Cost of Conciseness in Sponsored Search Auctions

Zoe Abrams
Yahoo! Inc
Santa Clara, CA
za@yahoo-inc.com

Arpita Ghosh
Yahoo! Research
Santa Clara, CA
arpita@yahoo-inc.com

Erik Vee
Yahoo! Research
Santa Clara, CA
erikvee@yahoo-inc.com

ABSTRACT

We study keyword auctions in a model where each advertiser has a value for every slot, which is not necessarily proportional to the number of clicks the advertiser expects to receive in that slot. In our model, advertisers need not only derive values from clicks on their ad, nor do they need to value clicks in all slots equally. This model encompasses a variety of advertising objectives, including conversions (user completion of some task such as a purchase) and branding (an increase in consumer awareness). Our goal is to understand the cost, in terms of efficiency, of reporting a single value to current bidding systems (modeled by the generalized second price auction), while actually having a full spectrum of values. Surprisingly, we find that there always exists an equilibrium with the VCG outcome if bidders report a single value to a version of the current bidding system that charges per impression. However, if bidders report a single value and are charged per click, this is not the case: an equilibrium need not exist, although efficient equilibria exist under restrictive assumptions on the valuations of bidders and clickthrough rates.

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 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. 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.

An obvious approach is simply to use the well-known VCG mechanism [8, 3, 5]. Although this has the advantage both of being a truthful mechanism and of 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 the advertisers and the system infrastructure, reporting a vector of bids is a significant departure from current bidding systems. And, how much would be gained? Our main focus here is to better understand the cost of using a *single-bid* system when bidders actually have a full spectrum of values for slots. 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

WWW 2007, May 8–12, 2007, Banff, Alberta, Canada.
ACM 978-1-59593-654-7/07/0005.

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 cost of using a GSP auction with a single report from each advertiser.

In §3, we show that oblivious bidding strategies, where an advertiser computes his single bid as a function of his spectrum of values alone (independent of other advertisers), is 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.

This leads us to study full information Nash equilibria in §4. When advertisers report a single bid per impression, we show the surprising result that *there exists a set of single bids for the GSP auction that leads to the efficient VCG outcome*. The proof of this result relies on showing that the VCG outcome in any auction where every bidder wins at most one item is envy free (earlier proofs of this result for special cases of our model [4, 7] do not generalize to this setting). (It was pointed out to us afterwards that this result also appears in [6]; our proof is different and self-contained.)

In §5, we study this same situation when advertisers pay per click received (as opposed to per impression), which turns out to be rather different: there are realistic examples where GSP has no equilibrium at all. We show certain conditions under which efficient equilibria exist. The difference between charging advertisers for impressions and for clicks has potential implications for the design of GSP auctions, as discussed in §6.

2. MODEL

2.1 Bidder Valuation Model

Our model is the following. There are n bidders, and k slots. Each bidder has a vector of private values for the k slots, $\vec{v}_i = (v_i^1, \ldots, v_i^k)$. Here v_i^j is bidder i's value for being shown in the jth slot, *i.e.*, it is a value per impression, not a value per click.

We say the advertisers' private values are *monotone* if they all agree on the order of the slots' values. That is, there is an ordering of the slots $1, 2, \ldots, k$ 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 $\ell = 1, \ldots, n$, $1 \leq i \leq j \leq k$. 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,...,k are *strictly* decreasing, until the value becomes 0, at which point the values may remain 0. Throughout this paper we assume private values are strictly monotone unless we explicitly say otherwise. We discuss non-monotone values in §6.

The current model subsumes past models in the following sense. The pay-per-click model, where a bidder i has a value u_i per click, is a special case with $\vec{v}_i = u_i\vec{\theta}$, where $\vec{\theta} = (\theta^1, \ldots, \theta^k)$ is the vector of clickthrough rates in the k slots. (Again, this assumes that the clickthrough rate for each slot is independent of which ad is placed in it.) When clickthrough rates are not bidder independent, this can be represented using $\vec{v}_i = u_i\vec{\theta}_i$, where $\vec{\theta}_i = (\theta_i^1, \ldots, \theta_i^k)$ is the vector of clickthrough rates for bidder i in each slot. Note that this allows for separable as well as non-separable clickthrough rates. The model in [1] is a special case of monotone values, where all advertisers have a value-per-click vector of the form $(u_i, \ldots, u_i, 0, \ldots, 0)$, *i.e.*, they value equally all positions up to some position t_i, and have a value of 0 for all slots below this slot.

