Revenue Enhancement in Ad Auctions

Michal Feldman[1,2], Reshef Meir[1,2], and Moshe Tennenholtz[1,3]

[1] Microsoft Research
[2] Hebrew University of Jerusalem
[3] Technion-Israel Institute of Technology

Abstract. We consider the revenue of the Generalized Second Price (GSP) auction, which is one of the most widely used mechanisms for ad auctions. While the standard model of ad auctions implies that the revenue of GSP in equilibrium is at least as high as the revenue of VCG, the literature suggests that it is not strictly higher due to the selection of a natural equilibrium that coincides with the VCG outcome. We propose a randomized modification of the GSP mechanism, which eliminates the low-revenue equilibria of the GSP mechanism under some natural restrictions. The proposed mechanism leads to a higher revenue to the seller.

1 Introduction

Ad auctions are perhaps the most widely studied economic setup in the literature on on-line markets. As ad auctions generate revenues of billions of dollars per year to publishers, every subtle feature of their design may have tremendous effect. The two most widely discussed mechanisms in the study of ad auctions are the Generalized Second Price (GSP) auction, versions of which are those typically used in practice, and the classical Vickrey-Clarke-Groves (VCG) auction. We consider the original ad auctions model introduced in the seminal work of Varian [12] and Edelman et al. [3], where bidders' valuations per click are fixed and independent of the ad slot. Previous work characterized a special family of equilibria of GSP auctions in that setting, termed *envy free* or *Symmetric* Nash Equilibria (SNE), and showed that the SNE leading to the lowest revenue for the seller, termed Lower Equilibrium [LE], coincides with the truth-revealing equilibrium of VCG. These connections between the GSP equilibria and the VCG outcome, as well as the related revenue ramifications, are central to the study of ad auctions.

While the above results suggest that GSP may lead to higher revenue than VCG, other arguments for revenue comparison between these mechanisms have been discussed. Kuminov and Tennenholtz model user behavior explicitly as glancing through the ads in a sequence [5]. Interestingly, in this setting the VCG outcome coincides with the GSP equilibrium that leads to the *highest* revenue. Closer to the study of the standard ad auctions setting, Edelman and Schwarz consider equilibrium selection in GSP by comparing the revenue in the standard static game to a dynamic variant of the game [4]. They suggest that the GSP equilibrium that leads to the highest revenue is less natural than the one that coincides with the VCG outcome since it generates too high revenue compared to the revenue obtained in the dynamic model. Recently, Lucier et al. [6] performed a detailed analysis of the revenue in the GSP auction, under both complete and incomplete information, using the VCG revenue as a baseline. In particular, they outlined the conditions under which the revenue in non-envy-free equilibria can be lower or higher than the revenue in SNE.

Thompson and Leyton-Brown [11] computed the revenue of GSP in equilibrium using several models from the literature. They found that while the expected revenue of GSP in Varian's model was slightly higher than the VCG baseline, in most models the revenue was profoundly affected by equilibrium selection.

The above suggests that GSP has attracted much attention from both researchers and practitioners, but it is unclear whether it has revenue advantage over VCG. Hence, one may wish to consider natural modifications for GSP that increase the auctioneer/publisher revenue. Notice that GSP is by now a standard practice and modifications to it should conform to having a relatively similar structure in the way bidders are assigned to ad slots and the way they are assigned payments; i.e., an advertiser's payment should be bounded by his bids and have some intuitive relations with other (less successful) bids. This is not a "mathematical" requirement, but a practical one, given the way advertisers perceive ad auctions.

Several modifications to the GSP mechanism have already been suggested in the literature. The most common modification to GSP is the addition of *reserve prices*. Indeed, field experiments (and to some extent also theory) suggest that reserve prices can substantially increase the revenue in GSP auctions [9]. A different modification deals with allocation efficiency. When some assumptions in Varian's model are violated, the GSP mechanism may not have efficient equilibria. Blumrosen et al. modify the GSP mechanism to guarantee the existence of an efficient equilibrium in a more general model [1].

Our contribution In this paper we take the study of the GSP revenue to the next stage, by suggesting natural modifications to the GSP mechanism which result in revenue boosting for many natural click-through rates. Recall that in the GSP mechanism every winning agent pays the *second* price, i.e., the bid of the bidder directly below her bid. Our revised mechanism selects randomly between GSP and a variant of it, in which an agent pays the *third* price, i.e., the bid that is just below the bid below her bid. [1] We show conditions under which the combined mechanism admits an ex-post envy-free equilibrium that achieves revenue that is arbitrarily close to the revenue obtained in the highest revenue equilibrium of GSP, while eliminating the low revenue equilibria outcomes. More generally, we introduce the family of m-price auctions by generalizing the GSP auctions as well as the random selection among them, study their ex-post envy-free equilibria, and prove that by random selection between a pair of such mechanisms we can boost the revenue of GSP.

