

Cost of Conciseness in Sponsored Search Auctions

Zoë Abrams
Yahoo!, Inc.

Arpita Ghosh
Yahoo! Research

Erik Vee
Yahoo! Research

2821 Mission College Blvd.
Santa Clara, CA, USA
{za,arpita,erikvee}@yahoo-inc.com

Abstract

We study keyword auctions in a setting where bidders have a vector of values for slots, and a bidder's value for a slot is not necessarily proportional to the expected number of clicks in that slot. Specifically, bidders need not only derive values from clicks on their ad, nor do they need to value clicks in all slots equally. This is the most general model of bidder valuations (without externalities) in sponsored search, and encompasses a variety of advertising objectives, including conversions and branding.

We study the properties of the generalized second price auction when bidders with such vector values report a single scalar bid to the auction (we assume that bidders and the auctioneer agree on a ranking of the slots). We compare against the efficient, welfare maximizing outcome of the VCG mechanism with input as the full vector of values. Surprisingly, we find that there always exists an equilibrium corresponding to the VCG outcome, when bids and prices are per-impression in the generalized second price auction. If bids and prices are per-click, this need not be the case: in fact, an equilibrium need not even exist. However, under monotonicity conditions on the valuations of bidders and clickthrough rates, the VCG outcome is indeed an equilibrium of the system.

Finally, we discuss the problem of bidding strategies leading to such efficient equilibria: contrary to the case when bidder values are one-dimensional, bidding strategies with reasonable restrictions on bid values do not exist.

1 Introduction

Internet advertisers spend billions of dollars every year, and large internet-search companies run keyword auctions millions of times a day. Despite this, our understanding of these auctions is built on incredibly simple assumptions. One of the most striking of these assumptions is that advertisers derive their value solely from users clicking on their ads. In reality, many advertisers care about *conversions* (that is, users actually completing some desired action such as buying items from their websites), while others care about *branding* (that is, simply familiarizing users with the name of the advertiser). Not all of these forms of deriving value from sponsored search advertisements can be mapped to a value per click, or the same value-per-click for all slots. Thus, while the value-per-click model is a reasonable first step, it leaves open many questions.

In this paper, we embark on the study of keyword auctions in which the value-per-click assumption is removed. Rather than modeling advertisers as having some fixed value per click, we allow them to have a full spectrum of values for slots, which are not necessarily proportional to the expected number of clicks received. This allows modeling, for example, advertisers who value slots based on some combination of factors including conversions, branding, and of course, clicks.

Our main focus in this paper is to understand the implications of using a *single-bid* system when bidders actually have a full spectrum of values for slots. While the VCG mechanism [8, 3, 5] applies to this setting and has the advantage of being a truthful mechanism and maximizing total efficiency (i.e. social welfare), it requires every advertiser to report a vector of bids rather than just one single bid. In addition to placing an

extra burden on advertisers and the system infrastructure, reporting a vector of bids is a significant departure from current bidding systems. And, how much would be gained? In particular, when the bidding system is a generalized second price auction (essentially the mechanism used in most sponsored search auctions), what can we say about efficiency and equilibria?

1.1 Related Work

If advertisers only derive value from clicks and value clicks in all slots equally, and if the publicly known clickthrough rate of each slot is independent of which advertiser is placed in it, the full spectrum of private values is simple to compute: it is the vector of clickthrough rates for the slots multiplied by the advertiser's value per click. The authors in [4] show that under these assumptions, the generalized second price auction (GSP) does indeed have at least one equilibrium, corresponding precisely to the VCG outcome. Further, they show that this equilibrium is envy-free, meaning that no advertiser would rather be in a different slot, paying that slot's associated price. Although the techniques used to prove these results can be generalized somewhat, e.g. to allow separable advertiser-dependent clickabilities, they cannot handle the more general case when the value-per-click assumption is removed. Similarly, Varian also makes the assumption that advertisers value each slot in proportion to its clickthrough rate [7].

Somewhat more general in terms of modeling bidder preferences, Aggarwal, Feldman, and Muthukrishnan [1] consider a *thresholded* value-per-click assumption: each advertiser has a threshold t such that she has value-per-click v for slots 1 through t , and value-per-click 0 for slots below t . In addition to capturing the normal value-per-click model, this model also allows some notion of branding. The authors propose a new auction mechanism which has the VCG outcome as an equilibrium. However, this work is quite different from ours: an advertiser is still assumed to derive positive value from clicks only, while our work allows any ranking of slots (specifically, an advertiser's value for an impression can be higher for a slot with a lower clickthrough rate). Second, the authors do not address the question of equilibria in the current GSP model; *i.e.*, the question of whether there are single bids corresponding to value vectors of this form that lead to efficient equilibria under GSP.

1.2 Our contributions

To the best of our knowledge, our work is the first to consider such a general form of advertiser valuations, and the effect of using a GSP auction with a single report from each advertiser.

The problems we address are the following:

- Oblivious strategies: How does the maximum social welfare achievable with oblivious bidding in this situation compare with the social welfare achieved by the VCG mechanism with the full value vectors as input? In §3, we show that oblivious bidding strategies are inefficient: for any (reasonable) oblivious bidding strategy, there is always a set of valuations that leads to $1/k$ of the optimal efficiency where k is the number of slots, even when all advertisers agree on the relative ranking of the slots.
- Equilibria: While oblivious strategies are inefficient, can the VCG outcome with vector inputs be achieved as an equilibrium of the GSP auction when bidders report a single value? In §??, we show the surprising result that when advertisers report a single bid per impression, *there exists a set of single bids for the GSP auction that leads to the efficient VCG outcome*. The same result holds, although under more stringent monotonicity conditions, when bids and prices are per-click (even though values need not be per-click).

Note that this result is more surprising than the corresponding results in [4, 7] where bidders have one-dimensional types and report a single scalar to the system; a similar situation exists in [1] where bidders have two-dimensional types, and are allowed to report two values to a modified GSP-like auction. In contrast, bidders in our model have vector valuations and yet the VCG outcome can be achieved with just a one-dimensional report to a GSP auction.