Our model also allows bidders to 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. In addition, it allows for 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. Finally, we note that advertisers' rankings of slots are allowed to be completely uncorrelated with the ordering of slots based on clickthrough rates: an advertiser interested most in branding, then conversions, then clicks, might value slots in an order that is different from the ordering by clickthrough rates.

2.2 Auction Model

Since bidders are assumed to have private values given as value-per-impression, we find it most natural to consider a pay-per-impression bidding system. Thus, much of our paper is devoted to the *pay-per-impression* GSP (PPI-GSP). In PPI-GSP, bidder i submits a single bid, b_i, for each $i = 1, ..., n$. The auctioneer has a ranking of the slots $1, \ldots, k$. (We assume this ordering is the same as those followed by the monotone bids and is known in advance to the auctioneer.) The bidders are ranked in decreasing order of b_i, and the jth highest bidder is assigned slot j at a price $b_{[j+1]}$ (where $b_{[j]}$ denotes the jth largest b_i). 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.

Most bidding systems in use today are more closely aligned with the pay-per-click GSP (PPC-GSP). We discuss this model more fully in Section 5.

We will sometimes refer to the above PPI-GSP auction as *single-bid PPI-GSP*, to emphasize that bidders actually have a full spectrum of values but are representing it using only a single bid; we will sometimes use *full spectrum VCG* to refer to the VCG outcome given the full vector of valuations v_i. For readers unfamiliar with the VCG mechanism, we briefly sketch the way it works in our context, in §4.1.

3. BOUNDING LOSS IN EFFICIENCY

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 §4, where maximum efficiency is attainable.

Let $g : \mathbf{R}^k \to \mathbf{R}$ be the function that a bidder uses to map his value vector into a single bid. We assume that the GSP auction will be applied to this single 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.

3.1 Non-monotone values

In this section, we show that it is not possible to make any guarantee on efficiency when values are not monotone.

THEOREM 1. *For any g, there are value vectors such that the single-bid efficiency under PPI-GSP is arbitrarily smaller than the maximum possible efficiency.*

PROOF. 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 PPI-GSP 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 PPI-GSP 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 PPI-GSP deterministically chooses based on bidder identity, there is an assignment of identities to vectors with efficiency 1. Otherwise, if PPI-GSP is randomized, 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.

□

3.2 Monotone values

Now we turn our attention to the case of monotone values and show that, 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.

For $i = 1, \ldots, k$, define $f_i(x) = g_i(x, \ldots, x, 0, \ldots, 0)$, where x is repeated i times. We will assume in this section that PPI-GSP breaks ties deterministically. However, we could alternatively assume that, for all x, there exists $x' > x$ such that $f_i(x') > f(x)$, *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 2 can easily be modified so that there are no ties in reported single-bids, and the lower bound results will apply.

THEOREM 2. *For any g such that $f_i(x)$, $i = 1, \ldots, k$ is continuous, there are value vectors such that the single-bid efficiency is as small as $1/k$ of the maximum efficiency; this bound is tight.*

Proof: To show this, suppose we can find $x_2, \ldots, x_{k-1}$ such that

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

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, $i = 2, \ldots, k-1$. Let $j = \arg\min_{i=1,\ldots,k-1} x_i$. Consider the set of value vectors

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

where $v_j = (x_j, \ldots, x_j, 0\ldots, 0)$ 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), with 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

$$x_j = \min_{i=1,\ldots,k} x_i \le \frac{1}{k} \sum_{i=1}^{k} x_i,$$

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. 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 $f_{i+1}(x) < f_i(x_i) \forall x$: Consider a set of bidders with the following values:

$$v_1 = v_2 = \ldots = v_{k-1} = (x_i, \ldots, x_i, 0, \ldots, 0),$$
$$v_k = (x, x, \ldots, x, 0, \ldots, 0),$$

  where $x \to \infty$ 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 $x \to \infty$, 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 $f_{i+1}(x) > f_i(x_i) \forall x$: Consider a set of bidders with the following values:

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

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

  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 $1, \ldots,, k-1$ for a total efficiency of at most $kx \to 0$ as $x \to 0$. So the single-bid efficiency can be arbitrarily bad here as well.

