constant factor - then we prove that GSP always obtains a constant fraction of the VCG revenue even in settings of partial information, extending a result of Lahaie [15] who considered welfare under this assumption on the CTRs. Our result holds even if we do not assume that agents avoid dominated strategies, as long as there are at least three participants in the auction.

In Sections 4 and 5 we consider the full information game. When auctions with the same participants, valuations, and quality scores are repeated many times each day, a stable outcome can be modeled by a full-information Nash equilibrium of the one-shot game. We prove that at any Nash equilibrium, the revenue generated by GSP is at least half of the VCG revenue, *excluding the single largest payment of a bidder*. Thus, as long as the VCG revenue is not concentrated on the payment of a single participant, the worst-case GSP revenue approximates the VCG revenue to within a constant factor. This result also holds with an arbitrary reserve price. We also provide an example illustrating that the factor of 2 in our analysis is tight, and the revenue of GSP at equilibrium may be arbitrarily less than the full revenue of VCG (without excluding a bidder).

In Section 5 we analyze the tradeoffs of the maximum revenue attainable by the full information GSP mechanism under different equilibrium notions. We demonstrate that there can exist inefficient, non-envy-free equilibria that obtain greater revenue than any envy-free equilibrium. However, we prove that if CTRs are *convex*, meaning that the marginal increase in CTR is monotone in slot position, then the optimal revenue always occurs at an envy-free equilibrium. This implies that when click-through rates are convex, the GSP auction optimizes revenue at an equilibrium that simultaneously maximizes the social welfare. The convexity assumption we introduce is quite natural and may be of independent interest. Note that this assumption is satisfied in the case when CTRs degrade by a constant factor from one slot to the next.

Related Work.

There has been considerable amount of work on the economic and algorithmic issues behind sponsored search auctions - see the survey of Lahaie et al [16] for an overview of the early work and the survey of Maille et al [19] for recent developments. The GSP model we adopt is due to Edelman et al [7], Varian [25] and Aggarwal et al [2]. Much of the previous work on the GSP auction considered social welfare properties of equilibria. Edelman et al [7] and Varian [25] define a subclass of Nash equilibria called envy-free equilibria and show that such equilibria always exist and are socially optimal, therefore showing that the price of stability is 1 for the full information game. Paes Leme and Tardos [22] showed a bound of 1.618 on the price of anarchy, which was recently improved to 1.282 by Caragiannis et al [4]. The Bayesian version of this game, when valuations are random and only the distribution is public knowledge, was first studied by Gomes and Sweeney [10] who showed that efficient equilibria may not exist in this setting. Paes Leme and Tardos [22] prove a price of anarchy of 8, which was recently improved by Lucier and Paes Leme [18] and Caragiannis et al [4] to 3.03.

Considering revenue properties of GSP, Edelman et al [7], and Varian [25] show that envy free equilibria have revenue at least as good as the revenue of VCG. Both [7] and [25]

present informal arguments to support the equilibrium selection for this class of equilibria, but there is no strong theoretical model that explains this selection [3, 9]. Further, the notion of envy-free equilibria applies only in the full information game. Varian [26] shows how to compute the revenue optimal envy-free Nash equilibrium, although his model allows agents to overbid (which is dominated strategy, and we consider it unnatural). We consider the question of maximum revenue equilibria without the assumption of envy-free outcomes, and show that in general inefficient equilibria can generate more revenue than efficient ones, but this is no longer the case under a natural convexity assumption.

Gomes and Sweeney [10] study GSP as a Bayesian game, and show that symmetric efficient equilibria may not exist in the Bayesian setting. They analyze the influence of click-through-rates on the revenue in equilibrium and observe the counter-intuitive phenomenon by which revenue decreases when click-through-rates increase.

A common tool for increasing revenue in settings of partial information is to apply *reserve prices*, where bids are rejected unless they meet some minimum bid r . When bidders' values are drawn from identically distributed satisfying the regularity condition¹, the revenue-optimal truthful auction for sponsored search is the VCG auction with an appropriate reserve price [21]. Edelman and Schwarz [8] show that in GSP auctions reserve prices have a surprisingly large effect on revenue.

Edelman and Schwarz [9] model the repeated auctions for a keyword as a repeated game, and show, using Myerson's optimal auction [21], that if valuations are drawn from an iid distribution then the Nash equilibria that arise as a stable limit of rational play in this repeated game cannot have revenue more than the optimal auction: VCG with an appropriately chosen reserve price. However, they don't consider whether GSP may generate revenue much less than the VCG auction, which is the main question we consider. Also, unlike Edelman and Schwarz [9], we do not use the repeated nature of this game for arguing for certain equilibrium selection: rather, we consider *all* stable outcomes of the auction, not only those that arise as limits of rational repeated play.

Finally, we mention that the well-known *revenue equivalence theorem*, which provides conditions under which alternative mechanisms generate the same revenue at equilibrium, does not apply in our settings. Revenue equivalence requires that agents have values drawn from identical distributions and that the mechanisms generate the same outcome. As a result, this equivalence does not apply in the full information setting, and cannot be used to compare inefficient equilibria of GSP to the VCG revenue.

2. PRELIMINARIES

An AdAuctions instance is composed of n players and n slots. In the full information model, each player has a value v_i for each click he gets, and each slot j has click-through-rate α_j . That means that if player i is allocated in slot j , he gets α_j clicks in expectation, where $\alpha_1 \geq \dots \geq \alpha_n$. In the full information setting, we will assume players are numbered such that $v_1 \geq v_2 \geq \dots \geq v_n$. Let $\alpha = (\alpha_1, \dots, \alpha_n)$ be the CTR vector and $\mathbf{v} = (v_1, \dots, v_n)$ be the type vector.

¹Many common distributions are regular, including all uniform, normal, and exponential distributions.

A mechanism for the AdAuctions problem is as follows: it begins by eliciting bids b_i from the players, which can be thought of as declared valuations. We write $\mathbf{b} = (b_1, \dots, b_n)$ for the bid vector. Using the \mathbf{b} and α , the mechanism chooses an allocation $\pi : [n] \rightarrow [n]$ which means that player $\pi(j)$ is allocated to slot j , and a price vector $\mathbf{p} = (p_1, \dots, p_n)$, where p_i is the price that player i pays per click. Player i then experiences utility $u_i(\mathbf{b}) = \alpha_{\sigma(i)}(v_i - p_i)$, where $\sigma(i) = \pi^{-1}(i)$.

The social welfare generated by the mechanism is given by $SW(\mathbf{v}, \pi) = \sum_i \alpha_i v_{\pi(i)}$ and the revenue is given by $\mathcal{R}(\mathbf{b}) = \sum_i \alpha_{\sigma(i)} p_i$. We focus on two mechanisms: GSP and VCG. In both mechanisms, the players are ordered by their bids, i.e., $\pi(j)$ is the player with the j^{th} largest bid, but they differ in the payments charged. GSP mimics the single-item second price auction by charging each player the bid of the next highest bidder, i.e. $p_i = b_{\pi(\sigma(i)+1)}$ if $\sigma(i) < n$ and zero otherwise. VCG charges each player the externality it imposes on the other players, which is $p_i^{\text{VCG}} = \frac{1}{\alpha_{\sigma(i)}} \sum_{j=\sigma(i)+1}^n (\alpha_{j-1} - \alpha_j) b_{\pi(j)}$.

