Position Auctions with Bidder-Specific Minimum Prices

[Extended Abstract]

Eyal Even-Dar
Google Research
evendar@google.com

Jon Feldman
Google Research
jonfeld@google.com

Yishay Mansour
Google Research and
Tel Aviv University
mansour@cs.tau.ac.il

S. Muthukrishnan
Google Research
muthu@google.com

ABSTRACT

Position auctions such as the Generalized Second Price (GSP) are commonly used for sponsored search, e.g., by Yahoo! and Google. We now have an understanding of the equilibria of these auctions, via game-theoretic concepts like Generalized English Auctions and the "locally envy-free" property, as well as through a relationship to the well-known, truthful Vickrey-Clarke-Groves (VCG) mechanism. In practice, however, position auctions are implemented with additional constraints, in particular, bidder-specific minimum prices. Such minimum prices are used to control the quality of the ads that appear on the page.

We study the effect of bidder-specific minimum prices in position auctions. Naïvely enforcing minimum prices in the VCG mechanism breaks the truthfulness of the auction; we describe two variants of VCG for which revealing the truth is a dominant strategy. The implications of bidder-specific minimum prices are more intricate for the GSP auction. Some properties proved for standard GSP no longer hold in this setting. For example, we show that the GSP allocation is now not always efficient (in terms of advertiser value). Also, the property of "envy-locality" enjoyed by GSP—which is essential in the prior analysis of strategies and equilibria—no longer holds. Our main result is to show that despite losing envy locality, GSP with bidder-specific minimum prices still has an envy-free equilibrium.

1. INTRODUCTION

The Internet economy has been revolutionized by the introduction of sponsored search links. Sponsored links are a small number of advertisements (*ads*, henceforth) that the search engine displays in addition to the standard search results. These ads are arranged in *positions* top to bottom, typically on the side. Normally, the advertiser pays only when the user clicks on the link (known as *pay per click (PPC)*). It is a difficult task to set a fixed price for each position because the search queries vary widely and with them the value of the positions. Hence, typically, auctions are used to determine the prices, and these are called *position auctions*. A major task for the search engine is to determine the rules of the position auction, and to select, rank and price the ads that will be displayed to the user, according to that auction.

Today, both Google and Yahoo! use a position auction called the *generalized second price auction* (GSP). The GSP auction ranks the ads by the product of the advertiser's bid with a quality score, which is often abstracted as the *click-through rate (ctr)*—the probability that the user will click on the advertisement. Then, the ad in position i is charged based on the bid of the ad on position $i+1$.

There is a great need to understand the behavior of these auctions since they are part of everyday life of many, with several billions being run each day, worldwide. Decades of research in economic, game and auction theories provide the tools to design and understand auctions. However, position auctions—and GSP in particular—have needed new specific methods such as the recent results of [25] and [3]. Specifically, they developed the notion of *Generalized English Auctions* to study GSP, introduced a new notion of "locally envy-free" equilibrium to characterize GSP, and related such equilibria of GSP to that of the well-known Vickrey-Clarke-Groves (VCG) mechanism (as applied to a position auction). In particular, the "locally envy-free" property captures the dynamics of advertisers trying to move up or down the list of positions and plays a crucial role in understanding the equilibrium properties of GSP.

The departure in our work from prior research begins with the observation that while Google and Yahoo! do implement GSP, they add other features. In particular, they are driven by the need to present high quality ads, and as a result, include features that encourage advertisers to make high-quality ads. One important such feature is the use of advertiser-specific factors for setting *minimum prices*[1]. Beyond the standard use of advertisers' bids and their quality score as in GSP, the search engines force the bid and

[1]See https://adwords.google.com/support/bin/answer.py?answer=49177 and http://help.yahoo.com/help/l/us/yahoo/ysm/sps/start/overview_qualityindex.html

price per click of advertiser i to be at least a minimum price R_i. The immediate impact is not only that advertisers may pay more than what is determined by GSP, but more importantly for sponsored search, because the "heavy tail" of infrequent keywords often has only a few advertisers per query, the minimum price determines whether or not the ad will appear for that query. So, advertiser-specific minimum prices have a profound effect on advertisers, users and search engines in practice.

This motivates the question that is the focus of this paper: What are the strategic changes in the outcome of GSP–and more generally, other position auctions—in presence of advertiser-specific minimum prices? For example, while the introduction of minimum prices looks innocuous, does it affect truthfulness or equilibrium properties of position auctions? A quick sanity check is to study VCG, and doing so immediately reveals that a naïve post-VCG enforcement of bidder-specific minimum bid prices can break the truthfulness property. Being more careful, we show suitably modified allocation and pricing that is a truthful variant of VCG; this modification shows the impact of minimum prices for VCG.