[1] k-price auctions have been shown to lead to some intriguing results in the more classical single-items setting; see [7, 8, 10]

2 Model and Preliminaries

2.1 Ad auctions

In an ad auction there are s slots to allocate, and n bidders, each with valuation v_i per click. Every slot $1 \leq j \leq s$ is associated with a click-through rate (CTR) $x_j > 0$, where $x_j \geq x_{j+1}$. For mathematical convenience, we define $x_j = 0$ for every $j > s$, and similarly $b_i = 0$ for all $i > n$. Throughout the paper we make the simplifying assumption that CTRs are strictly decreasing, i.e., $x_j > x_{j+1}$.

The mechanism receives as input a bid b_i from every bidder and determines an allocation of the slots to the bidders and a payment per click, p_i, for every bidder. We denote the slot allocated to bidder i by $\pi(i)$. A bidder i that has been allocated slot $\pi(i)$ gains v_i per click (regardless of the slot), and is charged p_i per click. Thus, his total utility is given by $u_i = (v_i - p_i)x_{\pi(i)}$. A mechanism that assigns better slots to higher bids is called *efficient*. A mechanism that never assigns lower payments to higher bids is called *monotone*. We restrict our attention to efficient and monotone mechanisms.

Nash equilibria. Let $v_1 > \ldots > v_n$ and $b_1, \ldots, b_n$ be the bidders' private values and submitted bids, respectively (both sorted according to decreasing valuations). Let f be an efficient and monotone auction mechanism, and let $p_1, \ldots, p_n$ be the payments assigned by f according to the bids. We say that the bids are in Nash equilibrium (NE), if no bidder wants to deviate by changing her bid. Let $p_f(b')$ be the payment assigned by f to bidder i if she changes her bid to b' (and all other bidders keep their current bids). From the efficiency of f, in order to get slot $j < \pi(i)$, bidder i must bid like the bidder currently occupying the slot, i.e., $b' = b_j$.[2] Similarly, in order to get slot $j > \pi(i)$, bidder i must bid below bidder j, i.e., $b' = b_{j+1}$. The stability requirements of NE can be divided into three parts. For every bidder i the following should hold:

bid neutral Bidder i has no incentive to change her bid if this deviation does not change the slot assigned to i.

up-Nash Bidder i does not want to get a better slot:

$$\forall j < \pi(i), \quad (v_i - p_i)x_{\pi(i)} \geq (v_i - p_f(b_{\pi^{-1}(j)}))x_j. \tag{1}$$

down-Nash Bidder i does not want to get a worse slot:

$$\forall j > \pi(i), \quad (v_i - p_i)x_{\pi(i)} \geq (v_i - p_f(b_{\pi^{-1}(j+1)}))x_j. \tag{2}$$

The first requirement is usually handled by mechanisms that ignore b_i when setting the payment p_i (i.e., b_i is only used to decide on the allocation and on p_j for $j \neq i$).

Efficiency A priori, even with an efficient mechanism we might end up with an inefficient outcome, where a bidder with low valuation bids higher than a bidder with higher valuation, thereby getting a better slot. We say that an equilibrium outcome is *efficient*, if $b_i \geq b_{i+1}$ for all i. Note that if both the mechanism f and the outcome $\mathbf{b}$ are efficient, then every bidder i gets slot i, i.e., $\pi(i) = i$. Efficient equilibria guarantee that the *social welfare*, i.e., the sum of utilities of all bidders and the auctioneer, is maximized.

[2] To guarantee that the deviator gets slot j, we assume that ties are broken in favor of the deviator. This assumption will not be required, however, when we analyze the GSP mechanism and its variations, as bidder i can bid $b_{j-1} > b' > b_j$, without affecting her own payment.

Envy freeness. A different notion of stability than Nash is captured by the *envy-freeness* requirement. An outcome is envy free if no bidder is interested in swapping slots (and payments) with any other bidder. Formally, it takes the following form:

$$\forall j \neq \pi(i), \quad (v_i - p_i)x_{\pi(i)} \geq (v_i - p_f(b_{\pi^{-1}(j+1)}))x_j. \tag{3}$$

It is easy to verify that (3) entails requirements (1) and (2). Thus, any envy-free outcome in a bid-neutral mechanism is also a NE. In fact, envy freeness is more restricting than (1), and thus we are left with a subset of the original set of NE.