VCG is a truthful mechanism: regardless of what the other players are doing, it is a weakly dominant strategy for player i to report his true valuation. The resulting outcome of VCG is therefore social-welfare optimal and the revenue is

$$\mathcal{R}^{\text{VCG}}(\mathbf{v}) = \sum_i \sum_{j>i} (\alpha_{j-1} - \alpha_j) v_j = \sum_{i=2}^n (i-1) (\alpha_{i-1} - \alpha_i) v_i.$$

The GSP auction, however, is not truthful. Thus, for GSP, we are interested in the set of bid profiles that constitute a Nash equilibrium, i.e. such that

$$u_i(b_i, \mathbf{b}_{-i}) \geq u_i(b'_i, \mathbf{b}_{-i}), \forall b'_i \in [0, v_i].$$

We will assume that players do not overbid (i.e. that $b_i \leq v_i$) since bidding more than one's true value is a weakly dominated strategy [22].

We say that an equilibrium is **efficient** if it maximizes social welfare, i.e. $\pi(i) = i, \forall i$ in the full information version.

We will also consider the comparison between VCG and GSP in the presence of a reserve price. Let VCG_r be the VCG mechanism with reserve price r , where we discard all players with bids smaller than r and run the VCG mechanism on the remaining players, who then pay price per click $\max\{p_i, r\}$. In the analogous variant of the GSP mechanism, which we call GSP with reserve price r (GSP_r), we also discard all players with bids smaller than r , the remaining players are allocated using GSP, and the last player to be allocated pays price r per click.

Below, we represent the special classes of equilibria that have been studied in the literature, which we call **equilibrium hierarchy** for GSP. We define and discuss them in Section 5 :

$$\left\{ \begin{array}{c} \text{VCG} \\ \text{outcome} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{envy-free} \\ \text{equilibria} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{efficient} \\ \text{Nash eq} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{all} \\ \text{Nash} \end{array} \right\}$$

2.1 Bayesian setting

In a Bayesian setting, each player knows her own valuation but only knows a distribution on the other players' valuations. In this model the values of the players are random variables, with type vector \mathbf{v} drawn from a known distribution F . After learning his own value v_i , a player chooses a bid $b_i(v_i)$ to play in the AdAuctions game. The strategies

are therefore bidding functions $b_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, and we will continue to assume that players do not overbid, i.e. $b_i(v) \leq v$ (again, since overbidding is weakly dominated). A set of bidding functions is a **Bayes-Nash equilibrium** if, for all i , v_i , and b'_i ,

$$\mathbb{E}[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))|v_i] \geq \mathbb{E}[u_i(b'_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))|v_i].$$

A useful tool for studying revenue in the Bayesian setting is Myerson’s Lemma, which can be rephrased in the AdAuctions setting as follows. Given a distribution F over agent values, the **virtual valuation** function is defined by $\phi(x) = x - \frac{1-F(x)}{f(x)}$.

Lemma 1 (Myerson [21]) *At any Bayes-Nash equilibrium of an AdAuction mechanism, we have that, for all i , $\mathbb{E}[\alpha_{\sigma(i)} p_i] = \mathbb{E}[\alpha_{\sigma(i)} \phi(v_i)]$ where p_i is the payment per click of player i and $\alpha_{\sigma(i)}$ is the number of clicks received by agent i , and expectation is with respect to $\mathbf{v} \sim \mathbf{F}$.*

We say that a distribution is **regular** if $\phi(x)$ is a monotone non-decreasing function. For regular distributions, the revenue-optimal mechanism for AdAuctions corresponds to running VCG with Myerson’s reserve price r , which is the largest value such that $\phi(r) = 0$. We will refer to this as *Myerson’s mechanism, VCG_r* .

Running GSP (or VCG) with **reserve price** r means not allocating any user with value $v_i < r$ and running GSP (or VCG) with the remaining agents. For the allocated agents, the mechanism charges per click the maximum between the GSP (VCG) price and r .

A special class of regular distributions is the **monotone hazard rate** distributions (**MHR**), which are the distributions for which $f(x)/(1 - F(x))$ is non-decreasing.

3. REVENUE IN THE BAYESIAN SETTING

In this section we consider the revenue properties of GSP at Bayes-Nash equilibrium. We prove that if agent values are drawn iid from a regular distribution and GSP is paired with an appropriate reserve price, the revenue generated at equilibrium will be within a constant factor of the VCG revenue with optimal reserve, the revenue-optimal mechanism over all Bayes-Nash implementations. So our result implies that GSP revenue is within a constant factor of the optimal. We will first consider the special case of MHR distributions, then prove our result in the more general setting where values are drawn from regular distributions.

We start by showing that the use of reserve prices is crucial: there are instances in which the GSP auction without reserve generates no revenue, whereas the VCG auction generates positive revenue. We do note, however, that one can bound the revenue of GSP without reserve prices when click-through rates are *well-separated*, meaning that there exists a δ such that $\alpha_{i+1} \leq \delta \alpha_i$ for all i . We present these bounds in Appendix A of the full version [17].

3.1 Revenue without Reserves: Bad Examples

We start by providing an example in the Bayesian setting where VCG generates positive revenue and GSP has a Bayes-Nash equilibrium that generates zero revenue. Consider three players with iid valuations drawn uniformly from $[1, 2]$ and three slots with $\alpha = [1, 0.5, 0.5]$. Let $v^{(i)}$ be the i^{th} largest valuation (which is naturally a random variable

defined by \mathbf{v}). We have

$$\mathbb{E}[\mathcal{R}^{VCG}(\mathbf{v})] = \mathbb{E}[0.5v^{(2)}] = \frac{3}{4}.$$

Now, consider the following equilibrium of GSP: $b_i(v_i) = 0$ for $i = 2, 3$ and $b_1(v_1) = v_1$. Clearly player 1 is in equilibrium. To see that players $i = 2, 3$ are in equilibrium, suppose player i has valuation $v_i > 0$. Then his expected utility when bidding any value in $[0, 1]$ is $0.5v_i$, whereas if he changed his bid to some $b > 1$ his utility would be

$$\begin{aligned} \mathbb{E}[u_i(b', b_{-i})|v_i] &= 0.5v_i + 0.5v_i \mathbb{P}(v_1 \leq b') - \int_0^{b'} v_1 d\mathbb{P}(v_1) = \\ &= 0.5v_i + 0.5v_i(b' - 1) - \frac{(b')^2 - 1}{2} \leq \\ &\leq 0.5v_i. \end{aligned}$$

Thus agent i cannot increase his expected utility by placing a non-zero bid.

3.2 Warmup: MHR Valuations

We now show that if valuations are drawn from a MHR distribution and GSP is paired with the Myerson reserve price, the resulting mechanism extracts a constant fraction of the optimal revenue.

