On Revenue in the Generalized Second Price Auction

Brendan Lucier
Dept of Computer Science
University of Toronto
blucier@cs.toronto.edu

Renato Paes Leme*
Dept of Computer Science
Cornell University
renatoppl@cs.cornell.edu

Éva Tardos†
Dept of Computer Science
Cornell University
eva@cs.cornell.edu

ABSTRACT

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. Most previous work of GSP focuses on envy free equilibria of the full information version of this game. Envy-free equilibria are known to obtain at least the revenue of the VCG auction. Here we consider revenue in equilibria that are not envy-free, and also consider revenue in the Bayesian version of the game.

We show that, at equilibrium, the GSP auction obtains at least half of the revenue of the VCG mechanism excluding the payment of a single participant. This bound is tight, and we give examples demonstrating that GSP cannot approximate the full revenue of the VCG mechanism either in the full information game, or in the Bayesian version (even if agent values are independently drawn from identical uniform distributions). We also show that the GSP revenue approximates the VCG revenue in the Bayesian game when the click-through rates are well separated.

We also consider revenue-maximizing equilibrium of GSP in the full information model. We show that if click-through rates satisfy a natural convexity assumption, then the revenue-maximizing equilibrium will necessarily be envy-free. In particular, it is possible to maximize revenue and social welfare simultaneously. On the other hand, without this convexity assumption, 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

*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.

1. INTRODUCTION

The sale of sponsored search advertising space is a primary source of revenue for Internet companies, and responsible for billions of dollars in annual advertising revenue [5]. Given the importance of this market for services on the Internet, it is crucial to choose a good mechanism. 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 recently adopted by Facebook for its AdAuction system [10]. In fact, Google also considered switching its advertising platform to a VCG auction some years ago, but eventually decided against it [17]. This apparent tension begs the question of 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 consider the comparison from the point of view of the seller and compare the revenue properties of the GSP and VCG auctions.

Let us first briefly describe a simple model of the market and the two auctions we consider. 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. If the user clicks on an advertisement, the advertiser pays a fee to the search provider; this is known as the "pay-per-click" pricing model. There are multiple possible positions (or "slots") in which an ad may appear, and the probability that a user clicks an ad depends on its slot: ads closer to the top of the page are likely to receive more clicks. In the simplest model, we have a click-through rate (CTR) associated with each slot (where a higher ad slot will have a larger click-through rate), where the CTR of a slot is 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 - the advertiser's maximum willingness to pay per click. Bids are used by the search engine to determine ad placements and prices. The mechanism used by Google, Bing, and Yahoo add variations to this basic scheme, such as adjusting the prices using ad quality, using a reserve price, and budgets, etc. In this paper, we will consider a game theoretical model of auctions without budgets, and will assume all ads have the same quality for simplicity of presentations.

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 CTRs. 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 behaviour. Indeed, strategic manipulation of bids is well-documented in historical GSP bidding data [4].

Since the VCG mechanism is truthful, the revenue of VCG is simply the revenue generated when all bidders declaring 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 by unilaterally deviating. Our goal, then, is a comparison between the revenue of the VCG auction and the revenue of GSP at equilibrium. Note that since there will not be a unique Nash equilibrium of GSP in general, there may be many possible revenue amounts generated by GSP.

This basic model of the GSP auction was first introduced by Edelman et al [5] and Varian [18]. Both papers consider a more restrictive notion of equilibrium than Nash, which they call envy-free or symmetric equilibrum respectively. They show that all *envy-free* equilibria are efficient and generate revenue at least as much as VCG. Both [5] and [18] 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, 7]. Further, the notion of envy-free equilibria applies only in the full information game. In fact [9] shows that an efficient equilibrium may not exist in the Bayesian game (not even when the valuations are drawn from identical uniform distributions). For the broader class of all Nash equilibria, the revenue properties of GSP are not understood. The primary focus of this paper is to study the revenue of the GSP auction, in relation to VCG, over the set of all Nash equilibria (including inefficient ones). We ask: how is the revenue of GSP affected if one cannot assume that agents necessarily converge to an envy-free equilibrium?

In sections 3 and 5 we will consider the full information game. For many keywords the search auction is repeated many times each day. In general, the repeated nature of this auction allows for complex strategic interactions. However, agents can be predicted to infer each others' bidding preferences over time, and thus we can approximate the general behaviour by assuming that players converge to a stationary equilibrium. If one assumes that, in fact, bidders learn each others' values over time, then results in a full-information Nash equilibrium of the one-shot game.

In section 4 we consider the Bayesian version of this game. For many keywords, the ability of a player to exactly predict his opponents types is impaired. Due to factors, such as the budgets, the algorithms used to compute quality factors (which depend on many characteristics of each query, such as origin, time, search history of the user), and the underlying ad allocation algorithm, each auction is different, even when it is triggered by the same search term. To capture this measure of uncertainty, we will also consider equilibria in a Bayesian partial information model.

Results. We begin by considering lower bounds on the revenue generated by GSP. One might wish to bound the revenue of GSP with respect to the revenue of VCG, but we demonstrate that the revenue of GSP at equilibrium may be arbitrarily less than that of VCG. However, we can bound the revenue of GSP with respect to a related benchmark: 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. Furthermore, this result also holds when an arbitrary reserve price is set upon the sale of a slot. We also provide an example illustrating that the factor of 2 in our analysis is tight.