We then turn our attention to GSP, which is the most widely used currently. Since GSP is not truthful to begin with, we study the effect of bidder-specific minimum prices on the equilibria. A simple example with just two bidders shows that minimum prices can cause a loss of efficiency. Furthermore, we see that an important property enjoyed by basic GSP no longer holds: namely, "envy locality." This property says that if a bidder in position i is in a state where she does not envy the bidders in adjacent positions ($i-1$ and $i+1$), then she does not envy any other bidders either. Envy locality is a strong property on its own, as it makes equilibrium discovery simpler for the bidder [3, 25]. Furthermore, it is essential in the existing proof that there is an envy-free equilibrium of GSP.

Our main result, which was also the most technically challenging one, is to show that despite losing envy locality, GSP with bidder-specific minimum prices still has an envy-free equilibrium. To derive the prices of this equilibrium, we define a specialized Tâtonnement process that takes a global view of the best-response relationship between bidders and positions. This global view was unnecessary in the basic GSP analysis such as in [3] because of envy locality. We prove that the process converges to a set of prices from which an envy-free equilibrium set of bids is derived.

In what follows, we will formally define position auctions in the presence of advertiser-specific minimum prices. We first present our observations for VCG to motivate the impact of such minimum prices before presenting our main results for GSP.

1.1 Related Work

Sponsored search has been an active area of research in the last several years after the early papers explored the foundational models [3, 12, 25, 16]. In general, the motivation for this work is that sponsored search in practice is much more complex than as described by the first models. Some papers have taken on the effect of advertiser budgets [5, 21, 2], as well as analyzing bidder strategy and dynamics [4, 23, 6, 10, 28, 27, 17]. There have also been several papers offering extensions to GSP, or entirely new models and mechanisms [13, 15, 19, 11, 24, 20, 1, 9].

The revenue maximization of a single good auction was characterized by Myerson [22] where is shown that in many cases the revenue maximizing auction is a second price auction, where the auctioneer adds a minimum price constraint. Bidder-specific minimum prices have also been studied in the more general context of maximizing revenue in combinatorial auctions [18]. In these models the bidders' valuations are drawn from a distribution, and the minimum prices are set (for the purpose of revenue maximization) when the distribution is known, but not the realization. In our model, the valuations are arbitrary, and the minimum prices are exogenous constraints—we are interested in the dynamics of mechanisms that are forced to satisfy these constraints.

Demange et al [8] consider a multi-item auction in which every buyer is interested in bundles of size at most one; using a Tâtonnement process similar to ours they show how to compute equilibrium prices (but without bidder-specific minima). However in the context of position auctions, their technique is somewhat appropriate for *position-specific* reserve prices, on which we will elaborate in the concluding remarks.

2. MODEL

A position auction is defined by a tuple (N, K, v, α, β). The set $N = \{1, \dots n\}$ is the set of bidders and the set $K = \{1, \dots, n\}$ is the set of positions. Each bidder $i \in N$ is associated with two values, v_i which is its valuation for a click and α_i which is its click trough rate (ctr). Each position $\ell \in K$ is associated with a click through multiplier β_ℓ. As a convention, $\beta_\ell > \beta_{\ell+1}$, $\beta_1 = 1$ and $\beta_{k+1} = \cdots = \beta_n = 0$ (therefore, effectively we have k positions).

We use the standard assumption that the actual click through rate of bidder i in position j is the product of the bidder's ctr α_i and the position click through multiplier β_j, i.e., if bidder i is placed at position j then she receives a click with probability $\alpha_i\beta_j$. We assume that the value of v_i is known only to bidder i while all the other parameters are publicly known.

In a position auction mechanism, each bidder i submits a bid b_i. Given all the bids b and the public information (N, K, α, β), the mechanism assigns each bidder i to a position $\text{loc}(i)$ and charges it a price $P_{\text{loc}(i)}(b, \alpha, \beta)$ per click. The utility of bidder i at position j is:

$$u_i(j) = \alpha_i\beta_j(v_i - p_j),$$

where $p_j = P_j(b, \alpha, \beta)$ is the price per click. The mechanism assigns to each position exactly one bidder, and therefore we have an inverse function $\text{loc}^{-1}(j)$ that returns the bidder that was assigned to position j. We call such an assignment loc a legal assignment.