In what follows we will write x^+ to denote $\max\{x, 0\}$.

Theorem 2 *If valuations are drawn iid from a MHR distribution F and r is the Myerson reserve price for F , then the expected revenue of GSP_r at any Bayes-Nash equilibrium is at least $\frac{1}{6}$ of the optimal revenue.*

Our proof will make use of the fact that, for MHR distributions, $\phi(x) \geq x - r$ for any $x \geq r$. To see this, note that $x - \phi(x) = \frac{1-F(x)}{f(x)} \leq \frac{1-F(r)}{f(r)} = r$ by monotonicity and the definition of Myerson’s reserve price.

PROOF. Let \mathbf{b} be a Bayes-Nash equilibrium of GSP_r , and let $\mathcal{R}_r(\mathbf{v})$ be the expected revenue of GSP_r at this equilibrium. Let $\mathcal{R}_r^{VCG}(\mathbf{v})$ be the VCG_r revenue. Let random variable $\mu(i)$ denote the slot occupied by player i in the optimal (i.e. efficient) allocation. By Myerson’s Lemma, $\mathbb{E}[\mathcal{R}_r^{VCG}(\mathbf{v})] = \mathbb{E}[\sum_i \alpha_{\mu(i)} \phi(v_i)^+]$. For each player i , let E_1^i denote the event that $b_{\pi(\mu(i))} < v_i/2$, and let E_2^i denote the event that $b_{\pi(\mu(i))} \geq v_i/2$. We will consider each of these events in turn. For the first event, we’ll show that player i contributes to the revenue at least $1/2$ his contribution in the optimum. Consider a player i with value v_i . We have

$$\begin{aligned} \mathbb{E}_{v_{-i}} \left[\alpha_{\mu(i)} \frac{v_i}{2} \mathbb{1}\{E_1^i\} \right] &\leq \mathbb{E}_{v_{-i}} \left[u_i \left(\frac{v_i}{2}, \mathbf{b}_{-i} \right) \right] \\ &\leq \mathbb{E}_{v_{-i}} [u_i(\mathbf{b})] \leq \mathbb{E}_{v_{-i}} [\alpha_{\sigma(i)} v_i] \end{aligned}$$

where the first inequality is due to the definition of E_1^i implying that a bid of $v_i/2$ would win slot $\mu(i)$ (or better) at price no more than $v_i/2$; the second follows since \mathbf{b} is a Bayes-Nash equilibrium, and the third comes from the definition of utility. Notice that all the expectations are taken over v_{-i} and v_i is a constant, so we can divide by v_i , multiply by $\phi(v_i)^+$, take expectations over v_i and sum over all players i to obtain

$$\begin{aligned} \sum_i \mathbb{E}_{\mathbf{v}} [\alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_1^i\}] &\leq 2 \sum_i \mathbb{E}_{\mathbf{v}} [\alpha_{\sigma(i)} \phi(v_i)^+] \\ &= 2\mathbb{E}_{\mathbf{v}} [\mathcal{R}_r(\mathbf{v})]. \end{aligned}$$

For the second event, consider again a player i with value v_i . We will show that the player who gets slot $\mu(i)$ contributes to the revenue. We have

$$\begin{aligned} \mathbb{E}_{v_{-i}} \left[\alpha_{\mu(i)} \frac{\phi(v_i)^+}{2} \mathbb{1}\{E_2^i\} \right] &\leq \mathbb{E}_{v_{-i}} \left[\alpha_{\mu(i)} \frac{v_i}{2} \mathbb{1}\{E_2^i\} \right] \\ &\leq \mathbb{E}_{v_{-i}} [\alpha_{\mu(i)} v_{\pi(\mu(i))}] \\ &\leq \mathbb{E}_{v_{-i}} [\alpha_{\mu(i)} (r + \phi(v_{\pi(\mu(i))})^+)] \end{aligned}$$

where we used the fact that $x \geq \phi(x)^+ \geq x - r$ for all x . Taking expectations over v_i , summing over all players, and noting that event E_2^i implies that $v_{\pi(\mu(i))} \geq r$, we obtain

$$\begin{aligned} \sum_i \mathbb{E}_{\mathbf{v}} [\alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_2^i\}] &\leq \\ &\leq 2 \sum_i \mathbb{E}_{\mathbf{v}} [\alpha_{\sigma(i)} \phi(v_i)^+] + 2 \sum_i \alpha_{\sigma(i)} r \mathbb{1}\{v_i \geq r\}. \end{aligned}$$

Since GSP_r extracts a revenue of at least r per click from every bidder with $v_i > r$, we have $\mathbb{E}_{\mathbf{v}} [\mathcal{R}_r] \geq \mathbb{E}_{\mathbf{v}} [\sum_i \alpha_{\sigma(i)} r \mathbb{1}\{v_i \geq r\}]$. We conclude that $\sum_i \mathbb{E}_{v_{-i}} [\alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_2^i\}] \leq 4\mathbb{E}[\mathcal{R}_r]$. Combining our analysis for the two events, we have

$$\begin{aligned} \mathbb{E}[\mathcal{R}_r^{VCG}(\mathbf{v})] &= \mathbb{E} \left[\sum_i \alpha_{\mu(i)} \phi(v_i)^+ (\mathbb{1}\{E_1^i\} + \mathbb{1}\{E_2^i\}) \right] \\ &\leq 2\mathbb{E}[\mathcal{R}_r(\mathbf{v})] + 4\mathbb{E}[\mathcal{R}_r(\mathbf{v})] = 6\mathbb{E}[\mathcal{R}_r(\mathbf{v})]. \end{aligned}$$

3.3 Regular valuations

We now show that if player valuations are drawn from a regular distribution, then there exists an r' such that running GSP with reserve r' extracts a constant fraction of the optimal revenue. The bound for the MHR bounding the contribution of the player at slot $\mu(i)$ took advantage of the fact that in a MHR distribution $\phi(x) \geq x - r$, which may not be true in a regular distribution. Instead, we will use that the player in slot $\mu(i) - 1$ pays at least the bid in slot $\mu(i)$. This leaves us with the added difficulty in bounding the revenue generated by the first slot. To address this issue, we make use of the well-studied Prophet Inequalities [13, 14, 11].

A simplified version of the Prophet Inequality is as follows. Suppose z_i are independent non-negative random variables. Given any $t \geq 0$, write y_t for the value of the first z_i (by index) satisfying $z_i > t$ (or 0 if there is no such z_i). Then the Prophet Inequality states that there exists some $t \geq 0$ such that $\mathbb{E}[y_t] \geq \frac{1}{2} \mathbb{E}[\max_i z_i]$. Since the proof is of this fact is very short, we include it for completeness in Appendix B in the full version [17].

As has been noted elsewhere [5], the Prophet Inequality has immediate consequences for the revenue of auctions with anonymous reserve prices. The following lemma encapsulates the observation we require.