One might hope that the gap between GSP and VCG revenue is an artifact of agents having very different values, or an artifact of the full-information setting. To the contrary, we demonstrate that this gap is essential in a broad setting: even in a partial information setting where agents' values are drawn independently from identical uniform distributions, the gap between the VCG revenue and GSP revenue at Bayesian-Nash equilibrium can be arbitrarily large. On the other hand, if the slot CTRs satisfy a certain well-separatedness condition - namely that the CTRs of adjacent slots differ by at least a certain constant factor - then we prove that GSP does obtain a constant fraction of the VCG revenue even in settings of partial information, extending a result of Lahaie [11] who considered welfare under this assumption on the CTRs. This result holds even without the assumption that agents avoid dominated strategies, as long as there are at least three participants in the auction.

We then turn to an analysis of the maximum revenue attainable by the GSP mechanism. 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. We feel that the convexity assumption is quite natural; note that it is weaker than the common assumption that 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 an early survey of Lahaie et al [12] for an overview. The GSP model we adopt is due to Edelman et al [5] and Varian [18]. Both papers consider a more restrictive notion of equilibrium than Nash. Edelman et al calls it *envy-free equilibrium* and Varian calls it *symmetric equilibrium*. Both authors show that this class of equilibria produce always optimal social welfare and revenue at least as good as the revenue of VCG.

Varian [19] shows how to compute the revenue optimal envy free Nash equilibrium, however in his model, he 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 outcome. We show that in general inefficient equilibria can generate more revenue than efficient ones. However, under a natural convexity assumption on click-through rates, we show that the maximum revenue equilibrium is envy free, and hence efficient, and show how to compute it efficiently.

Edelman and Schwarz [7] model the repeated auctions for a keyword as a repeated game, and show using Myerson's optimal auction [16] that 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. This leaves open the question whether GSP may generate revenue much less than the VCG auction, which is the main question we consider. However, unlike Edelman and Schwarz [7] we do not consider a repeated game, as rational play in a repeated game is too complex. Rather consider all stable outcomes of the auction, not only those that arise as limits of rational repeated play, which makes our results more general.

Gomes and Sweeney [9] study GSP as a Bayesian game – analyzing the symmetric efficient equilibria of this auction using the Revenue Equivalence Theorem as the main tool. The authors analyze the influence of click-through-rates in the revenue and observe the counter-intuitive phenomenon by which revenue decreases when click-through-rates increase. They also discuss the influence of reserve prices.

Paes Leme and Tardos [14] showed that the social welfare of GSP in equilibrium is within a constant factor of the optimal social welfare – which is composed by the engine revenue and the players total surplus. Lucier and Paes Leme [15] recently improved the bound for the Bayesian version. In the present work we tackle a natural question arising from their work: even though GSP guarantees reasonably high welfare, how does it break down in terms of revenue and total surplus?

There has been considerable work focused on studying revenue properties of GSP either by analyzing real auction data or by running simulations. Athey and Nekipelov [2] study the effect of quality-factors uncertainty in the revenue. Lahaie [11], Lahaie and Pennock [13] and Feng et al [8] study the effect of different ranking functions. Borgers et al [3] study revenue for alternative auction formats. Edelman and Schwarz [6] study the effect of reserve prices.

Our results compare the revenue of different mechanisms at equilibrium. It is worth noting 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 the mechanisms generate the same outcome. As a result, this equivalence does not apply in the full information setting. For our result in the partial information setting we also consider revenue properties of inefficient allocations, while the VCG mechanism is efficient.

2. PRELIMINARIES

An AdAuctions instance is composed of n players and n slots. Each player has a value v_i for each click he gets and quality factor γ_i. Slot j has click-through-rate α_j. That means that if player i is allocated in slot j, he gets $\gamma_i \alpha_j$ clicks in expectation. For the rest of the paper, we assume that $\gamma_i = 1$ for clarity of exposition. Assume we number players such that $v_1 \geq v_2 \geq \ldots \geq v_n$ and $\alpha_1 \geq \ldots \geq \alpha_n$. Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ be the CTR vector and $\mathbf{v} = (v_1, \ldots, v_n)$ be the type vector.

A mechanism for the AdAuctions problem has the following form: Since valuations v_i are private information, it begins by eliciting some bid b_i for the players, which works as his "declared valuation". We call $\mathbf{b} = (b_1, \ldots, b_n)$ be the bid vector. Using the $\mathbf{b}$ and α, the mechanism chooses an allocation $\pi : [n] \to [n]$ which means that player $\pi(j)$ is allocated to slot j, and a price vector $\mathbf{p} = (p_1, \ldots, p_n)$, where p_i is the price that player i pays for click. Player i then, experiences utility $u_i(\mathbf{b}) = \alpha_{\sigma(i)}(v_i - p_i)$, where $\sigma(i) = \pi^{-i}(i)$.

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:

$$p_i^{VCG} = \frac{1}{\alpha_{\sigma(i)}} \sum_{j=\sigma(i)+1}^{n} (\alpha_{j-1} - \alpha_j) b_{\pi(j)}$$

If the bidders truthfully declare their valuation in both VCG and GSP, then GSP generates strictly more revenue, as the revenue associated with player i in VCG is

$$p_i^{VCG} \alpha_{\sigma(i)} = \sum_{j=\sigma(i)+1}^{n} (\alpha_{j-1} - \alpha_j) b_{\pi(j)} \leq \alpha_{\sigma(i)} b_{\sigma(i)+1}$$

which is the GSP price paid by player i. VCG has the remarkable property that 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:

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

GSP, however, doesn't have this property. In general, the bid profile $\mathbf{b} = \mathbf{v}$ is not an equilibrium, i.e., there exists a player i that can improve his utility by misreporting his true valuation. We are interested in the set of bid profiles that constitute a Nash equilibrium, i.e.:

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

We assume for the rest of this paper that players don't overbid, i.e., $b_i \leq v_i$, since bidding $b_i > v_i$ it is a weakly dominated strategy (see [14]).