The two most studied position auction mechanisms are GSP [25, 3] and VCG [26, 7, 14]. The VCG mechanism ranks the bidders by $b_i\alpha_i$, which can be thought of as the expected advertiser value if $b_i = v_i$. Therefore, the VCG allocation maximizes the *social welfare*, which is the sum of the bidders' expected value, i.e., $\sum_{i \in N} v_i \alpha_i \beta_{loc(i)}$. The VCG mechanism charges each bidder the total value lost to other bidders caused by her presence in the auction (we give the exact price function below). The VCG mechanism has the property that the bidders' dominant strategy is to bid their true value, i.e., $b_i = v_i$.[2] Without loss of generality, assume that bidder i is assigned to position i; then the VCG price for position j is [25, 3]

$$P_j^{\mathrm{VCG}} = \sum_{i>j} \frac{b_i \alpha_i (\beta_{i-1} - \beta_i)}{\alpha_j \beta_j}.$$

Note that $\beta_j = \sum_{i>j} (\beta_{i-1} - \beta_i)$ and therefore $P_j^{\mathrm{VCG}} \leq b_{j+1}\alpha_{j+1}/\alpha_j$.

The GSP mechanism ranks the bidders by $b_i \alpha_i$.[3] Again, without loss of generality, assume that bidder i is assigned to position i. The price that the bidder at position i pays per click is $b_{i+1}\alpha_{i+1}/\alpha_i$. It was shown by [25, 3] that for any position auction, there exists an envy-free equilibrium (defined below) such that both the allocation and the payments of the GSP and the VCG mechanism are identical.

While so far we have described the traditional theoretical model, the position auctions used in practice contain an additional important feature, namely bidder-specific minimum prices. The minimum prices imply that each bidder i has a minimum price R_i. This forces both the bid b_i of bidder i to be at least R_i, and the price per click of bidder i to be at least R_i, i.e., its price per click at position $j = loc(i)$ is $\max\{R_i, P_j(b, \alpha, \beta)\} = \max\{R_i, b_{i+1}\alpha_{i+1}/\alpha_i\}$. The focus of our work is to study the effect of bidder-specific minimum prices. We will show that this small modification to the auction mechanism can dramatically influence the behavior of the bidders.

For completeness we define envy free prices specifically for our position auction setting with bidder-specific minimum prices.

DEFINITION 2.1. *Let $P = \{p_1, \ldots, p_n\}$ be a set of prices where $p_{k+1} = \cdots = p_n = 0$ and $loc : N \to K$ be a legal assignment. Then bidder i* envies *the bidder in position j if*

$$\begin{aligned} u_i(loc(i)) &= (v_i - \max\{p_{loc(i)}, R_i\})\beta_{loc(i)} \\ &< (v_i - \max\{p_j, R_i\})\beta_j = u_i(j). \end{aligned}$$

A set of prices P are envy free *if no bidder $i \in N$ envies any another bidder $i' \neq i$. A set of prices P are* locally envy free *if no bidder $i \in N$ envies an adjacent bidder i' where $|loc(i') - loc(i)| = 1$.*

We define an *envy-free equilibrium* as a set of bids for which a proposed mechanism outputs envy-free prices.

[2] A *dominant strategy* is a strategy that a bidder always prefers regardless of the other players' strategies. A mechanism is said to be truthful if revealing the true valuation is a dominant strategy for every bidder.

[3] One can derive the "rank by bid" mechanism by setting $\alpha_i = 1$ for all $i \in N$.

3. VCG AUCTIONS

The VCG mechanism gives a general methodology to implement truthful mechanisms. The mechanism is aimed at maximizing a social welfare function which is the sum of the bidders' utilities. The basic idea of the mechanism is that each bidder pays its marginal influence on the social welfare function of other bidders.

In this section we investigate possible modifications to the VCG payments so that the mechanism will maintain its basic properties (being truthful and efficient) and will also incorporate the bidder-specific minimum prices.

- Our first variation `Naive VCG` simply charges bidder j the maximum of P_j^{VCG} and R_j. This will unfortunately result in a mechanism in which bidders might have incentives to overbid their true valuation, and therefore it is not a truthful mechanism.
- The second approach is to introduce bidder-specific "virtual" bids; namely, when setting the price for bidder j, change the other bidders' bids to reflect the minimum price requirement of bidder j.
- Our third approach is to have a generic reduction, where we subtract the minimum price of a bidder from her bid, run VCG, and then add the minimum price of a bidder to her resulting VCG price. While this general reduction maintain truthfulness, it optimizes a different social welfare function.