Lemma 3 *If v_i are drawn iid from a regular distribution then there exists $r_2 \geq 0$ such that, writing Z for the event that $\max_i v_i \geq r_2$, $\mathbb{E}[\max_i \phi(v_i)^+ | Z] \mathbb{P}(Z) \geq \frac{1}{2} \mathbb{E}[\max_i \phi(v_i)^+]$.*

PROOF. (sketch) This follows by applying the Prophet Inequality to virtual values $z_i = \phi(v_i)$ and noting that regularity implies that $v_i \geq r_2$ iff $\phi(v_i) \geq \phi(r_2)$. ■

It is important to remark that the proof of the Prophet Inequality is constructive. If one is able to efficiently compute

$\mathbb{E}[(v_i - t)^+]$ for every t , then we can compute r_2 exactly using binary search. We refer to the appendix in the full version [17] for details.

Our approach will now be to analyze the revenue of GSP under two different reserve prices. An argument similar to Theorem 2 shows that GSP with Myerson reserve obtains a constant fraction of the optimal revenue for all slots other than the first slot. On the other hand, GSP with reserve r_2 from Lemma 3 will obtain at least half of the optimal revenue generated by the first slot. One of these two reserve prices must therefore generate a constant fraction of the optimal revenue.

Theorem 4 *If valuations v_i are drawn iid from a regular distribution F , then there is a reserve price r such that the expected revenue of GSP_r at any Bayes-Nash equilibrium is at least $\frac{1}{6}$ of the optimal revenue.*

PROOF. Define $\mathcal{R}_r^{VCG}(\mathbf{v})$, $\mathcal{R}_r(\mathbf{v})$, and $\mu(i)$ as in Theorem 2. Let r_1 denote the Myerson reserve price for F . By Myerson's Lemma, $\mathbb{E}[\mathcal{R}_r^{VCG}(\mathbf{v})] = \mathbb{E}[\sum_i \alpha_{\mu(i)} \phi(v_i)^+]$. For each player i , we define the following three events:

- $E_1^i = \{ b_{\pi(\mu(i))} < v_i/2 \text{ and } \mu(i) \neq 1 \}$
- $E_2^i = \{ b_{\pi(\mu(i))} \geq v_i/2 \text{ and } \mu(i) \neq 1 \}$
- $E_3^i = \{ \mu(i) = 1 \}$

We wish to bound the virtual value of the optimal allocation, conditioning on each of these events in turn. For the first event, we proceed precisely as in Theorem 2 to obtain

$$\sum_i \mathbb{E}_{v_{-i}} [\alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_1^i\}] \leq 2\mathbb{E}[\mathcal{R}_{r_1}].$$

For the second event, we use the revenue from slot $\mu(1) - 1$. Let random variable p_i denote the payment per click of the player in slot i . Then for all \mathbf{v} ,

$$\alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_2^i\} \leq \alpha_{\mu(i)} v_i \mathbb{1}\{E_2^i\} \leq 2\alpha_{\mu(i)-1} p_{\mu(i)-1} \mathbb{1}\{E_2^i\}$$

where the second inequality follows since E_2^i implies $p_{\mu(i)-1} = b_{\mu(i)} \geq v_i/2$. Therefore, summing over all agents i and taking expectations, we get

$$\mathbb{E}_{\mathbf{v}} \left[\sum_i \alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_2^i\} \right] \leq 2\mathbb{E}_{\mathbf{v}} \left[\sum_i \alpha_i p_i \right] = 2\mathbb{E}[\mathcal{R}_{r_1}].$$

Finally, for event E_3^i , consider setting the reserve price to be r_2 from the statement of Lemma 3 (with distribution F). Note that

$$\mathbb{E} \left[\sum_i \alpha_{\mu(i)} \phi(v_i)^+ \mathbb{1}\{E_3^i\} \right] = \alpha_1 \mathbb{E}[\max_i \phi(v_i)^+].$$

On the other hand, setting reserve price r_2 for GSP we get

$$\begin{aligned} \mathbb{E}[\mathcal{R}_{r_2}] &\geq \alpha_1 \mathbb{E}[\max_i \phi(v_i)^+ | \max_i v_i \geq r_2] \mathbb{P}(\max_i v_i \geq r_2) \\ &\geq \frac{1}{2} \alpha_1 \mathbb{E}[\max_i \phi(v_i)^+] \end{aligned}$$

where the first inequality follows by considering only the expected virtual value due to the first slot and the last inequality follows from Lemma 3. Combining our analysis for each of the three cases, we have

$$\begin{aligned} \mathbb{E}[\mathcal{R}_r^{VCG}] &= \mathbb{E} \left[\sum_i \alpha_{\mu(i)} \phi(v_i)^+ (\mathbb{1}\{E_1^i\} + \mathbb{1}\{E_2^i\} + \mathbb{1}\{E_3^i\}) \right] \\ &\leq 4\mathbb{E}[\mathcal{R}_{r_1}] + 2\mathbb{E}[\mathcal{R}_{r_2}] \end{aligned}$$

and hence $\max\{\mathbb{E}[\mathcal{R}_{r_1}], \mathbb{E}[\mathcal{R}_{r_2}]\} \geq \frac{1}{6}\mathbb{E}[\mathcal{R}_r^{VCG}]$. ■

4. REVENUE IN FULL INFORMATION GSP

We now wish to compare the revenue properties of GSP and VCG in the full information setting. We start by giving examples showing that there are no universal constants that bound these two quantities. Then we introduce a new benchmark related to VCG, and show that the GSP revenue is not too low relative to this benchmark.

4.1 Full Information Revenue: Examples

Unfortunately, there are no universal constants $c_1, c_2 > 0$ such that for every full information AdAuctions instance α, \mathbf{v} and for all equilibria b of GSP it holds that

$$c_1 \cdot \mathcal{R}^{VCG}(\mathbf{v}) \leq \mathcal{R}(\mathbf{b}) \leq c_2 \cdot \mathcal{R}^{VCG}(\mathbf{v}).$$

In fact, GSP can generate arbitrarily more revenue than VCG and vice-versa. For example, consider two players with $\alpha = \{1, 0\}$, $\mathbf{v} = \{2, 1\}$. Then VCG generates revenue 1, but GSP has the Nash equilibrium $\mathbf{b} = [2, 0]$ that generates no revenue.

As a counter-example for the second inequality, consider the following instance: $\alpha = \{1, 1 - \epsilon\}$, $\mathbf{v} = \{\epsilon^{-1}, 1\}$. Notice that the revenue produced by VCG is ϵ , while GSP has the equilibrium $\mathbf{b} = [1, 1]$ generating revenue 1.

4.2 Revenue Bound in Full Information

Next, we will prove that the GSP revenue cannot be much less than a revenue benchmark based on the VCG auction. Intuitively, the difficulty behind our bad examples is in extracting revenue from the player with the largest private value. Motivated by this, we consider the following benchmark:

$$\begin{aligned} \mathcal{B}(\mathbf{v}) &= \sum_{i=2}^n p_i^{VCG} \alpha_{\sigma(i)} = \sum_{i=2}^n \sum_{j>i} (\alpha_{j-1} - \alpha_j) v_j \\ &= \sum_{i=2}^n (i-2)(\alpha_{i-1} - \alpha_i) v_i \end{aligned}$$

which is the VCG revenue from players 2, 3, ..., n . Recall that in the full information setting we assumed that players are numbered such that $v_1 \geq v_2 \geq \dots$. We show that the GSP revenue is always at least half of this benchmark at any equilibrium. Thus, unless VCG gets most of its revenue from a single player, GSP revenue will be within a constant factor of the VCG revenue.