In our proof of this result, we additionally give a new, direct proof that the VCG outcome of any auction in which every bidder wants at most one item must be envy-free. Although this was shown in [6], we present a simpler, more accessible argument.

- Bidding strategies: Finally, while the VCG outcome is achievable in equilibrium, is there a natural bidding strategy that leads to this outcome? We present some negative results for this question in §5.

2 Model and problem statement

Our model is the following. We study a single instance of an auction for slots of a single keyword. There are n bidders (advertisers) competing in this auction, and k slots being auctioned. Each bidder has a vector of private values for the k slots, $\vec{v}_i = (v_i^1, \dots, v_i^k)$, where v_i^j is bidder i 's value for being shown in the j th slot (note that v_i^j is the value-per-impression in slot j , not the value-per-click in slot j ; values-per-impression can be computed from value-per-click or value-per-conversion using clickthrough rate or conversion rate information). We define $\vec{\theta} = (\theta_1, \dots, \theta_k)$ as the vector of (ad-independent) *clickthrough rates* in the k slots, and use these to define the value-per-click for bidder i at slot j can be seen to be v_i^j/θ_j . Note that these values-per-click may actually be meaningless quantities, such as for bidders who value impressions. The value vectors \vec{v} are not necessarily click-based, *i.e.*, $v_i^{j_1}/\theta_{j_1}$ is not necessarily equal to $v_i^{j_2}/\theta_{j_2}$, which would be the case if bidder i had a constant value-per-click in all slots, and derived value only from clicks.

This model is the most general model of private bidder valuations for the sponsored search setting. Specifically, it subsumes two important models used in prior work:

- $v_i^j = \theta_j v_{ii}$: This is the model used in [4, 7], and says that bidder i has a value-per-click v_{ii} , and her value (per impression) for slot j is the expected clickthrough rate of slot j times her value-per-click. It is easy to see that advertiser-dependent clickthrough rates, which we denote by μ_i , can be accounted for by multiplying it by the value-per-click.
- $v_i^j = \theta_j v_{ii}$ for $j \leq t_i$, and $v_i^j = 0$ otherwise: This is the thresholded value model studied in the work on position auctions in [1], and says that advertiser i has a uniform value-per-click v_{ii} until position t_i ; his value for clicks in slots beyond t_i is 0.

There are two important kinds of bidder valuations not covered in previous models: (1) Bidders may have different per-click values in different slots, *i.e.*, a bidder may value a click in one slot more than a click in other slots (for example, due to different conversion rates). (2) Bidders may have values not based on the number of clicks received; for example, advertisers concerned only with branding may not have a value proportional to the number of clicks received at all, but rather simply to the position they are displayed in. Our model allows for both these kinds of bidder values.

2.1 Monotonicity of bidder values

We say the advertisers' private values are *monotone* if they all agree on the ordering of the slots' values. That is, there is an ordering of the slots $1, 2, \dots, k$ such that every bidder values slot i at least as much as slot j for $i \leq j$, *i.e.*, $v_\ell^i \geq v_\ell^j$ for all $\ell = 1, \dots, n$, $1 \leq i \leq j \leq k$. We say private values are *strictly monotone* if in addition to being monotone, for all ℓ and $i \neq j$, $v_\ell^i \neq v_\ell^j$ unless $v_\ell^j = 0$. In other words, each advertiser's values for slots $1, 2, \dots, k$ are *strictly* decreasing, until the value becomes 0, at which point the values may remain 0.

Additionally, we say the advertisers' private value are *click-monotone* if they all agree on the ordering of the slots' values, in terms of value-per-click, and this ordering agrees with the clickthrough rate. That is, there is an ordering of the slots such that $\theta_1 \geq \theta_2 \geq \dots \geq \theta_k$ and for all $i \leq j$ and for every bidder ℓ , bidder ℓ 's value-per-click for slot i is at least as high as her value-per-click for slot j , *i.e.*, $v_\ell^i/\theta_i \geq v_\ell^j/\theta_j$. These values are *strictly click-monotone* if in addition to being click-monotone, for all ℓ and $i \neq j$, $v_\ell^i/\theta_i \neq v_\ell^j/\theta_j$ unless $v_\ell^j/\theta_j = 0$. We emphasize again the the per-click values might not be meaningful to advertisers.

Note that in previous work [4, 7, 1], the private values are all click-monotone, although it may not be immediately clear why click-monotonicity is a reasonable assumption. We discuss these issues further in §6.

2.2 Auction model

We consider two auction models. The first is based on bidding and paying per impression. This is preferable for a few reasons: first, advertisers’ values need be click-based at all; a per-impression bid simply says how much the advertiser values winning a slot in that auction, without saying that the value comes from clicks. For the same reason, the VCG mechanism computes allocations based on per-impression values and hence this is a more natural representation. More importantly, the per-impression bid model makes no assumptions about the clickability of ads or slots.

However, real internet auctions require bidders to report a value-per-click, not a value for impressions; the value-per-impression is assumed to be the product of the value-per-click times the clickthrough rate. Hence, we consider this auction as well, giving analogous results for both auctions throughout. We now describe both models, both of which we refer to as *single-bid auctions* to emphasize that they require bidders to use a single bid to represent a full spectrum of private values.

We refer to the pay-per-impression auction as $\text{GSP}_{\vec{v}}$, since it is the standard generalized second price (GSP) auction in which bidders pay for impressions. We assume that the auctioneer knows the ranking of the slots according to the monotone bidder values, and we assume that slots are numbered $1, \dots, k$, in this order.