Envy-free equilibria have been thoroughly studied in the literature of games in general, and ad auctions in particular. They are also known as *symmetric Nash equilibria* (SNE), due to the symmetry of up-Nash and down-Nash constraints. Varian [12] and Lucier et al. [6] further studied properties of SNEs in the GSP mechanism. For example, it is shown that every SNE is efficient, which is not true for arbitrary NE.

When a randomized mechanism is in use, we must distinguish between outcomes that are *envy free in expectation* from outcomes that are *envy free ex-post*. The latter definition means no bidder wants to change slots, even after the randomization has taken place and the outcome is known. We will be interested in this stronger interpretation of envy freeness.

The revenue interval. Suppose we are using some auction mechanism f. If f has multiple NE outcomes, then the auctioneer might end up with different revenues for the different equilibria. We define R_f^U (resp. R_f^L) as the highest (resp. lowest) revenue generated by mechanism f in some NE. We use a similar notation to denote the highest and lowest revenues in restricted subsets of NE, replacing R with ER (for *efficient NE*) or with SR (for *symmetric NE*). Clearly $[SR_f^L, SR_f^U], [ER_f^L, ER_f^U] \subseteq [R_f^L, R_f^U]$.

The revenue interval raises the natural question of *equilibrium selection*. Clearly, the auctioneer would like the bidders to end up playing an equilibrium with high revenue. However, the auctioneer is not a player in this game. The players are the bidders, and given an efficient allocation their joint incentive is basically the opposite - to end up paying the lowest possible amount.

2.2 VCG

The VCG mechanism sorts the bidders by their bids, and allocates the j'th slot, $j = 1, \ldots, s$, to the j'th highest bidder. Each bidder j is charged (per click) for the "harm" she poses to the other bidders, i.e., the difference between the welfare of bidders $k \neq j$ if j is omitted and their welfare when j exists; thus $p_j = \sum_{k \geq j+1} b_k(x_{k-1} - x_k)$. Note that VCG is efficient and monotone. It is well known that under the VCG mechanism, reporting the true values (i.e., $b_i = v_i$) is a dominant strategy, and in particular it is a Nash equilibrium. We denote the revenue in the truthful equilibrium of VCG by R_{VCG}^T.

2.3 GSP

The allocation of the GSP mechanism is efficient, i.e., identical to that of VCG. The charge of bidder $j = 1, \ldots, s$ equals the bid of the next bidder; i.e., $p_j = b_{j+1}$. For mathematical convenience, we define $b_{j+1} = 0$ for $j \geq n$. GSP is clearly efficient and monotone.

Varian [12] focuses on the analysis of envy-free equilibria (i.e., SNEs) in the GSP auction, due to their many attractive properties. The SNE requirement (3) takes the following form:

$$\forall j \neq i, (v_i - b_{i+1})x_i \geq (v_i - b_{j+1})x_j. \qquad (4)$$

As Varian shows, every SNE is an efficient equilibrium.

2.4 Known properties of GSP

Before presenting some known properties of the VCG and GSP mechanisms, we put forward the following basic definitions. Let $g_1, \ldots, g_m \in \mathbb{R}_+$ be the elements of a monotonically nonincreasing series. We sometimes refer to such series as *functions* (of the form $g : \{1, \ldots, m\} \to \mathbb{R}$). We say that g is *convex* if it has a decreasing marginal loss, i.e., $g_i - g_{i+1} \geq g_j - g_{j+1}$ for every $i < j$. Similarly, if g has an *increasing* marginal loss then it is *concave*. Notice that linear functions are both convex and concave.

A special case of convexity is when the marginal loss decreases exponentially fast. We say that g is *β-separated* (for some $0 < \beta < 1$) if $g_{i+1} \leq \beta g_i$ for every i. If the above holds with equality (rather than inequality), g is said to be *exponential*.

The last property refers to the *ratio* between consequent elements, rather than difference. A series/function g is said to be *log-concave*, if for every $i < j \leq m$, $\frac{g_i}{g_{i-1}} \geq \frac{g_j}{g_{j-1}}$.