Theorem 5 *Given an AdAuctions instance α, \mathbf{v} , and a Nash equilibrium \mathbf{b} of GSP, we have $\mathcal{R}(\mathbf{b}) \geq \frac{1}{2}\mathcal{B}(\mathbf{v})$, and this bound is tight.*

We prove Theorem 5 in two steps. First we define the concept of up-Nash² equilibrium for GSP, then we show that any inefficient Nash equilibria can be written as an efficient up-Nash equilibrium. In the second step, we prove the desired revenue bound for all efficient up-Nash equilibria.

Definition 6 *Given a bid profile \mathbf{b} , we say it is **up-Nash** for player i if he can't increase his utility by taking some slot*

²Our concepts of up-Nash and down-Nash equilibria are very similar to the concepts of upwards stable and downwards stable equilibria in Markakis and Telelis [20]

above, i.e.

$$\alpha_{\sigma(i)}(v_i - b_{\pi(\sigma(i)+1)}) \geq \alpha_j(v_i - b_{\pi(j)}), \forall j < \sigma(i).$$

Analogously, we say that \mathbf{b} is **down-Nash** for player i if he can't increase his utility by taking some slot below, i.e.

$$\alpha_{\sigma(i)}(v_i - b_{\pi(\sigma(i)+1)}) \geq \alpha_j(v_i - b_{\pi(j+1)}), \forall j > \sigma(i).$$

A bid profile is **up-Nash** (**down-Nash**) if it is up-Nash (**down-Nash**) for all players i . Clearly a bid profile \mathbf{b} is a Nash equilibrium iff it is both up-Nash and down-Nash.

Lemma 7 *If a bid profile \mathbf{b} is a Nash equilibrium, then the bid profile \mathbf{b}' where $b'_i = b_{\pi(i)}$ is up-Nash.*

PROOF. We will prove the lemma by modifying bid profile \mathbf{b} in a sequence of steps. Fix some $k \leq n$, and suppose that \mathbf{b} is a bid profile (with corresponding allocation π) such that

- players $j = 1, \dots, k$ satisfy the Nash conditions (i.e. both up-Nash and down-Nash) in \mathbf{b} ,
- players $j = k + 1, \dots, n$ are such that $\sigma(j) = j$ and they satisfy the up-Nash conditions in \mathbf{b} ,
- $\sigma(k) < k$.

We then define \mathbf{b}' by swapping the bids of players k and $\pi(k)$, that is setting $b'_i = b_i$ for $i \neq k, \pi(k)$, $b'_k = b_{\pi(k)}$, and $b'_{\pi(k)} = b_k$. We claim that \mathbf{b}' is up-Nash for players k, \dots, n and Nash for the remaining players. This then implies the desired result, since we can apply this operation for $k = n, \dots, 2$, resulting in the required bid profile.

Since our transformation does not alter the bids associated with given slots, we just need to check three things: the up and down-Nash inequalities for player $\pi(k)$, and the up-Nash inequality for player k .

Under bid profile \mathbf{b}' , player $\pi(k)$ gets slot $\sigma(k)$. This player doesn't want to change his bid to win any slot $j > \sigma(k)$ since in the bid profile \mathbf{b} player k with lower value didn't want to get these slots. We therefore have $\alpha_{\sigma(k)}(v_k - b_{\pi(\sigma(k)+1)}) \geq \alpha_j(v_k - b_{\pi(j+1)})$ and since $v_{\pi(k)} \geq v_k$, we conclude

$$\alpha_{\sigma(k)}(v_{\pi(k)} - b_{\pi(\sigma(k)+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j+1)}). \quad (1)$$

To see that player $\pi(k)$ would not prefer to take any slot $j < \sigma(k)$, notice that $\pi(k)$ didn't want to move to a higher slot in \mathbf{b} , so $\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$. This, combined with equation (1) for $j = k$ (stating that $\pi(k)$ prefers slot $\sigma(k)$ to k) gives the up-Nash inequality for player $\pi(k)$.

Next consider player k in bid profile \mathbf{b}' , where we gets slot k . We wish to prove the up-Nash inequality for k . Notice that, in \mathbf{b} , $\pi(k)$ had slot k and didn't want to switch to a higher slot, so we know $\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$.

Now, since $v_{\pi(k)} \geq v_k$, we have $\alpha_k(v_k - b_{\pi(k+1)}) \geq \alpha_j(v_k - b_{\pi(j)})$ which is the desired inequality. ■

Now to prove Theorem 5 we use the up-Nash profile \mathbf{b}' .

Proof of Theorem 5 : Given any Nash equilibrium \mathbf{b} , consider the bid profile \mathbf{b}' of Lemma 7, which is an up-Nash equilibrium in which each player k occupies slot k . By the up-Nash inequalities, for each k we have

$$\alpha_k(v_k - b'_{k+1}) \geq \alpha_{k-1}(v_k - b'_{k-1}).$$

We can rewrite this as

$$\alpha_{k-1}b'_{k-1} \geq (\alpha_{k-1} - \alpha_k)v_k + \alpha_k b'_{k+1}.$$

Then, since $\alpha_k \geq \alpha_{k+1}$,

$$\alpha_{k-1}b'_{k-1} \geq \sum_{j \in k+2\mathbb{N}} (\alpha_{j-1} - \alpha_j)v_j$$

where $k + 2\mathbb{N} = \{k, k + 2, k + 4, \dots\}$. Now we can bound the revenue of \mathbf{b} :

$$\begin{aligned} \mathcal{R}(\mathbf{b}) = \mathcal{R}(\mathbf{b}') &= \sum_k \alpha_k b'_{k+1} \geq \sum_k \alpha_{k+1} b'_{k+1} \geq \\ &\geq \sum_k \sum_{j \in k+2+2\mathbb{N}} (\alpha_{j-1} - \alpha_j)v_j \geq \\ &\geq \sum_{k=2}^n \frac{k-2}{2} (\alpha_{k-1} - \alpha_k)v_k = \frac{1}{2}\mathcal{B}(\mathbf{v}). \end{aligned}$$

■

To show that the bound in Theorem 5 is tight, consider the following example with n slots and n players, parametrized by $\delta > 0$:

$$\begin{aligned} \alpha &= [1, 1, \dots, 1, 1 - \delta, 0], \\ \mathbf{v} &= [1, 1, \dots, 1, 1, \delta], \\ \mathbf{b} &= [\delta, \delta, \dots, \delta, \delta, 0]. \end{aligned}$$

In this case $\mathcal{R}(\mathbf{b}) = (n - 2)\delta + \delta(1 - \delta)$ and $\mathcal{R}^{VCG}(\mathbf{v}) = (2\delta - \delta^2)(n - 3) + \delta(1 - \delta)$. Therefore $\lim_{n \rightarrow \infty} \frac{\mathcal{R}(\mathbf{b})}{\mathcal{B}(\mathbf{v})} = 2 - \delta$ and it tends to 2 as $\delta \rightarrow 0$.