In $\text{GSP}_{\vec{v}}$, every bidder submits a *single* bid b_i to the auctioneer, despite having a vector of values that may not be representable by a single value (note that b_i is a bid-per-impression, just like \vec{v}_i). Bidders are assigned to slots in decreasing order of bids b_i , *i.e.*, the bidder with the highest b_i is assigned slot 1, and so on. Slots are priced according to the generalized second price auction (GSP), *i.e.*, a bidder in slot j pays $b_{[j+1]}$ (where $b_{[j]}$ denotes the j th largest b_i), which is the bid of the bidder in the slot below. In the case of ties, we assume that the auctioneer is allowed to break ties in whatever way he chooses; generally, this will be at random. If the number of bidders is not greater than the number of slots, we will simply insert imaginary bidders, each bidding 0. So we may assume that $n > k$. Any bidder not assigned to one of the first k slots is simply not shown. For convenience, we allow bidders to be “assigned” to slots beyond the k -th— this is equivalent to not being shown. We do not consider reserve prices here.

We refer to the pay-per-click auction as $\text{GSP}_{\vec{c}}$. Again, we assume that the auctioneer knows the ranking of the slots according to the *click-monotone* bidder values, and we assume that slots are numbered $1, \dots, k$, in this order. The mechanism is as in [4], but we repeat it here for completeness.

In $\text{GSP}_{\vec{c}}$, each bidder submits a bid per-click, which the auctioneer puts in order. The i th-place bidder then pays a price equal to $b_{[i+1]}$ (*i.e.*, the $(i + 1)$ -st place bidder’s bid) every time a user *clicks* on her ad. In line with previous work, we assume that each slot has a fixed clickthrough rate, θ_j (for slot j). We say the clickthrough rates are *independent* if any advertiser placed in slot j will receive, in expectation, precisely θ_j clicks. (When calculating total value, this expected value is treated as a fixed constant.) Thus, if advertiser i is the j -th place bidder, she will be placed in slot j and pay a total of $\theta_j b_{[j+1]}$, netting a utility of $v_i^j - \theta_j b_{[j+1]}$.

Slightly more generally, clickthrough rates are said to be *separable* if for each $i = 1, 2, \dots, n$, advertiser i has an associated clickability α_i , and the expected number of clicks advertiser i receives for being placed in slot j is $\alpha_i \theta^j$. Several major auction systems are built on this separability assumption, and often bids and prices are adjusted based on clickability. Although our results extend to handle the separability assumption, we do not explicitly express it here, for simplicity.

2.3 VCG allocation and pricing

When bidders are allowed to report their full vector of valuations, the VCG mechanism can be applied to truthfully produce an efficient allocation. Throughout, we compare the performance of single-bid auctions (namely, $\text{GSP}_{\vec{v}}$ and $\text{GSP}_{\vec{c}}$) to this outcome, which we describe below.

Let G be a bipartite graph with advertisers on one side and slots on the other. The weight of edge (i, j) between bidder i and slot j has weight v_i^j . The VCG allocation computes the maximum weight matching

on this graph, and assigns advertisers to slots according to this matching. We will use M to denote the weight of the maximum matching on G , and number advertisers so that bidder i is assigned to slot i in this matching. Let M_{-i} denote the weight of the maximum weight matching on G when all edges incident to bidder i are removed. Then p_i , the VCG price for bidder i , is $p_i = M_{-i} + v_i^i - M$.

3 Inefficiency of oblivious strategies

In this section, we explore the loss in efficiency when restricting a bidder to oblivious bidding strategies (*i.e.*, her single bid is a function only of her full vector of values). We find that in terms of efficiency, no matter how bidders choose to encode their bids, performance is very poor amongst this class of strategies. This is in stark contrast with situations that allow bidders to express their full spectrum of private values in a single bid (e.g. see §2.3), where maximum efficiency is attainable.

Let $g : \mathbf{R}^k \rightarrow \mathbf{R}$ be the function that a bidder uses to map his value vector into a single bid. A GSP mechanism will be applied to this bid, so it is not effective for a bidder to simply encode their full vector of values into a single number (since the auctioneer cannot decode the single value into a full spectrum of bids). We show that when values are monotone, under mild assumptions on the behavior of g , the worst case loss in efficiency is lower-bounded by a factor k , even when the seller knows the ordering of the slots; *i.e.*, oblivious bidding strategies are not efficient. In fact, we show that this is tight; there are auctions for which this factor is obtained.

For $i = 1, \dots, k$, define $f_i(x) = g(x, \dots, x, 0, \dots, 0)$, where x is repeated i times. We will assume in this section that the auctioneer breaks ties deterministically. However, we could alternatively assume that, for all x , there exists $x' > x$ such that $f_i(x') > f_i(x)$, *i.e.*, f is eventually increasing in x for fixed i (this is a very weak restriction on a bidding strategy, and simply says that the reported bid is eventually increasing for a particular form of the value vector). Under this assumption, the proof of Theorem 1 can easily be modified so that there are no ties in reported single-bids, and the lower bound results will apply. Call g *block-continuous* if the induced functions f_i are continuous for all $i = 1, \dots, k$. (Note that block-continuity is a much weaker condition than requiring g to be continuous.) Then we have the following theorem.

Theorem 1. *For every oblivious bidding strategy g that is block-continuous, there are click-monotone [resp., monotone] value vectors such that the efficiency under $GSP_{\bar{c}}$ [resp., $GSP_{\bar{v}}$] is as small as $1/k$ of the maximum efficiency. Conversely, there is an oblivious block-continuous bidding strategy g such that the efficiency under $GSP_{\bar{c}}$ [resp., $GSP_{\bar{v}}$] is at least $1/k$ of the maximum efficiency, for all click-monotone [resp., monotone] private value vectors.*

Proof. For the first part of the proof, we consider $GSP_{\bar{v}}$. Note that the same argument shows that $GSP_{\bar{c}}$ can be inefficient as well, by setting clickthrough rates $\theta_1 = \dots = \theta_k = 1$. So, suppose we can find x_2, \dots, x_{k-1} such that

$$f_2(x_2) = f_1(1), f_3(x_3) = f_1(1), \dots, f_{k-1}(x_{k-1}) = f_1(1).$$

We will show that if we can find such x_i , then the efficiency can be as bad as $1/k$ of the maximum efficiency; if we cannot find such x_i , then we can make the efficiency be a vanishingly small fraction of the optimal efficiency, so that k is still a lower bound.