We say that an equilibrium is **efficient** if it maximizes social welfare, i.e., which happens when $\pi(i) = i$ for each slot i.

We will also consider this comparison 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 then 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 then r, the remaining players are allocated using GSP, and the last player to be allocated pays price r per click.

2.1 Equilibrium hierarchy for GSP

Edelman, Ostrovsky and Schwarz [5] and Varian [18] showed that the full information game always has a Pure Nash equilibrium, and moreover, there is a pure Nash equilibrium with same outcome and payments as VCG. And this happens when 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 equilibria** or **symmetric**. 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 the condition above implies that the bid profile is a Nash equilibrium (but doesn't capture all possible Nash equilibria). The bid profiles that are envy-free are always efficient and that the generated revenue is greater than or equal to that in VCG, i.e, for $\mathbf{b}$ envy-free, $\mathcal{R}(\mathbf{b}) \geq \mathcal{R}^{VCG}(\mathbf{v})$.

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. One can illustrate them as:

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

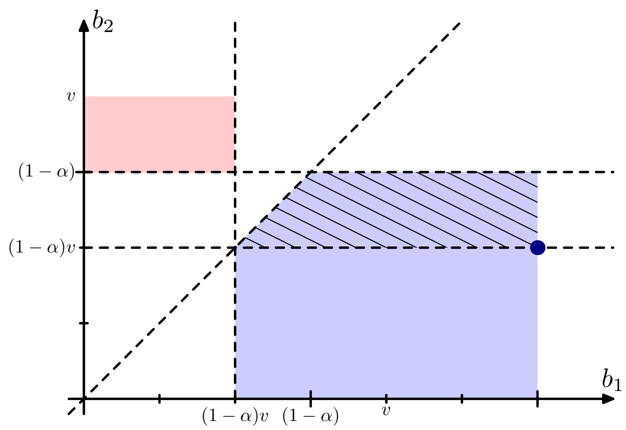

Figure 1: Equilibria hierarchy for GSP for $\alpha = [1, 1/2]$, $v = [1, 2/3]$: the strong blue dot represents the VCG outcome, the pattern region the envy-free equilibria, the blue region all the efficient equilibria and the red region the inefficient equilibria

2.2 Bayesian setting

The Bayesian setting models the uncertainty in the game and the fact the players know their own valuation but only know a distribution on the other players' valuations. In this model the values of the players are not fixed, but rather are random variables. The type vector $\mathbf{v}$ is drawn from a known distribution F. Each player learns his own value v_i and just knows the distribution of $\mathbf{v}_{-i}$. After learning his own value v_i, the player chooses a bid $b_i(v_i)$ to play in the AdAuctions game. The strategies are therefore bidding functions $b_i : \mathbb{R}_+ \to \mathbb{R}_+$, and we will continue to assume that player do not overbid, i.e., $b_i(v) \leq v$. A set of bidding functions is a Bayesian Nash equilibrium if:

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

.

3. REVENUE IN FULL INFORMATION GSP

The goal of this section is to compare the revenue properties of GSP and VCG. Unfortunately, there are no universal constants $c_1, c_2 > 0$ such that for every AdAuctions instance $\alpha, \mathbf{v}$ 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 the other way round, i.e., VCG can generate arbitrarily more revenue than GSP. A single item auction serves as a example for the first inequality: a single-item second price auction has many equilibria, some generating positive revenue and some generating zero. 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.

To give a bad 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.

However, we will prove that the GSP revenue cannot be much less than the VCG revenue in some sense. The main difficulty is to extract the revenue from the first player, a difficulty that is common through the revenue literature. 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 $2, 3, \ldots, n$. Next, we show that the GSP revenue is reasonably high against this benchmark, i.e., unless VCG gets most of its revenue from the first player, GSP revenue will be within a constant factor of VCG revenue.

Theorem 1 *Given an AdAuctions instance $\alpha, \mathbf{v}$, then for any Nash equilibrium of GSP, then $\mathcal{R}(\mathbf{b}) \geq \frac{1}{2}\mathcal{B}(\mathbf{v})$, and this bound is tight.*

We prove this theorem in two steps: first we define the concept of up-Nash equilibrium for GSP, then we show that all innefficient Nash equilibria can be written as an efficient up-Nash equilibrium. In the second step, we show the bound above for all efficient up-Nash equilibria.

Definition 2 *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 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 3 *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 show that if $\mathbf{b}$ is a bid profile (with corresponding allocation π) such that:

- players $j = k+1, \ldots, n$ are such that $\sigma(j) = j$ and they satisfy up-Nash
- players $j = 1, \ldots, k$ satisfy Nash (i.e. both up-Nash and down-Nash)
- $\sigma(k) < k$

then we define $\mathbf{b}'$ by swapping the bids of players k and $\pi(k)$, that is setting $b'_i = b_i$ for $i \neq k, \pi(k)$, $b'_k = b_{\pi(k)}$, $b'_{\pi(k)} = b_k$. We claim to get a profile that is up-Nash for players $k, \ldots, n$ and Nash for the remaining players. Then applying this construction for $k = n, \ldots, 2$ gives us the desired result.

Since each slot continues to get the same bid, 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.