Notice that those bounds also carry for the case where there is a reserve price r . We compare the revenue $\mathcal{R}_r(\mathbf{b})$ with reserve price r , against a slightly modified benchmark: $\mathcal{B}_r(\mathbf{v})$ which is the revenue VCG $_r$ extracts from players $2, \dots, n$.

Corollary 8 *Let \mathbf{b} be a Nash equilibrium of the GSP $_r$ game, then $\mathcal{R}_r(\mathbf{b}) \geq \frac{1}{2}\mathcal{B}_r(\mathbf{v})$.*

PROOF. We can assume wlog that $v_i, b_i \geq r$ (otherwise those players don't participate in any of the auctions). We can define an upper-Nash bid profile \mathbf{b}' as in Lemma 7. Now, notice that all players in \mathbf{b}' are paying at least r per click. We can divide the players in two groups: players $1 \dots k$ are paying more than r in VCG $_r$ and players $k + 1 \dots n$ are paying exactly r . It is trivial that for the players $k + 1 \dots n$ we extract at least the same revenue under VCG $_r$ then under GSP $_r$. For the rest of the players we need to do the exact same analysis as in the proof of Theorem 5. ■

5. TRADEOFF BETWEEN REVENUE AND EFFICIENCY

In this section we consider the tradeoff between efficiency and revenue, and ask if optimal efficiency and optimal revenue can always be achieved in the same equilibrium. We give a negative answer to this question, showing that for some AdAuction instances, one can increase revenue by selecting inefficient equilibria. First we recall the equilibrium hierarchy briefly discussed in the introduction.

Then we characterize the maximum revenue possible for envy free equilibrium (that is always efficient). Does this equilibrium class generate more or less revenue than other

classes, such as efficient equilibria or all pure equilibria? This question of comparing the revenue of VCG and envy-free equilibria of GSP was addressed by [7], who show that the revenue in any envy-free equilibrium is at least that of the VCG outcome (i.e. the VCG outcome is the envy-free equilibria generating smallest possible revenue). Moreover, as we've shown, an envy-free equilibrium can generate arbitrarily more revenue than the VCG outcome. Varian [26] shows how to compute the revenue optimal envy free Nash equilibrium, if we assume that agents will overbid. Allowing overbidding can result in very high revenue (eg., the maximum valuation in a single item auction). Here we determine the maximum revenue that can be obtained if we do not assume that agents bid at envy-free equilibria, and without requiring that agents apply the dominated strategy of overbidding.

Finally, we use this characterization to we give a natural sufficient condition under which there is a revenue-optimal equilibrium that is efficient.

5.1 Equilibrium hierarchy for GSP

Edelman, Ostrovsky and Schwarz [7] and Varian [25] showed that the full information game always has a Pure Nash equilibrium, and moreover, there is a pure Nash equilibrium which has same outcome and payments as VCG. At this equilibrium, players bid

$$b_i^V = \frac{1}{\alpha_{i-1}} \sum_{j=i}^n (\alpha_{j-1} - \alpha_j)v_j.$$

The authors also define a class of equilibria called **envy-free** or **symmetric equilibria**. This is the class of bid profiles \mathbf{b} such that

$$\alpha_{\sigma(i)}(v_i - b_{\sigma(i)+1}) \geq \alpha_j(v_i - b_{j+1}).$$

It is easy to see that all envy-free equilibria are Nash equilibria, though not all Nash equilibria are envy-free. The bid profiles that are envy-free are always efficient and the revenue of an envy-free equilibrium is always greater than or equal to the VCG revenue. That is, if \mathbf{b} is an envy-free equilibrium, $\mathcal{R}(\mathbf{b}) \geq \mathcal{R}^{VCG}(\mathbf{v})$.

Although all envy-free equilibria are efficient, there are efficient equilibria that are not envy-free, as one can see for example in Figure 1, as well as inefficient equilibria. We therefore have the following hierarchy:

$$\left\{ \begin{array}{c} \text{VCG} \\ \text{outcome} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{envy-free} \\ \text{equilibria} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{efficient} \\ \text{Nash eq} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{all} \\ \text{Nash} \end{array} \right\}$$

5.2 Envy-free and efficient equilibrium

As shown in the example of Figure 1, there are efficient equilibria that generate arbitrarily less revenue than any envy-free equilibrium. For the other direction, we show that all revenue-optimal efficient equilibria are envy-free.

Theorem 9 *For any AdAuctions instance such that $\alpha_i > \alpha_{i+1} \forall i$, all revenue-optimal efficient equilibria are envy-free. Moreover, we can write the revenue optimal efficient equilibrium explicitly as function of α, \mathbf{v} recursively as follows:*

$$\begin{aligned} b_n &= \min \left\{ v_n, \frac{\alpha_{n-1} - \alpha_n}{\alpha_{n-1}} v_{n-1} \right\}, \\ b_i &= \min \left\{ v_i, \frac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \frac{\alpha_i}{\alpha_{i-1}} b_{i+1} \right\} \quad \forall i < n. \end{aligned}$$




Figure 1: Equilibria hierarchy for GSP for $\alpha = [1, 1/2]$, $v = [1, 2/3]$: the strong blue dot represents the VCG outcome, the pattern region the envy-free equilibria, the blue region all the efficient equilibria and the red region the inefficient equilibria

PROOF. Given an efficient equilibrium \mathbf{b} , if it is not envy-free, we show that we can improve revenue by slightly increasing one of the bids. If the equilibrium is not envy-free, there is at least one player that envies the player above, i.e.

$$\alpha_i(v_i - b_{i+1}) < \alpha_{i-1}(v_i - b_i).$$

As pointed out in [7], if in an efficient equilibrium no player envies the above slot (i.e. no player i wants to take the above slot $i - 1$ by the price per click player i is paying) then the equilibrium is envy-free.

Let i be the player with the smallest index that envies slot $i - 1$. Consider the bid profile \mathbf{b}' such that $b'_j = b_j$ for $j \neq i$ and $b'_i = b_i + \epsilon$. We will verify that the Nash inequalities for player $i - 1$ still hold when $\epsilon > 0$ is sufficiently small. In other words, we will show that no Nash inequality for player $i - 1$ holds with equality in \mathbf{b} .