First suppose we can find such x_i , $i = 2, \dots, k-1$. Let $j = \arg \min_{i=1, \dots, k-1} x_i$. Consider the set of value vectors

$$\begin{aligned} v_1 &= (1, 0, \dots, 0), \\ v_2 &= (x_2, x_2, 0, \dots, 0), \\ &\vdots \\ v_k &= (x_k, \dots, x_k, 0, \dots, 0), \\ v_{k+1} &= v_j, \end{aligned}$$

where $v_j = (x_j, \dots, x_j, 0, \dots, 0)$ with x_j repeated j times.

Since the single bids reported by all bidders are the same, a possible allocation is one which assigns bidder $k + 1$ the first slot (we assumed that vectors are monotone) and bidders 1 through $k - 1$ to slots 2 through k , for a total efficiency x_j (since all other bidders except bidder k derive no value from their assignments). So the efficiency from a single bid auction can be as small as

$$x_j = \min_{i=1, \dots, k} x_i \leq \frac{1}{k} \sum_{i=1}^k x_i,$$

where $\sum_{i=1}^k x_i$ is the maximum efficiency, obtained by assigning bidder i to slot i .

Next we show that if such values cannot be found, then the efficiency can be arbitrarily bad, so the lower bound of k is trivially true. By our assumption that $f_i(x)$ is continuous in x , it is sufficient to consider the following two cases:

- There is an i , $2 \leq i \leq k$ such that $f_{i+1}(x) < f_i(x) \forall x$: Consider a set of bidders with the following values:

$$\begin{aligned} v_1 = v_2 = \dots = v_{k-1} &= (x_i, \dots, x_i, 0, \dots, 0), \\ v_k &= (x, x, \dots, x, 0, \dots, 0), \end{aligned}$$

where $x \rightarrow \infty$ is repeated $i + 1$ times (and x_i is repeated i times for bidders 1 through $k - 1$). Since the maximally efficient allocation assigns some slot between 1 and $i + 1$ to bidder k for total efficiency at least $x \rightarrow \infty$, whereas the single-bid allocation does not assign any of the top $i + 1$ slots to bidder k , since there are at $k - 1 \geq i + 1$ bidders with a higher single bid. Thus the single-bid efficiency is arbitrarily small.

- There is an i , $2 \leq i \leq k$ such that $f_{i+1}(x) > f_i(x) \forall x$: Consider a set of bidders with the following values:

$$v_1 = \dots = v_{k-1} = (x, \dots, x, 0, \dots, 0),$$

where $x \rightarrow 0$ is repeated $i + 1$ times, and $v_k = (1, 0, \dots, 0)$.

The most efficient allocation has an efficiency greater than 1, obtained by assigning bidder k to slot 1, but the single-bid allocation does not assign the first slot to bidder k but rather to some bidder $1, \dots, k - 1$ for a total efficiency of at most $kx \rightarrow 0$ as $x \rightarrow 0$. So the single-bid efficiency can be arbitrarily bad here as well.

Now, to show that there is a block-continuous bidding strategy whose efficiency is guaranteed to be within factor k of maximum, observe that the following function meets this lower bound for GSP $_{\bar{c}}$:

$$g(v_i^1, v_i^2, \dots, v_i^k) = \max_{j=1, \dots, k} v_i^j / \theta_j = v_i^1 / \theta_1 \quad (\text{by monotonicity})$$

This follows since the maximum efficiency for GSP $_{\bar{c}}$ is bounded by

$$\max_{S \subset \{1, \dots, n\}, |S|=k} \sum_{i \in S} v_i^1 / \theta_1 \leq k \max_{i=1, \dots, n} g(v_i),$$

and the efficiency obtained from using g is at least $\max_i g(v_i)$.

Likewise, we use $g(v_i^1, \dots, v_i^k) = v_i^1$ for the GSP $_{\bar{v}}$ auction. Again, we see that its efficiency is never smaller than $1/k$ times the maximum possible. \square

3.1 Inefficiency with non-monotone values

The situation is even worse when the private values are not monotone; we show that it is not possible to make any guarantee on efficiency in this case.

Theorem 2. *For any oblivious strategy g , there are value vectors such that the single-bid efficiency under single-bid auctions is arbitrarily smaller than the maximum possible efficiency.*

Proof. We consider the $\text{GSP}_{\bar{v}}$ auction. The $\text{GSP}_{\bar{c}}$ auction is similar. Let $v_1 = (0, x)$, $v_2 = (1, 0)$, $v_3 = (1, 0)$ and $1 \ll x$, and consider that there are two slots being auctioned off. The maximum efficiency is $1 + x$, obtained by placing either bidder 2 or 3 in the top slot and bidder 1 in the second slot.

- **Case I:** $g(0, x) > g(1, 0)$, then $\text{GSP}_{\bar{v}}$ will place bidder 1 in the first slot and one of bidders 2 and 3 in the second slot, for a total efficiency of 0.
- **Case II:** $g(0, x) < g(1, 0)$, then $\text{GSP}_{\bar{v}}$ will place bidders 2 and 3 in the first two slots, for a total efficiency of 1.
- **Case III:** $g(0, x) = g(1, 0)$, then if $\text{GSP}_{\bar{v}}$ deterministically chooses based on bidder identity, there is an assignment of identities to vectors with efficiency 1. Even if the tie is broken by a random coin flip, we may add an arbitrarily large number of bidders with valuation $v_i = (1, 0)$, so that the probability v_1 is assigned to the second slot is arbitrarily small, and the expected efficiency approaches 1.

□

4 Equilibrium with single bids

Having seen that oblivious strategies can be highly inefficient, we now investigate equilibria of single-bid auctions. In this section, we prove the following surprising result: the (full-spectrum) VCG outcome is also an envy-free equilibrium of $\text{GSP}_{\bar{c}}$ [and of $\text{GSP}_{\bar{v}}$] so long as the values are strictly click-monotone [respectively, strictly monotone], where an envy-free outcome means an outcome where for every bidder i (numbering bidders according to the slots they are assigned), $v_i^i - p_i \geq 0$, and for every slot j , $v_i^i - p_i \geq v_j^j - p_j$, where p_j is the (current) price for slot j . That is, bidder i (weakly) prefers slot i at price p_i to slot j at price p_j , for all $j \neq i$.