Observe that the max function meets this lower bound, *i.e.*,

$$g(v_i^1, v_i^2, \ldots, v_i^k) = \max_{j=1,\ldots,k} v_i^j$$

leads to an efficiency which is never smaller than $1/k$ times the maximum possible, which is bounded by

$$\max_{S \subset \{1,\ldots,n\}, |S|=k} \sum_{i \in S} \max v_i \leq k \max_{i=1,\ldots,n} g(v_i).$$

4. EQUILIBRIUM EXISTENCE UNDER PPI-GSP

In this section, we investigate equilibria of the single bid PPI-GSP. We prove the following surprising result: the full-spectrum VCG outcome is an envy-free equilibrium of PPI-GSP.

Our proof is comprised of two main parts. In the first part, we show that when the values are strictly monotone, any envy-free outcome (*i.e.*, assignment of advertisers to slots and prices) is an attainable equilibrium of PPI-GSP. The second part gives a direct proof that the VCG outcome is always envy-free in auctions that match bidders to at most item each, even when the spectrum of values are not monotone.

For clarity, we start with stating the outcome of the VCG mechanism in this setting.

4.1 VCG Mechanism

When bidders are allowed to report their full vector of valuations, the VCG mechanism can be applied to truthfully produce an efficient allocation. In this setting, let G be a bipartite graph with advertisers on one side and slots on the other. The weight of edge (l, i) between bidder l and slot i has weight v_l^i. 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 advertiser 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 i are removed. Then p_i, the VCG price for bidder i, is $p_i = M_{-i} + v_i^i - M$. We say an outcome is *envy-free* if for every bidder i, $v_i^i - p_i \geq 0$, and for every slot j, $v_i^i - p_i \geq v_i^j - p_j$, where p_j is the (current) price for slot j. That is, bidder i does not prefer slot j at price p_j, for all $j \neq i$.

4.2 Existence of efficient equilibrium

THEOREM 3. *Any envy-free outcome on k slots and $n > k$ bidders in which prices are nonnegative and values are strictly monotone is an attainable equilibrium in PPI-GSP.*[1]

PROOF. As above, label bidders so that the envy-free outcome we consider assigns bidder i to slot i, at price p_i. Suppose $i > j$; since all bidders have monotone values, $v_i^i \leq v_i^j$. By definition of an envy-free outcome,

$$v_i^i - p_i \geq v_i^j - p_j \Rightarrow p_j - p_i \geq v_i^j - v_i^i \geq 0,$$

i.e., $p_j \geq p_i$ for all $j < i$ (prices are higher for higher-ranked slots).

Consider the following set of bids: bidder 1 bids any amount larger than p_1, bidder j bids p_{j-1} for $j = 2$ through $k+1$, and the remaining bidders bid 0. According to PPI-GSP, bidder j is assigned slot j at a price p_j, for $j = 1, \ldots, k$, assuming that the auctioneer breaks ties in the "right" way. We now appeal to the fact that values are strictly monotone to remove this assumption.

Again assume that $i > j$, and note that either $p_j > p_i$, or $p_i = p_j$. In the second case, we see $0 = p_j - p_i \geq v_i^j - v_i^i \geq 0$, hence $v_i^j = v_i^i$. Since the values are strictly monotone, it must be that $v_i^j = v_i^i = 0$. Furthermore, $v_i^i - p_i \geq 0$, hence, $p_i = 0 = p_j$. In other words, there is some t such that the sequence of prices looks like $p_1 > p_2 > ... > p_t = p_{t+1} = ... = p_k = 0$.

Now, bidders $1, 2, ..., t$ all bid distinct positive values, so the auctioneer never needs to break a tie for these bidders. If $t = k$, we are done. If $t < k$, we argue that

[1]For technical reasons, we say an envy-free outcome is an attainable 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).

any bidder $i > t$ is equally happy to be placed in any spot $j > t$, paying price $p_j = 0$ (included being placed in no slot at all). Consider bidder $i > t$, bidding $p_{i-1} = 0$. In the envy-free outcome, bidder i pays $p_i = 0$, so $v_i^i - p_i \geq v_i^t - p_t \Rightarrow v_i^i \geq v_i^t$. Hence, $v_i^i = v_i^t$, since the values are decreasing. But the values are strictly monotone, so we see $v_i^i = 0$. Hence, $v_i^j - p_j = 0$ for all $j > t$. That is, bidder $i > t$ is indifferent as to which spot she is placed, including being placed in no slot at all, so long as it is beyond the t-th slot.