For slots $j > i - 1$, notice that

$$\alpha_j(v_i - b_{j+1}) \leq \alpha_i(v_i - b_{i+1}) < \alpha_{i-1}(v_i - b_i)$$

where the first is a standard Nash inequality and the second is the hypothesis that player i envies the above slot. Now, since $v_{i-1} > v_i$ in an efficient equilibrium, we have

$$\alpha_j(v_{i-1} - b_{j+1}) < \alpha_{i-1}(v_{i-1} - b_i).$$

For slots $j < i - 1$, we use the fact that player i is the first envious player. Also, without loss of generality, we can assume player 1 bids v_1 . Therefore we need to verify the Nash inequalities only for $j = 2, 3, \dots, k - 1$. We have

$$\alpha_{i-1}(v_i - b_i) \geq \alpha_j(v_i - b_{j+1}) > \alpha_j(v_i - b_j)$$

where the first inequality comes from the fact that player $i - 1$ doesn't envy any player j above him and the second inequality comes from the fact that $b_j > b_{j+1}$, since otherwise the player in slot j would envy the player in slot $j - 1$. This shows that the revenue optimal equilibrium is envy free.

To see that the bid profile defined in the theorem is optimal, we need to show the following things about this bid profile \mathbf{b} : (i) it is in Nash equilibrium, (ii) it is envy free, and (iii) no other efficient Nash equilibrium generates higher revenue. Begin by noticing that if \mathbf{b} is Nash, then player

$i - 1$ doesn't want to take slot i , for all i , and therefore $\alpha_{i-1}(v_{i-1} - b_i) \geq \alpha_i(v_{i-1} - b_{i+1})$ and this is satisfied by definition by the bid vector presented. Notice also that this series of inequalities implies an upper bound on the maximum revenue in an efficient equilibrium and this bound is achieved exactly by the bid profile defined above.

Furthermore, for all $j \leq i - 1$ we have $\alpha_{i-1}(v_j - b_i) \geq \alpha_i(v_j - b_{i+1})$ therefore by composing this expression with different values of i and j , it is straightforward to show that no player can profit by decreasing his bid. We prove that no player can profit by overbidding as a simple corollary of envy-freeness. For that, we need to prove that

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_{i-1}(v_i - b_i).$$

If $b_i = v_i$ than this is trivial. If not, then substitute the expression for b_i and notice it reduces to $v_{i-1} \geq v_i$. Now, this proved local envy-freeness, what implies that no player wants the slot above him by the price he player above him is paying. This in particular implies that no player wants to increase his bid to take a slot above. ■

5.3 Cost of efficiency: definition and example

Next we will analyze the relation between revenue and efficiency in GSP auctions.

We define the *cost of efficiency* for a given profile of click-through-rates as

$$\text{CoE}(\alpha) = \max_{\mathbf{v}} \frac{\max_{\mathbf{b} \in \text{NASH}(\alpha, \mathbf{v})} \mathcal{R}(\mathbf{b})}{\max_{\mathbf{b} \in \text{EFFNASH}(\alpha, \mathbf{v})} \mathcal{R}(\mathbf{b})}$$

where NASH is the set of all bid profiles in Nash equilibrium and EFFNASH is the set of all efficient Nash equilibrium.

First we give examples in which $\text{CoE}(\alpha) > 1$, in which case all revenue-optimal equilibria occur at inefficient equilibria. Our example will have $n = 3$ slots and advertisers. The click-through rates are given by $\alpha = [1, \frac{2}{3}, \frac{1}{6}]$ and the agent types are $v = [1, \frac{7}{8}, \frac{6}{8}]$. In this case, the best possible revenue generated by an efficient outcome is given by $\frac{1}{3} + \frac{7}{8} \approx 1.20833$ (this can be calculated using the formula in Theorem 9). However, for the (inefficient) allocation $\pi = [2, 1, 3]$, there is an equilibrium that generates revenue 1.21528.

In Figure 2 we calculate this value empirically for each $\alpha = [1, \alpha_2, \alpha_3]$, where each α_i is an integer multiple of 0.01 in $[0, 1]$. In all cases we found that $1 \leq \text{CoE}(\alpha) < 1.1$. The color of (α_1, α_2) in the graph corresponds to $\text{CoE}(1, \alpha_2, \alpha_3)$, where blue represents 1 and red represents 1.1. By solving a constrained non-linear optimization problem, one can show that the worst CoE for 3 slots is 1.09383.

5.4 Efficiency Versus Revenue when Click-Through-Rates are Convex

We now present a condition on α that implies $\text{CoE}(\alpha) = 1$. Our condition is that the click-through-rates are *convex*, meaning that $\alpha_i - \alpha_{i+1} \geq \alpha_{i+1} - \alpha_{i+2}$ for all i . We note that most models for CTRs studied in the literature satisfy convexity, such as exponential CTRs [15] and Markovian user models [1].

Theorem 10 *If click-through-rates α are convex then there is a revenue-maximizing Nash equilibrium that is also efficient.*




Figure 2: Cost of efficiency for $\alpha = [1, \alpha_2, \alpha_3]$: in the plot, blue means 1.0 and red means 1.1.

Our proof follows from a local improvement argument: given an instance with convex click-through-rates and an equilibrium that is not efficient, we show how to either improve its revenue or its welfare. A key step of the proof is bounding the maximum revenue possible in equilibrium for a given allocation, extending Theorem 9 to inefficient allocations.

PROOF. Let \mathbf{b} be the revenue maximizing efficient Nash equilibrium. Fix an allocation π and let \mathbf{b}' be an equilibrium under allocation π . We say that \mathbf{b} is **saturated** for slot i if $b_i = v_i$. We start by presenting the proof of the theorem under the simplifying assumption that no slot is saturated in the maximum revenue equilibrium..

Under the no-saturation assumption, Theorem 9 implies

$$\mathcal{R}(\mathbf{b}) = \sum_i \alpha_i b_{i+1} = \sum_i \sum_{j \geq i} (\alpha_j - \alpha_{j+1}) v_j. \quad (2)$$

Notice that we can view this expression as a dot product of two vectors where one has elements of the form v_i and other has elements in the form $\alpha_j - \alpha_{j+1}$. Notice also that due to the convexity assumption, this is a dot product of two sorted vectors. Now, for \mathbf{b}' , we will bound revenue as follows. Define $m(\pi, i, j) = \max\{\pi(i), \pi(i+1), \pi(i+2), \dots, \pi(j)\}$. Let p be such that the $k = i, i+1, \dots, i+p$ are all the indices such that $m(\pi, i, k) = \pi(i)$. Now, notice that the player in slot i doesn't want to take slot $i+p+1$, so

$$\alpha_i (v_{\pi(i)} - b'_{\pi(i+1)}) \geq \alpha_{i+p+1} (v_{\pi(i)} - b'_{\pi(i+p+2)}).$$

This implies

$$\begin{aligned} \alpha_i b'_{\pi(i+1)} &\leq \alpha_{i+p+1} b'_{\pi(i+p+2)} + (\alpha_i - \alpha_{i+p+1}) v_{\pi(i)} \\ &= \alpha_{i+p+1} b'_{\pi(i+p+2)} + \sum_{j=i}^{i+p} (\alpha_j - \alpha_{j+1}) v_{m(\pi, i, j)}. \end{aligned}$$