Our proof is comprised of two main parts. In the first part, we show that when the values are strictly click-monotone [respectively, strictly monotone], any envy-free outcome is a realizable equilibrium in $\text{GSP}_{\bar{c}}$ [resp., $\text{GSP}_{\bar{v}}$]. The second part gives a direct proof that the VCG outcome is always envy-free in auctions that match bidders to at most one item each (this proof holds even when the spectrum of values are not monotone, although we do not need it here).

Theorem 3. *Any envy-free outcome on k slots and $n > k$ bidders in which prices are nonnegative and values are strictly click-monotone [resp., strictly monotone] is a realizable equilibrium in $\text{GSP}_{\bar{c}}$ [resp., $\text{GSP}_{\bar{v}}$].¹*

Proof. We give the proof for the $\text{GSP}_{\bar{c}}$ auction. The proof for the $\text{GSP}_{\bar{v}}$ follows by setting $\theta_j = 1$ for all slots j .

As above, label bidders so that the envy-free outcome we consider assigns bidder i to slot i , at price p_i (note that p_i is a price per-impression, or simply the cost of slot i). Further, recall that $p_i = b_{i+1}\theta_i$. We first show that $p_i/\theta_i \geq p_j/\theta_j$ for all $i < j$ (which shows that $b_{i+1} \geq b_{j+1}$).

¹For technical reasons, we say an envy-free outcome is a realizable equilibrium so long as there is a set of bids leading to an equilibrium that agrees with the outcome for every bidder having nonzero value for her slot. We allow bidders that have zero value for their slots to be moved to other slots for which they have zero value (so long as they continue to pay 0 for the new slot).

Suppose $i < j$. By definition of an envy-free outcome,

$$\begin{aligned}
v_j^j - p_j &\geq v_j^i - p_i \\
\Rightarrow \theta_j \left(\frac{v_j^j}{\theta_j} - \frac{p_i}{\theta_j} \right) &\geq \theta_i \left(\frac{v_j^i}{\theta_i} - \frac{p_i}{\theta_i} \right) \\
&\geq \theta_j \left(\frac{v_j^i}{\theta_i} - \frac{p_i}{\theta_i} \right) \text{ by click-monotonicity} \\
&\geq \theta_j \left(\frac{v_j^j}{\theta_j} - \frac{p_i}{\theta_i} \right) \text{ by click-monotonicity} \\
\Rightarrow \frac{p_i}{\theta_i} &\geq \frac{p_i}{\theta_j}
\end{aligned}$$

Further, we see equality can occur only if $v_j^j = 0$. Now, consider the following set of bids: Bidder j bids p_{j-1}/θ_{j-1} for all j , where for convenience, we set $p_0 > p_1$, $\theta_0 < \theta_1$, and $p_j = 0, \theta_j = 1$ for all $j > k$. From the above argument, we see that the j th largest bid is indeed p_{j-1}/θ_{j-1} ; however, there may be ties. Hence, according to $\text{GSP}_{\bar{c}}$, bidder j is assigned slot j at a price $b_{j+1}\theta_j = p_j$, for $j = 1, \dots, k$, assuming that the auctioneer breaks ties in the “right” way. We now appeal to the fact that values are strictly click-monotone to remove this assumption.

We now show that bidding ties occur only when both bids are 0, and the bidders have no value for the slots they are tied for. To see this, first note that if $v_j^j = 0$, then $p_j = 0$, by the envy-free condition, hence $p_j/\theta_j = 0$. Hence, if $i < j$ and $p_i/\theta_i = p_j/\theta_j$, it must be the case that $v_j^j = 0$, implying $p_i/\theta_i = p_j/\theta_j = 0$. That is, bids are tied only when both bids are 0.

Now, suppose bidder j bids 0. We will show that $v_j^j = 0$. To this end, notice that bidder j bids $p_{j-1}/\theta_{j-1} = 0$, which implies $p_j/\theta_j = 0$, from the above argument. So $p_{j-1} = p_j = 0$. By the envy-free condition, $v_j^{j-1} - p_{j-1} \leq v_j^j - p_j$, hence $v_j^{j-1} \leq v_j^j$. But this violates strict click-monotonicity, unless $v_j^{j-1} = v_j^j = 0$. Hence, bidder j has value 0 for slot j .

Putting this together, we see that the auctioneer only breaks ties between bidders that bid 0, and have no value for any slot they are tied for. (Hence, these ties may be broken arbitrarily and still satisfy our goal.)

We finish by showing that this envy-free outcome is indeed an equilibrium. Suppose bidder i bids lower, so she gets slot $j > i$ rather than i . She then pays p_j , and by definition of envy-free, does not increase her utility. On the other hand, if bidder i bids higher, getting slot $j < i$, then she pays p_{j-1} , with utility $v_i^j - p_{j-1} \leq v_i^j - p_j \leq v_i^i - p_i$, since prices decrease with slot rankings. Therefore, this set of bids leads to an envy-free equilibrium in $\text{GSP}_{\bar{c}}$. \square

Next we show that the VCG outcome is envy-free in auctions that match each bidder to at most one item, even when the values are not monotone. We do not need clickthrough rates here at all, since the VCG price is simply a price per slot, or per-impression. We first prove a general property of maximum weight matchings on bipartite graphs. While this lemma is superficially similar to Lemma 2 in [1], we note that the proof is quite different: the proof in [1] crucially uses the fact that an advertiser’s value in a slot is the product of a value-per-click and a clickthrough rate which decreases with slot number, and simply does not work in our setting.

Lemma 1. *Let G , M , M_{-i} , and M_{-j} be as above. Then if advertiser j is assigned to slot j in a maximum matching, $M_{-j} \geq M_{-i} + v_i^j - v_j^j$.*

Proof. Fix a maximum matching on G , say \mathcal{M} , and without loss of generality, relabel the advertisers so that advertiser i is assigned to slot i for each $i = 1, \dots, k$ in \mathcal{M} . The proof of the lemma proceeds by taking a maximum matching of G with i removed, and using it to produce a matching of G with j removed. This new matching will have weight at least $M_{-i} + v_i^j - v_j^j$, showing that M_{-j} must also be at least this large.