Player $\pi(k)$ now gets slot $\sigma(k)$. This players doesn't want to get any slot $j > \sigma(k)$ since in the bid-profile b player k with lower value didn't want to get these slots, hence we have:

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

and since $v_{\pi(k)} \geq v_k$ then:

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

To see that he doesn't want to take any slot $j < \sigma(k)$, notice that $\pi(k)$ didn't want to move to a higher slot in b:

$$\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$$

That, combined with equation (1) for $j = k$ stating that $\pi(k)$ prefers slot $\sigma(k)$ to k gives us the up-Nash inequality.

Player k now gets slot k. For the up-Nash inequality for k, we need to show that he doesn't want to take any slot $j < k$. Notice that in b $\pi(k)$ had slot k and didn't want to switch to a higher slot:

$$\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$$

Now, since $v_{\pi(k)} \geq v_k$, we have:

$$\alpha_k(v_k - b_{\pi(k+1)}) \geq \alpha_j(v_k - b_{\pi(j)})$$

■

Proof of Theorem 1 : Given any Nash equilibrium $\mathbf{b}$, consider the bid profile $\mathbf{b}'$ of the lemma, which is an up-Nash equilibrium by the last lemma and player k occupies slot k. We use the fact that player k doesn't want to take slot $k-1$:

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

We can rewrite that as:

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

we have $\alpha_k \geq \alpha_{k+1}$, therefore:

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

where $k + 2\mathbb{N} = \{k, k+2, k+4, \ldots\}$. Now, we can bound:

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

■

To show that this bound is tight, consider the following example with n slots and n players parametrized by δ:

$$\alpha = [1, 1, 1, 1, \ldots, 1, 1-\delta, 0]$$

$$\mathbf{v} = [1, 1, 1, 1, \ldots, 1, 1, \delta]$$

$$\mathbf{b} = [\delta, \delta, \delta, \delta, \ldots, \delta, \delta, 0]$$

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

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

Corollary 4 *Let* $\mathbf{b}$ *be a Nash equilibrium of the GSP$_r$ game, then* $\mathcal{R}(\mathbf{b}) \geq \frac{1}{2}\mathcal{B}_r(\mathbf{v})$.

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 3. Now, notice that all players in $\mathbf{b}'$ are paying at least r per click. We can divide the players in two groups: players $1 \ldots k$ are paying more than r in VCG$_r$ and player $k+1 \ldots n$ are paying exactly r. It is trivial that for the players $k+1 \ldots n$ 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 1. ∎

4. REVENUE IN THE BAYESIAN SETTING

We showed in the full information setting that there are AdAuctions instances $\alpha, \mathbf{v}$ such that VCG generates positive revenue and there are GSP equilibria generating no revenue. One might ask if this can also happen in the Bayesian setting or if assuming, say iid players with some sort of well-behaved valuation distribution will eliminate the problem. Unfortunately, this is not the case.

Example. We show one example in the Bayesian setting where VCG generates positive revenue and GSP has a Bayesian-Nash equilibrium that generates zero revenue. Consider three players with iid valuations $v_i \sim$ Uniform$([1, 2])$ 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 $\mathbf{v}$). So, we know that:

$$\mathbb{E}[\mathcal{R}^{VCG}(\mathbf{v})] = \mathbb{E}[0.5v^{(2)}] > 0$$

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 valuation is $v_i > 0$ and notice that his current utility is $0.5v_i$ where if he changed their bid to $b > 0$, his utility would be:

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

for all bids $b' \geq 1$ (notice that $b' < 1$ doesn't change his utility since $b_1 \geq 1$ always).

The intuition behind this example is that players would like to stay away for slot 2 if its price is positive: so they want to go either for slot 1, since it has more clicks, or for slot 3 since it is cheaper. This generates this kind of anomaly. ∎

The main result of this section is that we can avoid the anomaly described above if we assume that the slot click-through-rates are well separated, in the sense of [11]. We say that click-through-rates are δ-well separated if $\alpha_{i+1} \leq \delta\alpha_i$ for all i.

Lemma 5 *If click-through-rates are δ-well separated, then bidding $b_i(v_i) < (1-\delta)v_i$ is dominated by playing $(1-\delta)v_i$.*

PROOF. If a player is playing $b_i < (1-\delta)v_i$, if he increases his bid to $b'_i = (1-\delta)v_i$ then with some probability he still gets the same slot (event S) and with some probability he gets a better slot (event B). Then clearly $\mathbb{E}[u_i(b_i, b_{-i})|v_i] \leq \mathbb{E}[u_i(b'_i, b_{-i})|v_i]$ since the expectation conditioned to S is the same and conditioned to B it can only increase by changing the bid to b'_i. To see that, let $\alpha_{\pi(i)}$ be the slot player i gets under b_i and $\alpha_{\pi'(i)}$ the slot he gets under b'_i. Conditioned on B we know that $\alpha_{\pi'(i)} \geq \delta^{-1}\alpha_{\pi(i)}$, and this generates revenue at least $\alpha_{\pi'(i)}(v_i - b'_i)$, while the revenue with bid b_i was at most $\alpha_{\pi(i)}v_i$, which implies the claim:

$$\begin{aligned}E[u_i(b_i, b_{-i})|v_i, B] &\leq \mathbb{E}[\alpha_{\pi(i)}v_i|v_i, B] \leq \mathbb{E}[\delta\alpha_{\pi'(i)}v_i|v_i, B] = \\ &= \mathbb{E}[\alpha_{\pi'(i)}(v_i - (1-\delta)v_i)|v_i, B] \leq \\ &\leq \mathbb{E}[u_i(b'_i, b_{-i})|v_i, B]\end{aligned}$$

∎

If one eliminates the strategies $b_i(v_i) < (1-\delta)v_i$ from the players strategy set, then it is easy to see that under any Bayesian-Nash equilibrium eliminating those dominated strategies, $\mathbb{E}_{\mathbf{v}}\mathcal{R}(b(\mathbf{v})) \geq (1-\delta)\mathbb{E}_{\mathbf{v}}\mathcal{R}^{VCG}(\mathbf{v})$.