Naïve implementation of VCG. We start by defining the most natural implementation of VCG when minimum prices are enforced, `Naïve VCG`. We first compute the VCG prices, and then the price for bidder i is the maximum of its VCG computed price and its minimum price, i.e, the bidder j in position $loc(j)$ pays $\max\{R_j, P_{loc(j)}^{VCG}\}$. Unfortunately, the resulting mechanism is not truthful.

THEOREM 3.1. *The* `Naïve VCG` *mechanism is not truthful.*

PROOF. Assume that all the ctrs are identical, i.e., $\alpha_i = \alpha_j$ for $i, j \in N$. We have four bidders with values $v = (3/2, 5/4, 1/2, 1/4)$ and minimum prices $R = (0, 1, 0, 0)$. Let the position multipliers be $\beta = (1, 1/2, 1/4, 0)$, i.e., three effective positions. For contradiction, assume that the auction is truthful and each bidder bids her value, i.e., $b_i = v_i$. The VCG auction will order them in the sorted order of valuations (since the ctrs are identical) and will price them $13/16, 3/8, 1/4, 0$. Note that only the second bidder has a non-zero minimum price (which is $R_2 = 1$). For the second bidder, the VCG price for the first position, assuming the other bidders bid truthfully, is $(3/2)/2 + (1/2)/4 + (1/4)/4 = 15/16$ and for the second position it is $[(1/2)/4 + (1/4)/4]/(1/2) = 3/8$. Since the second bidder has a minimum price of $R_2 = 1$, she will pay in either position a price of 1. This implies that the second bidder would prefer to overbid and be assigned the first position. □

Virtual Values. Since the naïve approach to incorporating bidder-specific minimum prices fails, we would like to explore another approach. We first make the observation that if for some bidder i, every other bidder i' with $\mathrm{loc}(i') > \mathrm{loc}(i)$ had $b_{i'}\alpha_{i'} \geq R_i\alpha_i$, then the (unmodified) VCG price for i would be at least R_i. (This follows from simple manipulations of the VCG price definition.) This observation motivates introduction of bidder-specific "virtual" values: When computing the price for bidder i, we use $\max\{b_{i'}\alpha_{i'}, R_i\alpha_i\}$ as a substitute for $b_{i'}\alpha_{i'}$ for all applicable i'. This implies that the bid of a bidder i' is interpreted differently when computing prices of different bidders. We define the mechanism formally as follows:

Virtual Values$(b_1, ..., b_n, R_1, ..., R_n)$
Sort $b_i\alpha_i$ in descending order and assign bidder i to its position in the sorted list;
For every bidder j at $loc(j)$ charge

$$P^{VV}_{loc(j)} = \sum_{i:\mathrm{loc}(i)>\mathrm{loc}(j)} \frac{\max\{b_i\alpha_i, R_j\alpha_j\}(\beta_{\mathrm{loc}(i)-1} - \beta_{\mathrm{loc}(i)})}{\alpha_j\beta_{loc(j)}}$$

THEOREM 3.2. *The* `Virtual Values` *mechanism is efficient and truthful.*

PROOF. Consider a bidder j placed in position $\mathrm{loc}(j)$ when she bids her true valuation v_j. We will show that position $\mathrm{loc}(j)$ maximizes her utility and thus revealing the true valuation is a weakly dominant strategy. First, we show that for any position $m > \mathrm{loc}(j)$ bidder j's utility from position $m-1$ is at least her utility from position m.

Let $P^{VV}_{m,j}$ be the price of bidder j for position m, assuming all the other bidders do not change their bid. Therefore, the utility of bidder j from position m is $u_j(m) = \alpha_j\beta_m(v_j - P^{VV}_{m,j})$. We would like to show that $u_j(m) \leq u_j(m-1)$ for $m > \mathrm{loc}(j)$. This is equivalent to,

$$\begin{aligned} u_j(m-1) &= v_j\alpha_j\beta_{m-1} - (\beta_{m-1} - \beta_m)\max\{b_m\alpha_m, R_j\alpha_j\} \\ &\quad -B \geq v_j\alpha_j\beta_m - B = u_j(m), \end{aligned}$$

where $B = \sum_{i>m} \max\{b_i\alpha_i, R_j\alpha_j\}(\beta_{i-1} - \beta_i)$. Now rearranging the terms we obtain,

$$v_j\alpha_j(\beta_{m-1} - \beta_m) \geq (\beta_{m-1} - \beta_m)\max\{b_m\alpha_m, R_j\alpha_j\}\ .$$

Since the list is sorted we have $v_j\alpha_j \geq b_m\alpha_m$. In addition $v_j \geq R_j$ otherwise bidder j will never participate. Therefore the inequality holds and bidder j prefers position $\mathrm{loc}(j)$ over any position $m > \mathrm{loc}(j)$.