Fix a maximum matching of G with i removed, call it \mathcal{M}_{-i} . If there is more than one such maximum matching, we will take \mathcal{M}_{-i} to be one in which advertiser j is matched to slot j , if such a maximum matching exists. Either advertiser j is matched to slot j in \mathcal{M}_{-i} or not. We consider each case in turn.

- **Case I.** Bidder j is assigned to slot j in \mathcal{M}_{-i} .

In this case, simply remove the edge from advertiser j to slot j in \mathcal{M}_{-i} , and add the edge from advertiser i to slot j . This is now a matching on G without j , and its total weight is $M_{-i} + v_i^j - v_j^j$. Hence, $M_{-j} \geq M_{-i} + v_i^j - v_j^j$.

- **Case II.** Bidder j is not assigned to slot j in \mathcal{M}_{-i} .

Again, we will construct a matching for G without j , in a somewhat more complicated way. We first observe that removing advertiser i from a maximum matching on G creates a “chain of replacements.” More precisely, let i_1 be the advertiser that is matched to slot i in \mathcal{M}_{-i} . Notice that $i_1 \neq i$. If $i_1 \leq k$, then let i_2 be the advertiser matched to slot i_1 in \mathcal{M}_{-i} . And in general, if $i_\ell \leq k$, let $i_{\ell+1}$ be the advertiser matched to slot i_ℓ in \mathcal{M}_{-i} . Let t be the smallest index such that $i_t > k$. We first claim that for some $s < t$ that $j = i_s$. To see this, let A be the set of advertisers $\{i, i_1, \dots, i_{t-1}\}$, and let S be the set of slots $\{i, i-1, \dots, i_{t-1}\}$. Let \mathcal{M}' be the maximum matching \mathcal{M} restricted to advertisers not in A and to slots not in S , and let \mathcal{M}'_{-i} be matching \mathcal{M}_{-i} restricted to advertisers not in A and slots not in S . Clearly, the weight of \mathcal{M}' and \mathcal{M}'_{-i} must be the same. But if j is not in A , then \mathcal{M}' matches advertiser j to slot j . By our choice of \mathcal{M}_{-i} , this means that \mathcal{M}_{-i} matches advertiser j to slot j , and we are back to case 1.

So, $j = i_s$ for some $s < t$. For convenience, let $i_0 = i$. Change \mathcal{M}_{-i} as follows: for each $\ell = 1, 2, \dots, s$, remove the edge from advertiser i_ℓ to slot $i_{\ell-1}$, and replace it with the edge from $i_{\ell-1}$ to $i_{\ell-1}$. Notice that in this new matching, advertiser $i_s = j$ is matched to no one. Furthermore, it is easy to see that its weight is $M_{-i} + \sum_{\ell=1}^s v_{i_{\ell-1}}^{i_{\ell-1}} - \sum_{\ell=1}^s v_{i_\ell}^{i_{\ell-1}}$. Since \mathcal{M} is a maximum matching, we see

$$\sum_{\ell=1}^s v_{i_{\ell-1}}^{i_{\ell-1}} + v_{i_s}^{i_s} \geq \sum_{\ell=1}^s v_{i_\ell}^{i_{\ell-1}} + v_{i_0}^{i_0} \Rightarrow \sum_{\ell=1}^s v_{i_{\ell-1}}^{i_{\ell-1}} - \sum_{\ell=1}^s v_{i_\ell}^{i_{\ell-1}} \geq v_i^j - v_j^j$$

Substituting, we have that $M_{-j} \geq M_{-i} + v_i^j - v_j^j$.

□

Theorem 4. *The VCG outcome is envy-free, even when bidder values are not monotone.*

Proof. As always, we assume without loss of generality that bidder i is assigned to slot i in the VCG outcome, for $i = 1, \dots, k$. Recall that the price set in the VCG outcome for bidder j is $p_j = M_{-j} - M + v_j^j$. Hence, from our lemma above, we have that for all i, j

$$\begin{aligned} v_i^j - p_j &= v_i^j - M_{-j} + M - v_j^j \\ &\leq v_i^j - (M_{-i} + v_i^j - v_j^j) + M - v_j^j \\ &= M - M_{-i} = v_i^i - p_i \end{aligned}$$

Furthermore,

$$v_i^i - p_i = M - M_{-i} \geq 0.$$

That is, the VCG outcome is envy-free.

□

Combining our two main theorems yields the following.

Theorem 5. *When bidder values are strictly click-monotone [resp., strictly monotone], there exists an equilibrium of $GSP_{\bar{c}}$ [resp., $GSP_{\bar{v}}$] that corresponds to the welfare-maximizing outcome, i.e., the VCG outcome.*

Notice that this generalizes the result of [4], but the techniques we use are quite different; although the proof of [4] can be extended to cover a slightly more general case, the technique relies on an assumption that is not true in our more general setting. Also, notice that our result immediately implies that under the thresholded value-per-click assumption of [1], the VCG outcome is an attainable equilibrium in $GSP_{\bar{c}}$.

5 Bidding strategies

Now, we consider the following question: while the VCG outcome is indeed realizable as an efficient equilibrium of single-bid auctions, is there a natural bidding strategy for bidders, such that when all bidders bid according to this strategy in repeated auctions for the same keyword (so that the private values do not change), they converge to this efficient outcome?

We define an *history-independent bidding strategy*, g , to be a set of functions g_1, \dots, g_n , one for each bidder i , where g_i takes the bids \vec{b} from the previous round of the auction, together with the vector of private values \vec{v}_i , and outputs a nonnegative real number, which is the bid of player i .² We call the history-independent strategy *myopic* if the value of g_i does not depend on the bid of player i from last round. Finally, we say a vector of bids \vec{b} is a *fixed point* for history-independent strategy g , if after bidding \vec{b} , the bidding strategy continues to output that same vector of bids \vec{b} , *i.e.*, $g_i(\vec{b}; \vec{v}_i) = \vec{b}$. Throughout this section, we will focus primarily on $\text{GSP}_{\vec{v}}$ for clarity, although the results also hold for $\text{GSP}_{\vec{c}}$.