Corollary 6 *If click-through-rates are δ-well separated, and all players play un-dominated strategies, than*

$$\mathcal{R}(\mathbf{b}) \geq (1-\delta)\mathcal{R}^{VCG}(\mathbf{v})$$

Next, we consider whether it is really necessary to eliminate dominated strategies, as players may not know what are all their dominated strategies. If we allow players to use dominated strategies, then we might have equilibria with very bad revenue compared to VCG, as one can see in the following example:

Example. Consider two players with iid valuations $v_i \sim$ Uniform$([0, 1])$ and two slots with $\alpha = [1, 1-\epsilon]$. Then VCG generates revenue $\mathbb{E}[\mathcal{R}^{VCG}(\mathbf{v})] = \mathbb{E}[\epsilon \min\{v_1, v_2\}] = O(\epsilon)$. However, consider the following equilibrium:

$$b_1(v_1) = \begin{cases} \epsilon(1-\delta), & v_1 \geq \epsilon(1-\delta) \\ \epsilon v_1, & v_1 < \epsilon(1-\delta) \end{cases}$$

$$b_2(v_2) = \begin{cases} \epsilon, & v_2 \geq 1-\delta \\ \epsilon^2(1-\delta), & \epsilon(1-\delta) \leq v_2 < 1-\delta \\ \epsilon v_2, & v_2 < \epsilon(1-\delta) \end{cases}$$

It is not hard to check that this is an equilibrium. In fact, for two player GSP in the Bayesian setting, playing $(\alpha_1 -$

$\alpha_2)v_i/\alpha_1$ is a best reply - and any bid that gives the player the same outcome is also a best reply. So, in the above example, one can simply check that the bids generate the same utility as bidding $b_i(v_i) = \epsilon v_i$. This example generates revenue $\mathbb{E}\mathcal{R}(\mathbf{b}) = O(\epsilon(\epsilon + \delta))$, so taking $\delta = O(\epsilon)$ in the above example give us $O(\epsilon^2)$ revenue. ■

The following Theorem is a version of Corollary 6 that doesn't depend on eliminating dominated strategies:

Theorem 7 *With n players with iid valuations v_i and δ-well separated click-through-rates, then for all (non-overbidding) Bayesian-Nash equilibria $\mathbf{b}$:*

$$\mathbb{E}\mathcal{R}(\mathbf{b}) \geq \frac{n-2}{n}(1-\delta)\mathbb{E}\mathcal{R}^{VCG}(\mathbf{v})$$

Proof. Given a profile $\mathbf{b}$ in Bayesian-Nash equilibrium and fixed two players i and j, we have that:

$$\mathbb{P}_{v\sim F}[b_i(v) < (1-\delta)v - \epsilon, b_j(v) < (1-\delta)v - \epsilon] = 0$$

in fact, suppose the contrary. Then there is $\epsilon' \ll \epsilon$ such that if we take $F' = F|_{[v^0-\epsilon', v^0+\epsilon']}$ then:

$$\mathbb{P}_{v\sim F'}[b_i(v) < (1-\delta)v - \epsilon, b_j(v) < (1-\delta)v - \epsilon] > 0$$

For ϵ' small enough $\underline{v_0} = v^0 - \epsilon$ and some $\epsilon'' < \epsilon$, then:

$$\mathbb{P}_{v\sim F'}[b_i(v) < (1-\delta)\underline{v_0} - \epsilon'', b_j(v) < (1-\delta)\underline{v_0} - \epsilon''] > 0$$

Now pick v^i, v^j in this interval such that $\mathbb{P}_{v\sim F'}[b_i(v^i) \leq b_i(v) < (1-\delta)\underline{v_0}] > 0$ and the same for j. By lemma 5, playing $(1-\delta)v^i$ is a best response, then for player j for example, it can't be the case that any of the other players play between $b_j(v^j)$ and $(1-\delta)v^j$ with positive probability. Therefore:

$$\mathbb{P}_{v\sim F'}[b_j(v) \in [b_i(v^i), (1-\alpha)v^i)] = 0$$

$$\mathbb{P}_{v\sim F'}[b_i(v) \in [b_j(v^j), (1-\alpha)v^j)] = 0$$

but notice this is a contradiction.

Now, we can think of the procedure of sampling $\mathbf{v}$ iid from F in the following way: sample $v_i'' \sim F$ iid , let v_i' be the sorted valuations, and then apply a random permutation $\tau \in S_n$ to the values so that $v_i = v'_{\tau(i)}$. Notice that $\mathbf{v}$ is iid and now, notice that with $\geq 1 - \frac{2}{n}$ probability, v_i' and v_{i+1}' will generate $(1-\delta)v_i'$ and $(1-\delta)v_{i+1}'$ bids producing $(1-\delta)\alpha_i v_{i+1}'$ revenue, therefore:

$$\mathbb{E}\mathcal{R}(\mathbf{v}) \geq \mathbb{E}\sum_i \left(1 - \frac{2}{n}\right)(1-\delta)\alpha_i v'_{i+1} \geq \frac{n-2}{n}(1-\delta)\mathbb{E}\mathcal{R}^V(\mathbf{v})$$

■

5. REVENUE IN THE GSP HIERARCHY

Last section was mainly concerned in comparing the VCG outcome (which can be emulated by one particular equilibrium of GSP) with all possible equilibria of GSP. In section 3, we showed that there are equilibria in GSP that can generate arbitrarily more and arbitrarily less revenue than the VCG outcome. Now, we come back to the full information setting to compare the revenue extraction properties of the different classes of GSP equilibria. Can one equilibrium class generate more or less revenue than other?