Now consider a position $m < \mathrm{loc}(j)$. We show that bidder j utility from position $m+1$ is larger than her utility from position m, i.e., $u_j(m+1) \geq u_j(m)$. In this case we observe that $b_m\alpha_m \geq v_j\alpha_j \geq R_j\alpha_j$ and therefore $\max\{b_m\alpha_m, R_j\alpha_j\} = b_m\alpha_m \geq v_j\alpha_j$. Therefore, using essentially the same inequalities, we have that $u_j(m) \geq u_j(m-1)$. This implies that bidder j prefers position $\mathrm{loc}(j)$ over any position $m < \mathrm{loc}(j)$, which completes the proof. □

Note that although the `Virtual Values` mechanism is truthful the equilibrium prices are not envy-free, and even worse, they are not monotone in the position. Consider the following example: we have two positions with multipliers $\beta = (1.0, 0.5)$ and two bidders with valuations $(10, 5)$ and minimum prices $(.05, 4.95)$. Assume that the bidders have identical ctr, i.e., $\alpha_1 = \alpha_2 = 1$. The `Virtual Value` mechanism assigns bidder 2 to position 2 and charges her 4.95 while bidder 1 is assigned to position 1 and charged only 2.525.

Offsetting Bid by Minimum Price. We now present a generic approach of incorporating minimum prices. The idea is to reduce the problem to a setting in which there are no minimum prices. Here is the simple reduction, where we assume that A is a truthful auction:

Subtract Min Price$(A, b_1, ..., b_n, R_1, ..., R_n)$
Let $b'_i = b_i - R_i$.
Run $A(b'_1, \ldots, b'_n)$ and get a legal assignment loc_A and set of prices P^A.
The price of bidder i is $P_i = P^A_i + R_i$ and its position is $\mathrm{loc}_A(i)$.

THEOREM 3.3. *Assume that A is a truthful mechanism when bidder i's valuation is $v'_i = v_i - R_i$. Then* `Subtract Min Price` *is a truthful mechanism.*

PROOF. The important observation is that for any bidder, her utility with valuation v_i in `Subtract Min Price` is identical to her utility in A when she has valuation $v'_i = v_i - R_i$. Namely her utility from being assigned to position $\ell = \mathrm{loc}_A(i)$ is $\alpha_i\beta_\ell(v_i - R_i - P^A_i)$. Therefore the truthfulness of `Subtract Min Price` follows from that of A. □

We remark that `Subtract Min Price` is efficient with respect to the social welfare function $\alpha_i(v_i - R_i)$, which can be very different from $\alpha_i v_i$. Note also that the revenue of the auctioneer can also be very small in some scenarios. For instance, assume two bidders with valuations $(100, 3)$, minimum prices $(99, 1)$ and identical ctrs, and there is only one effective position. The `Subtract Min Price` assigns bidder 2 to position 1 and charges her 2.

The VCG mechanism above optimizes a function interpreted as follows. Consider the auctioneer as an additional player in the game, with utility $-R_i$ for any click of bidder i. The `Subtract Min Price` mechanism implements an efficient solution for the $N+1$ players thus obtained. This is due to the fact that the value of player i at position j plus the value the auctioneer draws from that placement is $v_i\alpha_i\beta_j - R_i\alpha_i\beta_j = (v_i - R_i)\alpha_i\beta_j$. This utility function is reasonable when there is a cost to the auctioneer for any click on the ith advertiser, which is the future effect of allowing low quality ads.

4. GENERALIZED SECOND PRICE (GSP) AUCTIONS

In this section we will discuss the effect of introducing minimum bid prices to Generalized Second Price (GSP) auctions. It is well known that GSP is not a truthful mechanism, so our main focus is to show that there exists an equilibrium. In fact we will show a stronger result, that there are envy-free prices for the GSP.

The existence of envy-free prices for GSP with no minimum bids (or a uniform minimum bid which is identical to all the bidders) was shown in [3, 25]. Along the way, these analyses show a few interesting properties of GSP. The first is that there are envy-free prices which result in an efficient allocation (i.e., maximize the sum of bidders valuations). The second is the fact that local envy-free prices imply (global) envy-free prices, the property of "envy locality" we discussed earlier. We show that both of those properties do not hold once bidder-specific minimum prices are introduced.