First, we show that for any bidding strategy that always has an envy-free fixed point, there are value vectors where bidders must bid more than their maximum value over all slots. This puts them at risk to lose money (*i.e.*, have negative utility) if other bidders change their bids, or new bidders join the auction. We call a strategy *safe* if it never requires bidders to bid beyond their maximum value. Although it would not be surprising to see bidders occasionally bid values that are not 'safe', it seems questionable whether they would continue to bid this way repeatedly.

Theorem 6. *There is no safe, history-independent bidding strategy that always has an envy-free fixed point whenever the private values are strictly monotonic.*

Proof. Consider the following auction ($\text{GSP}_{\vec{v}}$), in which three bidders vie for two slots: Bidder 1 values slot 1 at 20 and slot 2 at 0; bidder 2 values slot 1 at 15 and slot 2 at 14; bidder 3 values slot 1 at 19 and slot 2 at 0. Let us examine an envy-free fixed point, showing that bidder 2 must bid above her maximum value (*i.e.*, 15).

If bidder 1 does not get slot 1, then clearly she will be envious (unless the cost is greater than 20, in which case the assignment cannot be envy-free). If bidder 1 does not pay at least 19 for slot 1, then bidder 3 will be envious. Further, slot 2 must be assigned to bidder 2, for otherwise she will be envious (or someone will have to pay more than their value for slot 2, violating the envy-free condition). Hence, bidder 2 must bid at least 19, showing that she must bid above her maximum value. \square

We further show that no bidding strategy has the VCG outcome as a fixed point, under a few minor assumptions. These assumptions stem from our need to stop bidders from encoding information in their bids, essentially artificially allowing the players to calculate the fixed point and bidding accordingly. As such, recall that in the VCG outcome, when the values are strictly monotone, all non-zero bids are distinct, and determined. Further, the highest bidder is free to bid any value larger than some quantity (determined by the VCG prices). Here, we refer to the VCG bids as the set of bids supporting the VCG outcome in which the highest bidder bids her utility for the top slot. Then we have the following:

Theorem 7. *There is no myopic bidding strategy that always attains the VCG outcome as a fixed point, as defined above, even if that VCG outcome does not require bidders to bid above their maximum value over all slots.*

Proof. First consider the following auction ($\text{GSP}_{\vec{v}}$), in which three bidders vie for two slots: Bidder 1 values slot 1 at 30 and slot 2 at 16; bidder 2 values slot 1 at 15 and slot 2 at 14; bidder 3 values slot 1 at 13 and slot 2 at 10. The VCG outcome of this auction has bidder 1 bid 30, bidder 2 bid 13, and bidder 3 bid 10.

Now, suppose bidder 3 instead valued slot 1 at 12 and slot 2 at 10. Then the VCG outcome of this auction has bidder 1 bid 30, bidder 2 bid 12, and bidder 3 bid 10. So they cannot both be fixed points of the myopic strategy, since the input to g_2 is the same in both cases—bidder 1 bids 30 and bidder 3 bids 10, and private values are slot 1 at 15 and slot 2 at 14 (the bid of bidder 2 is ignored). \square

²Notice that our definition actually allows a kind of collaboration between bidders, despite the fact that in practice, we do not expect this to happen.

In [2], the authors show that under the model of a single value-per-click for all slots as in [4, 7], there is a myopic strategy, based on envy-free bidding, that does indeed have the VCG outcome as a fixed point. The result above shows that this does not work in our more general setting.

We might wonder whether we can prove even stronger results on bidding strategy. Unfortunately, there is an *unnatural* bidding strategy that converges to the VCG outcome. Specifically, suppose that the private vector values in a $\text{GSP}_{\vec{c}}$ [resp., $\text{GSP}_{\vec{v}}$] auction are strictly click-monotone [resp., strictly monotone]. Then there is an history-independent bidding strategy that converges to the VCG outcome. The proof of this, which relies heavily on the ability to encode private values in bids, can be found in the full version of this paper. The strategy allows bidders to encode their private values in their bids, which are then used by the other players to calculate the VCG outcome and bid accordingly. This is unnatural for (at least) two reasons: bidders should not be able to encode much information in a bid, and bidders will not cooperate with each other consistently. The two negative results we give above use the fact that bidders are restricted in their ability to artificially encode information. We leave open the question of whether we can prove similar negative results by restricting the cooperation between players.

6 The monotonicity assumption

Throughout this paper, we show a number of results that rely on the private values to be monotone, or click-monotone. Although previous work has (implicitly) assumed click-monotonicity [4, 7, 1], this assumption may not always be true. Here, we give several examples demonstrating the problems that can occur with non-monotone values (in addition to §3.1, which shows the inefficiency that can occur with oblivious bidding).

Instability in $\text{GSP}_{\vec{v}}$. We first give a simple, but compelling, reason for assuming monotone private values. Put simply, the auction simply does not work in general without this assumption: the auctioneer must pick some ordering of the slots to assign bidders in decreasing order of bid, but this is not meaningful when all bidders do not agree on the ordering.

For example, suppose there are just three slots, and bidders 1,2,3 value slot 1 at 100,90,80, respectively, and slots 2,3 at 0. Further, suppose bidders 4,5,6 value slot 3 at 100,90,80, respectively, and slots 1,2 at 0. Let us assume that the auctioneer orders slots so that the highest bidder gets slot 1, the second highest gets slot 2, and the third gets slot 3. Bidder 1 can bid high in order to obtain slot 1. Everyone else has no incentive but to bid 0: bidders 2,3 do not want slots 2 or 3, so will not pay; bidders 4,5,6 want slot 3, but bidding above 0 puts them into slot 2, which they do not want.