Distinguished equilibria Of particular interest are the two equilibria of GSP that reside on the boundaries of the SNE set, referred to as *Lower Equilibrium* (LE) and *Upper Equilibrium* (UE). We denote the LE and UE profiles by $\mathbf{b}^L = \{b_i^L\}_{i \in N}$ and $\mathbf{b}^U = \{b_i^U\}_{i \in N}$, respectively. The bids in the LE, for every $2 \leq i \leq s$, are given by

$$b_i^L x_{i-1} = v_i(x_{i-1} - x_i) + b_{i+1}^L x_i = \sum_{s+1 \geq t \geq i} v_t(x_{t-1} - x_t).$$

This equilibrium induces payments, utilities, and revenue equal to those in VCG.

Proposition 1 (Varian [12]) *The payments of all bidders in the LE of GSP are the same as under the truthful bidding in VCG. In particular, it follows that $SR_{GSP}^L = R_{VCG}^T$.*

The bids in the UE, for every $2 \leq i \leq s$, are given by

$$b_i^U x_{i-1} = v_{i-1}(x_{i-1} - x_i) + b_{i+1}^U x_i = \sum_{s+1 \geq t \geq i} v_{t-1}(x_{t-1} - x_t).$$

Note that since $x_{i-1} > x_i$, no two bidders submit the same bid in LE (or in UE).

Another outcome of GSP which one may wish to consider is obtained by truthful bidding (i.e., $b_i = v_i$), which may or may not be an equilibrium. We further discuss truthful bidding in Section 3.3.

Since SNE is always efficient, we have that

$$[SR_{GSP}^L, SR_{GSP}^U] \subseteq [ER_{GSP}^L, ER_{GSP}^U] \subseteq [R_{GSP}^L, R_{GSP}^U],$$

where in the most general case, these inclusions may be strict.

Lucier et al. [6] study the conditions on the CTR function under which the boundaries of these sets become close or equal. They show that if $x_i > x_{i+1}$ for all i (as we assume here), then the first inclusion becomes an equality. That is, we can get a revenue r in an efficient equilibrium iff there is an SNE with revenue r.

Justifying the Lower equilibrium. We argue that when faced with the equilibrium selection problem of GSP, bidders are likely to play the lower envy-free equilibrium $\mathbf{b}^L$. The concept of envy freeness itself is well justified in various settings. In addition, in our ad auction setting, bidders have a particular interest in an efficient allocation, which makes the set of SNEs even more prominent. Within this set, there are two natural focal points, introduced in the previous paragraph.

Observation 1 *Let* $\mathbf{b}$ *be an SNE other than LE; then* $\mathbf{b}$ *is Pareto dominated by* $\mathbf{b}^L$.

To see this, observe that if all bidders change their bids from b_i to b_i^L, then the allocation remains the same, and every bidder i pays the same, or strictly less (if $b_{i+1}^L < b_{i+1}$).

To sum up: if bidders try to influence the outcome equilibrium, they are most likely to play the lowest envy-free equilibrium (LE in the GSP case),[3] and the revenue of the auctioneer will be $SR_{GSP}^L = ER_{GSP}^L$.

3 Boosting the revenue

In this section we propose a variant of the GSP mechanism that eliminates the LE profile (but not the UE profile) as an envy-free equilibrium, thereby incentivizing the bidders to end up in higher-revenue equilibrium.

3.1 Generalized Next Price auctions

Consider a modified GSP mechanism, termed m-price auction, in which $p_i = b_{i+m-1}$. For example, the 2-Price auction is GSP, as $p_i = b_{i+1}$. We argue that the SNE outcomes of the auction are essentially the same for all $m \geq 2$.

Lemma 1. *Let* $\mathbf{b}$ *be a (sorted) bid vector, and let* $k, m \geq 2$ *s.t.* $k < m \leq n + 1 - s$. $\mathbf{b}$ *is an SNE of* m*-price if and only if* $\mathbf{b}'$ *is an SNE of* k*-price, where* $b'_i = b_{i+k-m}$.[4]

Proof. The two bid vectors in both mechanisms induce exactly the same allocation (in decreasing order of valuations), and exactly the same payments. This is because $p_i(m) = b_{i+m-1} = b'_{i+k-1} = p_i(k)$ (note that the payment is well defined). Thus, if in one auction there is a bidder i envious of a bidder j, then the same bidder is envious in the other auction. Since $k, m \geq 2$, the payment p_i does not depend on the bid b_i itself, thus no envy also entails Nash equilibrium, i.e., SNE. □

[3] We emphasize that we do not consider *collusion*, i.e., binding contracts among agents.

[4] If $i + k - m < 1$ then b'_i can be completed arbitrarily, as long as bids' order is kept, i.e., $b'_i > b'_{i+1}$.