This question of comparing the VCG outcome and envy-free equilibria was answered by [5], that show that the revenue in all envy-free equilibria is at least as good as the VCG outcome (i.e. the VCG outcome is the envy-free equilibria generating smallest possible revenue). It is easy to see that envy-free equilibria can generate arbitrarily more revenue than the VCG outcome - for example, when we showed that GSP can generate arbitrarily more revenue than VCG, the GSP equilibrium was envy free. Varian [19] shows how to compute the revenue optimal envy free Nash equilibrium, if agents are allowed to overbid. Here we consider the question of maximum revenue equilibria without the assumption of envy-free outcome, and do not allow overbidding.

5.1 Envy-free and efficient equilibrium

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

Theorem 8 *For any AdAuctions instance such that $\alpha_i > \alpha_{i+1}, \forall i$, all the revenue-optimal efficient equilibria are envy-free. Moreover, we can write the revenue optimal efficient equilibrium explicitly as function of $\alpha, \mathbf{v}$.*

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 [5], if an efficient equilibrium is such that 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. Take i to be the player with smallest index that has this property.

Now, consider the bid profile $\mathbf{b}'$ such that $b_j' = b_j$ for $j \neq i$ and $b_i' = b_i + \epsilon$. We need to check the Nash inequalities for player $i-1$ still hold for a sufficiently small $\epsilon > 0$. In other words, we need to show that no Nash inequality for player $i-1$ holded with equality for $\mathbf{b}$.

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 $v_{i-1} > v_i$ 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, wlog, we can assume player 1 bids v_1. Therefore we need to prove it just for $j = 2, 3, \ldots, k-1$:

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

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 $b_j > b_{j+1}$, otherwise the player in slot j would envy the player in slot $j-1$. ■

In fact, we can give a more explicit proof of Theorem 8 by showing the bid profile that generates largest revenue and verifying it is an envy-free equilibrium. Given $(\alpha, \mathbf{v})$ define a bid profile $\mathbf{b}$ in a bottom up fashion:

$$b_n = \min\left\{v_n, \frac{\alpha_{n-1} - \alpha_n}{\alpha_{n-1}} v_{n-1}\right\}$$

$$b_i = \min\left\{v_i, \frac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \frac{\alpha_i}{\alpha_{i-1}} b_{i+1}\right\}, \forall i = n-1, \ldots, 1$$

Now, we need to show that (i) it is in Nash equilibrium; (ii) it is envy free and (iii) no efficient Nash has higher revenue than the one above. Begin by noticing that if $\mathbf{b}$ is Nash, then player $i-1$ doesn't want to take the slot 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 it gives 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.

5.2 Cost of efficiency

Here, we analyze the relation between revenue and efficiency in GSP auctions. One might ask if it is possible to have optimal efficiency and optimal revenue in the same equilibrium. In other words, among all GSP equilibria is the revenue-maximizing one efficient? We give a negative answer to this question, showing that for some AdAuction instances, we can increase revenue by selecting an innefficient equilibrium. However, we give a natural sufficient condition so that the revenue-optimal equilibrium is efficient.

We define the cost of efficiency for a given click-through-rate as the ratio:

$$\textsc{CoE}(\alpha) = \max_{\mathbf{v}} \frac{\max_{\mathbf{b} \in \textsc{Nash}(\alpha,\mathbf{v})} \mathcal{R}(\mathbf{b})}{\max_{\mathbf{b} \in \textsc{EffNash}(\alpha,\mathbf{v})} \mathcal{R}(\mathbf{b})}$$

where Nash is the set of all bid profiles in Nash equilibrium and EffNash is the set of all efficient Nash equilibrium. In figure 2 we calculate for each $\alpha = [1, \alpha_2, \alpha_3]$, where α_i is an integer multiple of 0.01 we calculate the $\textsc{CoE}(\alpha)$. For all α calculated, we got $1 \leq \textsc{CoE}(\alpha) < 1.1$. The color of (α_1, α_2) in the graph corresponds to $\textsc{CoE}(1, \alpha_2, \alpha_3)$ 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.

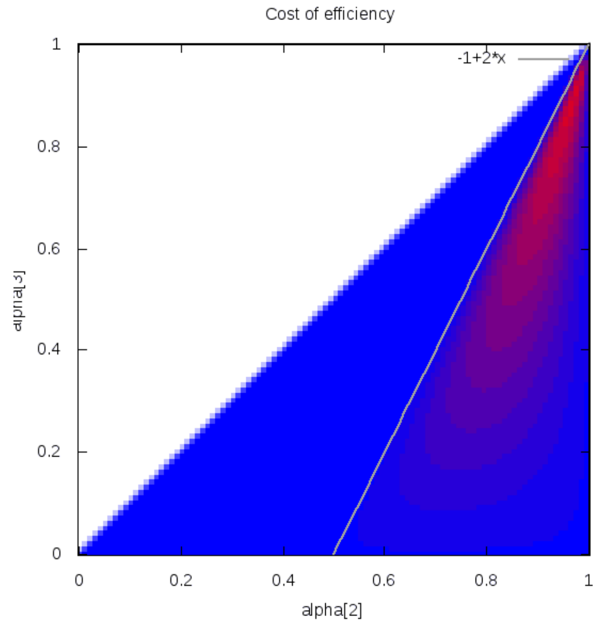

Figure 2: Cost of efficiency for $\alpha = [1, \alpha_2, \alpha_3]$: in the plot, blue means 1.0 and red means 1.1.