Instability in $\text{GSP}_{\vec{c}}$. Unfortunately, when private values for slots are monotone, but not click-monotone, there are situations in which $\text{GSP}_{\vec{c}}$ has no equilibrium outcome at all. This is not surprising given that bidder's values are not really per-click values: the intuition behind our counterexample is that when the 'value' for a click is higher in slots with lower clickthrough rate, then the most desirable slot is not the slot for which the most money is charged. Therefore, everyone will vie for this slot, and we cannot reach an equilibrium because the auction mechanism does not allow us to charge more for the more desirable item.

Theorem 8. *$\text{GSP}_{\vec{c}}$ does not always have an equilibrium, even when the private values are strictly monotone.*

Proof. We prove the theorem by giving an example for which there is no equilibrium outcome under $\text{GSP}_{\vec{c}}$. The bidder valuations are given in the table below. We assume clickthrough rates of $\theta_1 = 1$ and $\theta_2 = 0.1$, which are bidder independent.

	Slot 1	Slot 2
Clickthrough rate	1	0.1
Bidder A	5	1
Bidder B	5	1
Bidder C	5	1

Consider two possible cases: **Case I.** Assume $b_{[1]} > b_{[2]}$. Then, in order for the excluded bidder to have no incentive to bid some value in between $b_{[1]}$ and $b_{[2]}$, we must have $1 - (0.1)b_{[2]} \geq 0 \Rightarrow 10 \leq b_{[2]}$. But then the utility for the bidder with the highest bid is $5 - 10 = -5$. Since a bidder cannot have negative utility in an equilibrium, we have reached a contradiction. **Case II.** Assume $b_{[1]} = b_{[2]}$. Whether the tie is broken by a random coin flip, or by some other process, we must ensure that the highest-ranked bidder has no incentive to undercut the second highest bidder. Hence, we must have that $5 - b_{[2]} \geq 1 - (0.1)b_{[3]} \geq 0$, where the last inequality follows because the utility of the second highest bidder must be nonnegative. To ensure the excluded bidder does not vie for the top slot, we must have $5 - b_{[1]} \leq 0 \Rightarrow b_{[1]} \geq 5$. Combining the previous inequalities, $b_{[2]} = b_{[1]} = 5$ and $1 - (0.1)b_{[3]} = 0 \Rightarrow b_{[3]} = 10$. Thus, we have reached a contradiction since $b_{[1]} \geq b_{[2]} \geq b_{[3]}$. \square

Observe that in this example, the per-click value for slot 1 is 5, while the per-click value for slot 2 is 10. Therefore, the ordering of values per-click is not the same as the ordering of the slot clickabilities, violating click-monotonicity.

How reasonable is the monotonicity assumption anyway? It is quite possible that bidders all value impressions in slots in the same order: branding advertisers value top slots highly; advertisers deriving value from clicks alone (*i.e.*, the same value from clicks in all slots) also value higher ranked slots more than lower ranked ones. Further, the decrease in clickthrough rates is perhaps sufficiently steep to offset any minor increases in rates of conversion (after an ad is clicked on) from slots other than the top slots. So it is reasonable to assume that bidders all value impressions in slots in the same order. The click-monotonicity assumption, on the other hand, may not be quite as well-founded; for instance if the ordering of slots according to post-click conversion rates is indeed quite different from the ordering according to clickthrough rates.

7 Discussion

The results in our paper provide insights into single-bid auctions in relation to bidder valuations. The results in [4, 7, 1] all point out that despite not being incentive compatible, the GSP auction has nice properties: the VCG outcome is an envy-free equilibrium of the GSP auction, and there are simple greedy bidding strategies that converge to the VCG outcome when all players bid according to these strategies.

When advertisers have private value vectors, rather than a single private value, the first result is still true: the VCG outcome is still attainable as an envy free equilibrium of the (single-bid) GSP under appropriate monotonicity conditions. However, the negative result on bidding strategies is unsettling: it is not clear whether the efficient outcome will indeed be attained in practice. Note that the negative result on bidding strategies had nothing to do with bidding per-impression or per-click: the difference with the single value assumption is that although bidder 1 may value slot 1 much more than bidder 2, bidder 2 may value slot 2 much more than bidder 1, which cannot happen when values for all slots are the product of a single value-per-click and a clickthrough rate.

This raises two natural directions for further work. First, what are bidder values really like? How much do they deviate from the one-dimensional assumption? And, do they satisfy the conditions under which we show $\text{GSP}_{\bar{c}}$ to have an efficient equilibrium? Note also that our conditions for existence of efficient equilibria are sufficient, but not necessary: the question of necessary conditions for the existence of efficient equilibria in $\text{GSP}_{\bar{c}}$ is open. Second, under what conditions on bidder values do there exist simple bidding strategies that converge to the efficient outcome?

There is a tradeoff between expressiveness, and the overhead imposed on bidders and the mechanism: at one extreme, a full vector of bids can be accepted, but this imposes a severe burden on both bidders, and the bidding system, which must compute allocations and prices for millions of auctions everyday. At the other extreme is the current system, which arguably is not expressive enough- bidders with widely varying types might be forced to report a single number to the system, leading to possible instability and inefficiency. The question of choosing the right tradeoff (and the resulting mechanism design questions) is a fascinating problem with plenty of opportunity for both experimental and theoretical research.

References

- [1] G. Aggarwal, J. Feldman, and S. Muthukrishnan. Bidding to the top: VCG and equilibria of position-based auctions. *Workshop on Approximation and Online Algorithms*, 2006.
- [2] M. Carry, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. *ACM Conference on Electronic Commerce*, 2007.
- [3] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17-33, 1971.
- [4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalizaed second price auction: Selling billions of dollars worth of keywords. *Second Workshop on Sponsored Search Auctions*, Ann Arbor, MI. June, 2006.
- [5] T. Groves. Incentives in teams. *Econometrica*, 41:617-631, 1973.
- [6] H. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *the Journal of Political Economy*, 91(3):461-479, 1983.
- [7] H. R. Varian. Position auctions. *To appear in International Journal of Industrial Organization*, 2006.
- [8] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8-37, 1961.