In particular, from the perspective of the bidders the outcome is the same whether GSP or any other m-price auction is used (although they will submit different bids). Clearly the revenue of the auctioneer is the same as well.

Due to Lemma 1, we can easily derive the lower and upper bounds on SNE bids in m-price auctions, which we denote by $b_i^L(m), b_i^U(m)$. Indeed, for every $m \geq 2$,

$$b_i^L(m) = b_{i+2-m}^L = \sum_{t=i+2-m}^{s+1} v_t \frac{x_{t-1} - x_t}{x_{i+1-m}},$$

and similarly for the upper equilibrium.

We will focus on 2-price (i.e., GSP) and 3-price auctions; we refer to the latter as the *generalized third-price* auction [GTP].

3.2 Boosting revenue via randomization

Due to Lemma 1, it seems that there is no benefit in using the m-price auction rather than the original GSP. Quite surprisingly, it turns out that combining these mechanisms enables us to improve the revenue interval.

We introduce a randomized mechanism that boosts the auctioneer's revenue. To this end, we identify a bid profile that constitutes an SNE in two mechanisms simultaneously, enabling us to preserve the envy-free equilibrium when we choose randomly between them (the proof is given below, together with the proof of Theorem 3).

Theorem 2. *Suppose that the CTR function is convex* and *log-concave, and for every* $2 \geq i \geq s$ *it holds that* $\frac{v_i}{v_{i-1}} \leq \frac{x_i}{x_{i-1}}$. *Then the UE of GSP is also an SNE of GTP.*

Note that there is a wide family of log-concave functions that are also convex. These include linear, polynomial, quasi-polynomial, and exponential functions. The more delicate condition is the requirement on the valuation function, i.e., that it be *steeper* than the CTR function. While this may not always hold, such steepness is a likely assumption in domains where bidders are heterogeneous.

Log-concavity is a necessary condition. That is, if the CTR function is *not* log-concave, then there is no bid profile that is simultaneously envy-free in both auctions. Further, we can sharpen our conditions by focusing on the "least log-concave" function.

Theorem 3. *If the CTR function is exponential, then the UE of GSP is equal to the LE of GTP. This is true for arbitrary valuations.*

This means that (with exponential CTR) there is a *unique* bid profile that is simultaneously envy-free in both auctions.

Proof. By Varian's paper (see [12], p. 1167), the bids $\mathbf{b}^*$ are SNE iff

$$\alpha_i b_{i+1}^* + (1 - \alpha_i)v_{i-1} \geq b_i^* \geq \alpha_i b_{i+1}^* + (1 - \alpha_i)v_i,$$

where $\alpha_i = \frac{x_i}{x_{i-1}} < 1$. For example $b_i^U(2) = \alpha_i b_{i+1}^U(1) + (1 - \alpha_i)v_{i-1}$.

Similarly, for GTP, we get from Lemma 1 that the bids $\mathbf{b}^*$ are SNE iff

$$\alpha_{i-1} b_{i+1}^* + (1 - \alpha_{i-1})v_{i-2} \geq b_i^* \geq \alpha_{i-1} b_{i+1}^* + (1 - \alpha_{i-1})v_{i-1}.$$

or, equivalently, if

$$x_{i-1}b^*_{i+1} + (x_{i-2} - x_{i-1})v_{i-2} \geq x_{i-2}b^*_i \geq x_{i-1}b^*_{i+1} + (x_{i-2} - x_{i-1})v_{i-1}. \quad (5)$$

We can substitute $\mathbf{b}^*$ for any given set of bids (e.g., $\mathbf{b}^U$) to see when these bids form an equilibrium.

The lower bound holds if

$$\alpha_i b^U_{i+1} + (1-\alpha_i)v_{i-1} = b^U_i \geq \alpha_{i-1}b^U_{i+1} + (1-\alpha_{i-1})v_{i-1} \iff$$
$$1-\alpha_i \geq 1-\alpha_{i-1} \iff$$
$$\frac{x_i}{x_{i-1}} = \alpha_i \leq \alpha_{i-1} = \frac{x_{i-1}}{x_{i-2}},$$

i.e., when $\mathbf{x}$ is log-concave. The first transition holds because it is a convex combination of b^U_{i+1} and v_{i-1}, where the latter is larger.

For exponential CTR functions, α_i is a constant and we get an equality with the lower bound. In particular, the upper bound holds as well. Moreover, any lower bids will not be SNE in GTP, as they must violate the lower bound. This proves Theorem 3.