Example. One example where an innefficient equilibrium generates strictly more revenue then all efficient ones is $\alpha = [1, \frac{2}{3}, \frac{1}{6}]$ and $v = [1, \frac{7}{8}, \frac{6}{8}]$: the best efficient revenue is given by $\frac{1}{3} + \frac{7}{8} \approx 1.20833$ (which can be calculated using the formula in the last section), but for the allocation $\pi = [2, 1, 3]$ there is an equilibrium generating revenue: 1.21528. ∎

A remarkable fact we can observe from the graph is that when $\alpha_1 - \alpha_2 \geq \alpha_2 - \alpha_3$, then $\textsc{CoE}(\alpha) = 1$, what motivates us to look at AdAuctions instances with convex click-through-rates, i.e., $\alpha_i - \alpha_{i+1} \geq \alpha_{i+1} - \alpha_{i+2}$. Notice that this is a natural assumption, since most models for CTR follow convexity, as exponential CTR as in [11], Markovian users [1], ... The main theorem is this section shows that this is a sufficient condition for having $\textsc{CoE} = 1$:

Theorem 9 *If click-through-rates α are convex (i.e. $\alpha_i - \alpha_{i+1} \geq \alpha_{i+1} - \alpha_{i+2}, \forall i$), then there is a revenue maximizing equilibrium that is also efficient.*

Proof. Let $\mathbf{b}$ be the revenue maximizing Nash equilibrium, which can be calculated according to the formula in the last section. Now, 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 $b_i = v_i$. First we prove the theorem if no slot is saturated in the maximum revenue equilibrium. This part will be simpler and capture the main spirit of the proof. Then we prove the general case, which is more technical.

Under the no-saturation assumption, then:

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

Notice that we can see it is a dot product of two vectors where one has elements of the for v_i and other has elements in the form $\alpha_j - \alpha_{j+1}$. Notice also that due to the convexity assumption, we can think of it as a dot product of two sorted vectors. Now, for $\mathbf{b}'$, we can use the fact that no player in slot i wants to take some slot $j > i$ to bound:

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

where $m(\pi, i, j) = \max\{\pi(i), \pi(i+1), \pi(i+2), \ldots, \pi(j)\}$. To see that, let $k = i, i+1, \ldots, i+p$ be 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 $i+p+1$, therefore:

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

so:

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

Now, we just apply recursion. Now, 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 and 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 that, simply count how many times we have one of $v_1, \ldots, v_i$ appear in both vectors for each index i. For equation (2) they appear $\sum_{j=1}^{i} j$ times. For equation (3), they appear at most:

$$\sum_{j=1}^{i} 1 + \max\{p; m(\pi(j, j+p)) \leq i\} \leq \sum_{j=1}^{i} j$$

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, then it must be the case that $\mathcal{R}(\mathbf{b}) \geq \mathcal{R}(\mathbf{b}')$.

Case with saturations: Now, let $\mathbf{b}$ be the optimal efficient equilibrium and let $S \subseteq [n+1]$ be the set of saturated bids and $n+1$ (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

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

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).

Claim 10 *If bid profile* $\mathbf{b}$ *induces allocation* π *at equilibrium, then* $\mathcal{R}(\mathbf{b}) \leq \overline{\mathcal{R}}_\pi$.

Claim 11 *There exists an efficient equilibrium with revenue* $\overline{\mathcal{R}}_{id}$.

We want to argue that id is the permutation that maximizes $\overline{\mathcal{R}}_\pi$ and therefore we can show that for all innefficient bid profile $\mathbf{b}'$ we have:

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

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. For the other players. 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 12 $\overline{\mathcal{R}}_{\pi'} \geq \overline{\mathcal{R}}_\pi$.

■

Proof of Claim 10 : We need to show that for all $\mathbf{b}'$ inducing allocation π, then $\alpha_{\sigma(j)} b'_{\sigma(j)+1} \leq B_\pi(j)$. For $\sigma(j) = S(j) - 1$, we simply use the fact that $b'_{\sigma(j)+1} = b'_{S(j)} \leq v_{S(j)}$. Now, for $\sigma(j) < S(j) - 1$ we simply use the same proof used in the unsaturated case. For $\sigma(j) > S(j) - 1$, we use the fact that player j doesn't want to take slot j and therefore:

$$\alpha_{\sigma(j)} (v_j - b'_{\sigma(j)+1}) \geq \alpha_{S(j)-1} (v_j - b'_{S(j)-1}) \geq \alpha_{S(j)-1} (v_j - v_{S(j)})$$

since

$$b'_{S(j)} \leq \min\{v_{\pi(1)}, \ldots, v_{\pi(S(j)-1)}\} \leq v_{S(j)}$$

and $\sigma(j) > S(j) - 1$ so one of the players with value $\leq v_{S(j)}$ must be among the first $S(j) - 1$ slots. Reordering the Nash inequalities above gives us the desired result. ■

Proof of Claim 11 : That is true simply by the formula defininig the optimal-revenue efficient-equilibrium in the last section and the definition of saturation. ■