We can now apply recursion to conclude that $\alpha_i b'_{\pi(i+1)} \leq$

$\sum_{j \geq i} (\alpha_j - \alpha_{j+1}) v_{m(\pi, i, j)}$, and hence

$$\mathcal{R}(\mathbf{b}') = \sum_i \alpha_i b'_{\pi(i+1)} \leq \sum_i \sum_{j \geq i} (\alpha_j - \alpha_{j+1}) v_{m(\pi, i, j)}. \quad (3)$$

Notice that equation (3) can also be written as a dot product between two vectors of type v_i and $\alpha_j - \alpha_{j+1}$. If we sort the vectors, we see that the $(\alpha_j - \alpha_{j+1})$ -vector is the same in both (2) and (3). Moreover, the sorted vector of v_j for equation (3) is dominated by that of equation (2), in the sense that it is pointwise smaller. To see this, simply count how many times we have one of v_1, \dots, v_i appear in both vectors for each index i : for equation (2) they appear $\sum_{j=1}^i j$ times, whereas for equation (3) they appear at most $\sum_{j=1}^i 1 + \max\{p \mid m(\pi(j, j+p)) \leq i\} \leq \sum_{j=1}^i j$ times. Since the $(\alpha_j - \alpha_{j+1})$ -vectors are the same in both equations, the v_i vector in the first equation dominates the order and in the first equation both vectors are sorted in the same order, so it must be the case that $\mathcal{R}(\mathbf{b}) \geq \mathcal{R}(\mathbf{b}')$.

It remains to remove our simplifying assumption about saturation and prove the general result. Let \mathbf{b} be the optimal efficient equilibrium and let $S \subseteq [n+1]$ be the set of saturated bids, including $n+1$ (where we consider a "fake" player $n+1$ with $b_{n+1} = v_{n+1} = 0$), i.e., $i \in S$ iff $b_i = v_i$. Let $S(i) = \min\{j \in S; j > i\}$.

Given an allocation π , we wish to define an upper bound, $\overline{\mathcal{R}}_\pi$, on the revenue of a bid profile that induces allocation π at equilibrium. To this end, we define

$$B_\pi(j) = \begin{cases} \alpha_{S(j)-1} v_{S(j)} + \sum_{i=\sigma(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi, \sigma(j), i)} & \text{if } \sigma(j) \leq S(j) - 1 \\ \alpha_{S(j)-1} v_{S(j)} - v_j (\alpha_{S(j)-1} - \alpha_{\sigma(j)}) & \text{if } \sigma(j) \geq S(j) - 1 \end{cases}$$

We then define

$$\overline{\mathcal{R}}_\pi = \sum_j B_\pi(j).$$

We claim that this is, indeed, an upper bound on revenue. Moreover, this bound is tight for revenue at efficient equilibria (i.e. when π is the identity id).

Claim 11 *If bid profile \mathbf{b} induces allocation π at equilibrium, then $\mathcal{R}(\mathbf{b}) \leq \overline{\mathcal{R}}_\pi$.*

Claim 12 *There exists an efficient equilibrium with revenue $\overline{\mathcal{R}}_{id}$.*

Using these two claims we want to argue that id is the permutation that maximizes $\overline{\mathcal{R}}_\pi$ and therefore we can show that for all inefficient bid profile \mathbf{b}' we have

$$\mathcal{R}(\mathbf{b}') \leq \overline{\mathcal{R}}_\pi \leq \overline{\mathcal{R}}_{id} = \mathcal{R}(\mathbf{b}).$$

To show this, consider some permutation π . Let $j = \max\{k : \pi(k) \neq k\}$ and define a permutation π' such that $\pi'(k) = k$ for $k \geq j$ and $\pi'(k) = \pi(k)$ for $k < \sigma(j)$ and $\pi'(k) = \pi(k+1)$ for $\sigma(j) \leq k < j$. Essentially this is picking the last player that is not allocated to his correct slot and bring him there. Now, if we prove that $\overline{\mathcal{R}}_{\pi'} \geq \overline{\mathcal{R}}_\pi$, then we are done, since we can repeat this procedure many times and get to id .

Claim 13 $\overline{\mathcal{R}}_{\pi'} \geq \overline{\mathcal{R}}_\pi$.

This completes the proof, subject to our claims, which we prove in Appendix C of the full version [17]. ■

6. REFERENCES

- [1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pál. Sponsored search auctions with markovian users. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, WINE '08, pages 621–628, Berlin, Heidelberg, 2008. Springer-Verlag.
- [2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, New York, NY, USA, 2006. ACM.
- [3] T. Borgers, I. Cox, M. Pesendorfer, and V. Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. Technical report, 2007.
- [4] I. Caragiannis, P. Kanellopoulos, C. Kaklamanis, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *EC'11*, 2011.
- [5] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 311–320, New York, NY, USA, 2010. ACM.
- [6] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, (43):192–198, 2007.
- [7] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, March 2007.
- [8] B. Edelman and M. Schwarz. Optimal auction design in a multi-unit environment: The case of sponsored search auctions. Technical report, 2006.
- [9] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, (2), 2010.
- [10] R. D. Gomes and K. S. Sweeney. Bayes-nash equilibria of the generalized second price auction. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 107–108, New York, NY, USA, 2009. ACM.
- [11] M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd national conference on Artificial intelligence - Volume 1*, pages 58–65. AAAI Press, 2007.
- [12] J. Hegeman. Facebook's ad auction. Talk at Ad Auctions Workshop, May 2010, 2010.
- [13] U. Krengel and L. Sucheston. Semiamarts and finite values. In *Bull. Amer. Math. Soc.* 83, 1977.
- [14] U. Krengel and L. Sucheston. On semiamarts and amarts, and processes with finite value. In *Probability on Banach Spaces*, 1978.
- [15] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.
- [16] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.
- [17] B. Lucier, R. P. Leme, and Éva Tardos. On revenue in the generalized second price auction (*full version*). www.cs.cornell.edu/~renatoppl/gsp-revenue-www2012.pdf, 2012.
- [18] B. Lucier and R. Paes Leme. Gsp auctions with correlated types. In *EC'11*, 2011.
- [19] P. Maille, E. Markakis, M. Naldi, G. D. Stamoulis, and B. Tuffin. Sponsored search auctions: An overview of research with emphasis on game theoretic aspects. To appear in the Electronic Commerce Research journal (ECR), 2012.
- [20] E. Markakis and O. Telelis. Discrete strategies in keyword auctions and their inefficiency for locally aware bidders. In *Workshop on Internet and Network Economics (WINE)*, 2010.
- [21] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [22] R. Paes Leme and E. Tardos. Pure and bayes-nash price of anarchy for generalized second price auctions. In *FOCS10: 51st Annual IEEE Symposium on Foundations of Computer Science*, 2010.
- [23] F. Pin and P. Key. Stochastic variability in sponsored search auctions: observations and models. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 61–70, New York, NY, USA, 2011. ACM.
- [24] H. Varian. Mechanism design double feature. Talk at New York Computer Science and Economics Day, October 2010, 2010.
- [25] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.
- [26] H. R. Varian. Online ad auctions. *American Economic Review Papers and Proceedings*, 2009.