For other log-concave CTR functions the inequality may be strict, and therefore we must verify that the upper bound holds as well. Under our assumptions on $\mathbf{x}$ and $\mathbf{v}$,

$$\begin{aligned}
&(x_{i-1}b^U_{i+1} + (x_{i-2} - x_{i-1})v_{i-2}) - (x_{i-2}b^U_i) \\
&\quad = x_{i-1}b^U_{i+1} + (x_{i-2} - x_{i-1})v_{i-2} - x_{i-2}\frac{b^U_{i+1}x_i + (x_{i-1} - x_i)v_{i-1}}{x_{i-1}} \\
&\quad = b^U_{i+1}\left(\frac{x_{i-2}x_i}{x_{i-1}} - x_{i-1}\right) + (x_{i-2} - x_{i-1})v_{i-2} - \frac{x_{i-2}}{x_{i-1}}(x_{i-1} - x_i)v_{i-1} \\
&\quad \geq (x_{i-2} - x_{i-1})v_{i-2} - \frac{x_{i-2}}{x_{i-1}}(x_{i-1} - x_i)v_{i-1} && (6) \\
&\quad \geq v_{i-2} - \frac{x_{i-2}}{x_{i-1}}v_{i-1} && (7) \\
&\quad \geq \frac{x_{i-2}}{x_{i-1}}v_{i-1} - \frac{x_{i-2}}{x_{i-1}}v_{i-1} = 0, && (8)
\end{aligned}$$

and thus $x_{i-1}b^U_{i+1} + (x_{i-2} - x_{i-1})v_{i-2} \geq x_{i-2}b^U_i$. Inequality (6) follows from log-concavity of $\mathbf{x}$, (7) from convexity, and (8) from the steepness of $\mathbf{v}$. □

In addition to showing that the UE profile is an SNE in both auctions, we also want to show that worse profiles (for the auctioneer) are not. Theorem 3 guarantees that this is the case with exponential CTRs, but in the less restricted case of Theorem 2 there may be other profiles with this property. However, we show that any profile with this property must include bids that are strictly above the LE profile.

Observation 2 *The LE of GSP is* not *an SNE of GTP.*

This follows directly from Lemma 1. The lower SNE bound of bidder i in GTP (i.e., $b^L_i(2)$) equals the LE bid of bidder $i-1$ in GSP, and is therefore strictly higher than b^L_i. Thus b^L_i violates the SNE constraints of GTP (for all bidders).

Finally, we bring together the results above in order to construct a mechanism that beats VCG in terms of revenue. This mechanism, denoted $M(q)$, runs GSP w.p. q (where $q \in [0,1]$) and otherwise runs GTP.

Algorithm 1 THE RANDOM NEXT-PRICE MECHANISM ($M(q)$)

Collect bids from all agents.
Allocate slots according to bids in decreasing order.
 w.p. q, each bidder i pays b_{i+1}. *// GSP is applied*
 w.p. $1-q$, each bidder i pays b_{i+2}. *// GTP is applied*

Recall our discussion from the Introduction, where we argued that the most reasonable bid profile that will be played is the envy-free equilibrium with the lowest payments (and lowest revenue). For the GSP mechanism, that was the LE profile. The next result follows immediately from Theorem 2 and Observation 2.

Corollary 1. *Suppose that the conditions of Theorem 2 hold; then* $\mathbf{b}^U = \{b_i^U\}_{i \in N}$ *is an ex-post envy-free Nash equilibrium of* $M(q)$*, for all* $0 < q < 1$*. Moreover, for any ex-post envy-free Nash equilibrium* $\mathbf{b}$*, we have* $b_i \geq b_i^L$ *for all* $i \leq s+2$*, with at least one strict inequality.*

Similarly, from Theorem 3, we get the following:

Corollary 2. *If CTRs are decreasing exponentially, then* $\mathbf{b}^U$ *is the* unique *ex-post envy-free Nash equilibrium of* $M(q)$*, for all* $0 < q < 1$.

We conclude that in either case the expected revenue of the $M(q)$ mechanism exceeds that of GSP and VCG, and with exponential CTR it is approaching the highest possible revenue of the GSP auction (in any Nash equilibrium).