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 $v_i^j - p_{j-1} \leq v_i^j - p_j \leq v_i^i - p_i^i$, since prices decrease with slot rankings. Therefore, this set of bids leads to an envy-free equilibrium under PPI-GSP. □

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 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 $i = 1, ..., k$ 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 $M_{-i} + v_i^j - v_j^j$, thus 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,

$$M_{-j} \geq M_{-j} + v_i^j - v_j^j.$$

- **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 $\{i, i_1, ..., i_{t-1}\}$, and let S be the set of slots $\{i, i-1, ..., i_{t-1}\}$. Let $\mathcal{M}'$ be the maximum matching $\mathcal{M}$ restricted to advertisers not in A and to slots not in S, and 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 $\ell = 1, 2, ..., s$, 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

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

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

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 the bidders' 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 $i = 1, \ldots, k$. 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 monotone, there exists an equilibrium of PPI-GSP which corresponds to the welfare-maximizing outcome, i.e., the VCG outcome.*

We note that this result subsumes a number of cases, including nonseparable clickthrough rates: even when clickthrough rates are non-separable (although known to bidders), there is an efficient equilibrium under PPI-GSP.

It should not be surprising that this efficient equilibrium is not necessarily unique. Consider the following example.

	Slot 1	Slot 2	Slot 3	Equilibrium Bid
Bidder A	102	100	0	102
Bidder B	99	51	0	51
Bidder C	51	50	0	50

The efficient solution is to assign Bidder B to the first slot and Bidder A to the second slot. Yet, given their current bids, Bidder B is assigned to the second slot and Bidder A to the first slot, and no bidder has an incentive to change their bid. In general, there can be equilibria where the solution is not efficient, and equilibria where the solution is efficient and has strictly larger revenue than the equilibrium based on VCG prices.

5. EQUILIBRIA UNDER PAY-PER-CLICK GSP

The results of the previous sections show that there is an equilibrium corresponding to the full spectrum VCG outcome when the PPI-GSP auction is used. However, most current systems implement PPC-GSP. This section considers bidder behavior when using PPC-GSP. Not surprisingly, we show that in general, there are situations in which no equilibrium solution exists. Conversely, under certain additional assumptions, the full-spectrum VCG outcome is an attainable equilibrium. Before presenting these results, we first summarize the PPC-GSP model from [4] for completeness.

The PPC-GSP auction follows essentially the same rules as PPI-GSP; each bidder submits a bid, which the auctioneer puts in order. The ith-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 for this section that each slot has a fixed *clickthrough rate*, say with slot j having clickthrough rate θ^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 $i = 1, 2, .., n$, 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. We simply note that our results easily extend to PPC-GSP modified in this way.

5.1 Pay-Per-Click GSP Instability

Unfortunately, there are situations in which pay-per-impression values are monotone, but the PPC-GSP does not have an equilibrium solution. 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 6. *If bidders' values are per-impression, pay-per-click GSP does not always have an equilibrium solution.*

PROOF. We prove the theorem by giving an example for which no equilibrium solution exists using PPC-GSP. 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

$5 - b_{[1]} \le 0 \Rightarrow b_{[1]} \ge 5$. 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 $b_{[1]} \ge b_{[2]} \ge b_{[3]}$.

□

Observe that in this example, the pay-per-click value for slot 1 is 5, while the pay-per-click value for slot 2 is 10. Therefore, the ordering of pay-per-click values is not the same as the ordering of the slot clickabilities. We will see in that this violates the additional assumptions we make in the next section.

5.2 Equilibrium Conditions for PPC-GSP

In this section, we show that if clickthrough rates are independent and distinct, the VCG outcome is attainable in the pay-per-click GSP auction, so long as the values-per-click are monotone and are ranked in the same order as the clickthrough rates for slots. Note that in the Section 4, we used the result that any envy-free outcome was attainable in PPI-GSP. This same result does not hold here, as we saw in the previous subsection. Our proof relies on particular properties of the VCG mechanism.