Proof of Claim 12 : Note first that $B_\pi(k) = B_{\pi'}(k)$ for all $k > j$. Moreover, for any k with $\sigma(k) < \sigma(j)$, we will have $\sigma'(k) = \sigma(k)$. In this case, either $S(k) < \sigma(k)$ in which case $B_{\pi'}(k) = B_\pi(k)$, or else

$$\begin{aligned} B_{\pi'}(k) &= \alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi,\sigma(k),i)} \\ &\geq \alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma'(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi',\sigma'(k),i)} \\ &= B_{\pi'}(k). \end{aligned}$$

It remains to consider k is such that $\sigma(j) \leq \sigma(k) \leq j$; that is, those players k such that $\sigma(k) \neq \sigma'(k)$. For each such player, we will consider the difference between $B_\pi(k)$ and $B_{\pi'}(k)$. First note that, for player j, we have

$$\begin{aligned}
&B_\pi(j) - B_{\pi'}(j) \\
&= \left(\alpha_{S(j)-1}v_{S(j)} + \sum_{i=\sigma(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1})v_{m(\pi,\sigma(j),i)}\right) \\
&\quad - \left(\alpha_{S(j)-1}v_{S(j)} + \sum_{i=\sigma'(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1})v_{m(\pi',\sigma'(j),i)}\right) \\
&= \sum_{i=\sigma(j)}^{j-1} (\alpha_i - \alpha_{i+1})v_j
\end{aligned}$$

For $k \neq j$, we claim that $B_{\pi'}(k) - B_\pi(k) \geq v_j(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})$. We proceed by two cases. First, if $S(k) \leq \sigma(k)$, we have

$$\begin{aligned}
B_{\pi'}(k) - B_\pi(k) &= \left(\alpha_{S(k)-1}v_{S(k)} - v_k(\alpha_{S(k)-1} - \alpha_{\sigma'(k)})\right) \\
&\quad - \left(\alpha_{S(k)-1}v_{S(k)} - v_k(\alpha_{S(k)-1} - \alpha_{\sigma(k)})\right) \\
&= v_k(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)}) \\
&\geq v_j(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})
\end{aligned}$$

Second, if $S(k) - 1 > \sigma(k)$, then we have

$$\begin{aligned}
&B_{\pi'}(k) - B_\pi(k) \\
&= \left(\alpha_{S(k)-1}v_{S(k)} + \sum_{i=\sigma'(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1})v_{m(\pi',\sigma'(k),i)}\right) \\
&\quad - \left(\alpha_{S(k)-1}v_{S(k)} + \sum_{i=\sigma(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1})v_{m(\pi,\sigma(k),i)}\right) \\
&= (\alpha_{S(k)-2} - \alpha_{S(k)-1})v_{m(\pi',\sigma'(k),S(k)-2)} \\
&\quad + \sum_{i=\sigma'(k)}^{S(k)-3} v_{m(\pi',\sigma'(k),i)}[(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2})] \\
&\geq v_j(\alpha_{S(k)-2} - \alpha_{S(k)-1}) \\
&\quad + \sum_{i=\sigma'(k)}^{S(k)-3} v_j[(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2})] \\
&= v_j(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})
\end{aligned}$$

Notice that we strongly use the fact that click-through-rates are convex in the last inequality to ensure that $(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2}) \geq 0$.

Therefore, taking the sum over all k with $\sigma(j) \leq \sigma(k) \leq j$, we have

$$\begin{aligned}
\sum_{k:\sigma(j)<\sigma(k)\leq j} (B_{\pi'}(k) - B_\pi(k)) &\geq \sum_{i=\sigma(j)}^{j-1} v_j(\alpha_i - \alpha_{i+1}) \\
&= B_\pi(j) - B_{\pi'}(j)
\end{aligned}$$

So that

$$\sum_{k:\sigma(j)\leq\sigma(k)\leq j} (B_{\pi'}(k) - B_\pi(k)) \geq 0.$$

Combining this with the fact that $B_{\pi'}(k) \geq B_\pi(k)$ for all k with $\sigma(k) < \sigma(j)$ or $\sigma(k) > j$, we conclude

$$\overline{\mathcal{R}}_{\pi'} = \sum_k B_{\pi'}(k) \geq \sum_k B_\pi(k) = \overline{\mathcal{R}}_\pi$$

as desired. ∎

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] S. Athey and D. Nekipelov. Equilibrium and uncertainty in sponsored search advertising. In *AdAuctions Workshop 2010*, 2010.

[3] T. Borgers, I. Cox, M. Pesendorfer, and V. Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. Technical report, 2007.

[4] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, (43):192–198, 2007.

[5] 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.

[6] B. Edelman and M. Schwarz. Optimal auction design in a multi-unit environment: The case of sponsored search auctions. Technical report, 2006.

[7] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, (2), 2010.

[8] J. Feng, H. K. Bhargava, and D. M. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS J. on Computing*, 19(1):137–148, 2007.

[9] 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.

[10] J. Hegeman. Facebook's ad auction. Talk at Ad Auctions Workshop, May 2010, 2010.

[11] 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.

[12] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.

[13] S. Lahaie and D. M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 50–56, New York, NY, USA, 2007. ACM.

[14] R. P. Leme and Éva 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.

[15] B. Lucier and R. Paes Leme. Improved social welfare bounds for gsp at equilibrium (new title: Gsp with correlated types). Technical report, arXiv.org, 2010.

[16] R. Myerson. Optimal auction design. pp 58-73, Mathematics of Operations Research, Vol 1/6, pp 58-73, 1981.

[17] H. Varian. Mechanism design double feature. Talk at New York Computer Science and Economics Day, October 2010, 2010.

[18] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.

[19] H. R. Varian. Online ad auctions. *American Economic Review Papers and Proceedings*, 2009.