Theorem 4. *Under the conditions of either Theorem 2 or 3, for any* $0 < q < 1$ *it holds that* $SR^L_{M(q)} > SR^L_{GSP}$ *(and thus* $SR^L_{M(q)} > R^T_{VCG}$*). Moreover, under the conditions of Theorem 3,*

$$\lim_{q \to 1} SR^L_{M(q)} = SR^U_{GSP} = ER^U_{GSP}.$$

Proof. By Corollary 1, there exists a bidding profile $\mathbf{b}$ that is ex-post envy-free in $M(q)$. The payment of bidder i is given by $q(b_{i+1}) + (1-q)(b_{i+2})$. Since $b_{i'}$ cannot be lower than the LE bid of GTP (or otherwise it would not be envy-free), it follows that $b_{i'} \geq b^L_{i'}(3)$. By Lemma 1, $b^L_{i'}(3) = b^L_{i'-1}$. Therefore,

$$\mathbb{E}[p_i] \geq q b^L_{i+1}(3) + (1-q) b^L_{i+2}(3) = q b^L_i + (1-q) b^L_{i+1} > b^L_{i+1} = p^L_i.$$

That is, every bidder is paying (in expectation) strictly more in the $M(q)$ mechanism than in GSP. Further, when the CTRs are exponential, the payment is

$$\mathbb{E}[p_i] = q b^L_{i+1}(3) + (1-q) b^L_{i+2}(3) = q b^U_{i+1} + (1-q) b^U_{i+2} = q p^U_i + (1-q) p^U_{i+1},$$

and as q gets closer to 1, the revenue gets arbitrarily close to SR^U_{GSP}, i.e., to the highest revenue possible in the GSP mechanism *in envy-free equilibrium*.

Finally, since exponential CTR is in particular strictly decreasing, it follows (according to Lucier et al. [6]) that we get that $SR^U_{GSP} = ER^U_{GSP}$, as required. □

As a concluding remark, note that by Lemma 1 we can in fact randomize between any pair of mechanisms m-price and $(m+1)$-price, for $m \geq 2$ (provided that there are at least $s+m-1$ bidders). This is because the ranges of equilibrium payments (and revenues) will coincide with those of GSP and GTP, respectively.

3.3 Truthful bidding

A possible focal point for bidders might be to bid their true values, if this happens to be an equilibrium. We ask when is truthful bidding an SNE (i.e., envy-free equilibrium) of either of the mechanisms we studied. It always holds that $v_i \geq b^L_i$, since the latter is a convex combination of v_i and a lower value (b^L_{i+1}). Thus truthful bidding is never too low. According to Varian [12], SNE bids must satisfy the following:

$$\begin{aligned}
&b_i x_{i-1} \leq v_{i-1}(x_{i-1} - x_i) + b_{i+1}x_i && \iff \\
&v_i x_{i-1} \leq v_{i-1}(x_{i-1} - x_i) + v_{i+1}x_i && \\
&\quad = v_{i-1}x_{i-1} - x_i(v_{i-1} - v_{i+1}) && \iff && (b_i = v_i) \\
&x_i(v_{i-1} - v_{i+1}) \leq x_{i-1}(v_{i-1} - v_i) && \iff \\
&\frac{x_i}{x_{i-1}} \leq \frac{v_{i-1} - v_i}{v_{i-1} - v_{i+1}}
\end{aligned}$$

We get that truth-telling is an SNE in GSP, iff the last inequality occurs.

For truthful bidding to be an SNE of GTP there is an additional constraint (Eq.(5)):

$$\begin{aligned}
&x_{i-2}v_i \geq x_{i-1}v_{i+1} + (x_{i-2} - x_{i-1})v_{i-1} && \iff \\
&x_{i-1}(v_{i-1} - v_{i+1}) \geq x_{i-2}(v_{i-1} - v_i) && \iff \\
&\frac{x_{i-1}}{x_{i-2}} \geq \frac{v_{i-1} - v_i}{v_{i-1} - v_{i+1}}
\end{aligned}$$

Joining both constraints, we get that truthful bidding is ex-post envy-free (in the randomized mechanism) if and only if for all i

$$\alpha_{i-1} = \frac{x_{i-1}}{x_{i-2}} \geq \frac{v_{i-1} - v_i}{v_{i-1} - v_{i+1}} \geq \frac{x_i}{x_{i-1}} = \alpha_i. \tag{9}$$

Note that in particular this condition entails that the CTR function is log-concave, as in the conditions of Theorem 2.