THEOREM 7. *Suppose that $\theta^1 > \theta^2 > ... > \theta^k$, and that for every bidder ℓ, we have $v^i_\ell/\theta^i \ge v^j_\ell/\theta^j$ for all $i < j$. (i.e., bidder values-per-click and slot clickthrough rates follow the same ordering). Then there exists a set of single bids which leads to the welfare maximizing solution using PPC-GSP, assuming clickthrough rates are independent.*

PROOF. By Lemma 1, we see that (swapping the i and j) $M_{-i} \ge M_{-j} + v^i_j - v^i_i$. Thus, we have for any $i < j$,

$$\begin{aligned}
\frac{p_i}{\theta^i} &= \frac{M_{-i} - M + v^i_i}{\theta^i} \ge \frac{M_{-j} - M}{\theta^i} + \frac{v^i_j}{\theta^i} \quad \text{by Lemma 1} \\
&> \frac{M_{-j} - M}{\theta^j} + \frac{v^i_j}{\theta^i} \quad \text{since } M_{-j} - M \le 0 \text{ and } \theta^i > \theta^j. \\
&= \frac{M_{-j} - M + v^j_j}{\theta^j} - \frac{v^j_j}{\theta^j} + \frac{v^i_j}{\theta^i} \\
&= \frac{p_j}{\theta^j} + \frac{v^j_j}{\theta^j} - \frac{v^i_j}{\theta^i} \ge \frac{p_j}{\theta^j}
\end{aligned}$$

So consider the auction in which bidder 1 bids any value greater than p_1/θ^1, and for $i > 1$, bidder i bids p_{i+1}/θ^{i+1}. Since $p_i/\theta^i > p_{i+1}/\theta^{i+1}$, this produces the correct order with no tie-breaking necessary. Furthermore, bidder i pays precisely p_i, as we wanted. □

Notice that this generalizes the result of [4]. The techniques we use are quite different; although the proof of [4] can be extended to cover a slightly more general case, it 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 attainable by PPC-GSP.

6. DISCUSSION

In this paper, we began a study of equilibria and efficiency in generalized second price auctions when bidders do not have a single value per click for all slots. We have shown that PPI-GSP has several important properties, including existence of equilibria under mild assumptions. Additionally, PPI-GSP is robust in the sense that, regardless of whether a bidder values impressions, clicks, or conversions, then a pay-per-impression system has an equilibrium solution (corresponding to VCG). But the converse is not true: if the GSP auction charges by click or conversion, then there is not necessarily an equilibrium if bidders value slots according to impression. This potentially implies a trade-off. On one hand, accurately assessing value for a slot may be a large burden to place on advertisers as there are many complicated factors involved, such as estimating clickthrough (conversion) rates, risk aversion, and seasonality effects, to name a few (although advertisers must already contend with assessing their values per click). On the other hand, the robustness of the system is an important factor to consider. Charging per impression appears to maintain desirable stability properties under more variable valuation functions.

There are many interesting future directions to explore, which we have not touched on in this paper. The first is a better understanding of other equilibria in the PPI-GSP model when bidders have a full spectrum of values, and their revenue and efficiency properties. Another related direction is that of bidding strategies for bidders: when a bidder has a full spectrum of values, what bidding strategy should he use? A natural bidding strategy which converges to an efficient outcome would be interesting. In [2] a bidding strategy is proposed and it is shown that the strategy converges to a VCG Nash equilibrium. Unfortunately, this bidding strategy does not converge to an equilibrium point in our setting.

So far, we have restricted ourselves to analyzing the existing mechanism (GSP) for a full spectrum of bidder valuations. We have not investigated the design of mechanisms accepting multiple reports, or the efficiency of oblivious strategies as a function of the number of reports allowed. This is a promising area for future work, and has been considered for a very specific form of full spectra in [1].

Finally, it will be interesting to investigate a particular class of non-monotone valuations, where bidders rankings belong to one of two orderings on the slots, one coming from, for example, from valuing clicks, and the other coming from valuing conversions.

7. ACKNOWLEDGEMENTS

We are extremely grateful to Andrei Broder for stimulating conversations that led to this problem. We thank Jon Levin, Paul Milgrom, Mohammad Mahdian, Michael Ostrovsky and Michael Schwarz for insightful comments and pointers to relevant work.

8. REFERENCES

[1] Gagan Aggarwal, Jon Feldman, and S. Muthukrishnan. Bidding to the top: Vcg and equilibria of position-based auctions. *Unpublished Manuscript*.

[2] Matthew Carry, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. *Unpublished Manuscript*, 2007.

[3] Edward H. Clarke. Multipart pricing of public goods. *Public Choice, 11:17-33*, 1971.

[4] Benjamin Edelman, Michael Ostrovsky, and Michael 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] Theodore 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] Hal R. Varian. Position auctions. *To appear in International Journal of Industrial Organization*.

[8] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance, 16:8-37*, 1961.