The possible revenue ramification of this result is not straight-forward. On the one hand, truthful bidding is always above the lower envy-free equilibrium, and so it might

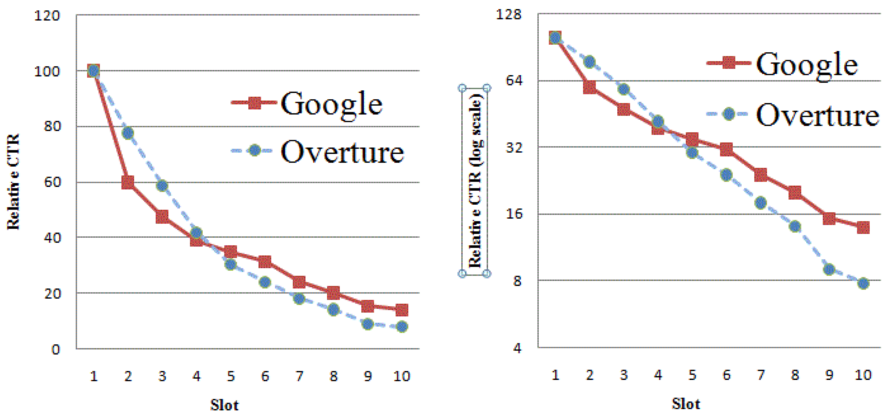

(a) Average CTR for each slot (linear scale) (b) Average CTR for each slot (log scale)

Fig. 1. The average click-through rate for ads positioned in any of the first ten slots are shown in Fig. 1(a) (numbers are normalized so that the CTR of slot 1 is 100). We can see that the shape of the CTR function is convex in both Google and Overture. In Fig. 1(b) we see the same data in log scale. Interestingly, the Overture CTR function is very close to exponential (with $\beta \cong 1.3$).

increase revenue (both in GSP and in the random next-price auction). On the other hand, if bidders need some focal point to coordinate (i.e., some natural set of bids), and truthful bids are *not an SNE* then the most likely outcome in the randomized auction is $\mathbf{b}^U$ (which leads to the maximal revenue). In such a case, stability of the truthful bids might hurt the revenue of the auctioneer.

4 Discussion

We proposed a randomized modification to the GSP auction, and analyzed the set of ex-post envy-free equilibria in this new mechanism. We showed natural conditions under which the revenue to the auctioneer strictly increases compared to the revenue in the "natural" envy-free equilibrium of the original GSP auction (which equals to the revenue in VCG under truthful bidding). When the CTRs are exponentially decreasing, our mechanism effectively eliminates all ex-post envy-free equilibria, except the one leading to the maximal possible equilibrium revenue of the GSP auction. It is important to note that convex and even exponential CTRs are common in the real world, as can be seen in Figure 1.[5]

Future research directions include the extension of our analysis to generalizations of the basic model (e.g., by considering the "ad's quality," as in [12]), and a study of the randomized mechanism in a model with incomplete information. Field experiments with the randomized mechanism (in the spirit of Ostrovsky and Schwarz [9]) will help to determine the practical value of our approach.

[5] Stats taken from Atlas Institutes rank report [2]. We present the data on the "click potential" attribute, which corresponds to the actual CTR in our model.

References

1. Liad Blumrosen, Jason D. Hartline, and Shuzhen Nong. Position auctions and non-uniform conversion rates. In *Proceedings of the 4th ACM EC Workshop on Ad Auctions*, 2008.
2. Nico Brooks. The Atlas rank report: How search engine rank impacts traffic. http://atlassolutions.com/uploadedFiles/Atlas/Atlas_Institute/Published_Content/RankReport.pdf.
3. Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
4. Benjamin Edelman and Michael Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, 100:597–602, 2010.
5. Danny Kuminov and Moshe Tennenholtz. User modeling in position auctions: re-considering the GSP and VCG mechanisms. In *AAMAS'09: Proceedings of the 8th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 273–280, 2009.
6. Brendan Lucier, Renato Paes Leme, and Éva Tardos. On revenue in the generalized second price auction. In *Proceedings of the 7th ACM EC Workshop on Ad Auctions*, 2011.
7. Dov Monderer and Moshe Tennenholtz. K-price auctions. *Games and Economic Behavior*, 31:220–244, 2000.
8. Dov Monderer and Moshe Tennenholtz. K-price auctions: revenue inequalities, utility equivalence, and competition in auction design. *Economic Theory*, 24(2):255–270, 2002.
9. Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *EC '11: Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 59–60, 2011.
10. Yair Tauman. A note on k-price auctions with complete information. *Games and Economic Behavior*, 41:161–164, 2002.
11. David R. M. Thompson and Kevin Leyton-Brown. Computational analysis of perfect-information position auctions. In *EC '09: Proceedings of the 10th ACM Conference on Electronic Commerce*, 2009.
12. Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.