For the proof of the existence of envy-free prices we use a specific Tâtonnement process. Our process increases prices while ensuring that certain properties of the allocation are maintained.

4.1 GSP Efficiency

The following theorem shows that there is an example where every equilibrium in GSP mechanism is not efficient.

THEOREM 4.1. *The GSP mechanism with bidder specific minimum prices is not necessarily efficient.*

PROOF. Consider the following case of two bidders and two positions. Let $\beta = (1, 1/2)$, $v = (1, 2/3 + \epsilon)$, $R = (0, 2/3)$, and identical ctrs. We will show that there is no set of equilibrium bids which will maintain efficiency. Since each bid has to be at least the bidder specific minimum price, we have $b_2 \geq R_2 = 2/3$. Bidder 1's utility at the first position is $(1 - b_2) \leq 1/3$ and in the second position is $1/2$, since the price would be zero. Therefore, player 1 would underbid player 2 in any equilibrium. □

4.2 GSP: Local envy-free prices

We provide the following important observation that highlights the differences between the behavior in GSP with and without bidder-specific minimum prices. While in basic GSP it was shown that locally envy free prices imply globally envy free prices [3, 25], this is not true anymore when minimum prices are used. This means that we will need to argue more globally and cannot rely on local analysis to be sufficient. As a consequence, the proof techniques that were used in the previous analysis of the basic GSP cannot be applied to GSP with bidder specific minimum prices.

THEOREM 4.2. *In GSP with bidder specific minimum prices, locally envy free prices do not imply envy free prices.*

PROOF. Consider an instance where three bidders with valuations $v = (12, 11, 20)$, minimum prices $R = (10, 10, 0)$, position multipliers $\beta = (1, 1/2, 1/4)$, and identical ctrs. Consider bids $(10.5, 10.5, 10)$. This generates prices $P = (10.5, 10, 0)$ and utilities $u = (3/2, 1/2, 5)$. First we need to verify that the prices are local envy-free prices. The utility of bidder 1 at the second position is $2 \cdot 1/2 = 1$ and thus she does not envy. The utility of bidder 2 at the first position is still $1/2$ and thus she does not envy bidder 1. We still need to show that bidder 2 and bidder 3 do not envy each other. The utility that bidder 2 would get at position 3 is $(11-10)/4 = 1/4$ and thus she does not envy bidder 3. The utility that bidder 3 would get at position 2 is $(20-10)/2 = 5$ and thus she does not envy bidder 2. Hence the prices are locally envy-free prices.

However bidder 3 is better off getting the first position at price 10.5 as her utility would be 9.5. Therefore the prices are not envy-free prices.[4] □

4.3 GSP: Envy-free prices

The remainder of the section will be devoted to proving our main theorem, extending the existence of envy-free equilibria to the case of bidder-specific minimum prices:

THEOREM 4.3. *The GSP mechanism with bidder-specific minimum prices has an envy-free prices equilibrium.*

The proof technique that we will use to show Theorem 4.3 is to define a specific Tâtonnement process, and show that it converges to a set of envy-free prices.

4.3.1 *Definitions and notation.*

We start by giving the definitions and notation that will be used in the proof. Let $K' = \{1, \ldots, k\}$ be the set of positions with non-zero multiplier, i.e., $\beta_i > 0$. Given a price vector P for any subset of bidders $B \subset N$ and subset of positions $S \subseteq K'$ we define the *best response graph*, $G(P, B, S) = (B, S, E)$. The graph $G(P, B, S)$ is a bipartite graph where $(b, s) \in E$ if and only if position $s \in S$ is a best response for bidder $b \in B$. We say that positions $i \in S$ and $j \in S$ are *connected* if there exists a path between i and j in $G(P, B, S)$. We denote by $\nu_G(v)$ the neighbors of a node v in $G = G(P, B, S)$. We use the notation $P' = (P, \epsilon, j)$ to denote a price update of position j by ϵ, i.e., $p'_i = p_i$ for every $i \neq j$ and $p'_j = p_j + \epsilon$. We also let $S_{\rm NE}(P) \subseteq K'$ be the set of positions that are a best response for at least one bidder at the prices P; equivalently $S_{\rm NE}(P)$ is the set of position nodes $s \in K'$ in $G = G(P, N, K')$ with at least one incident edge, i.e., $\nu_G(s) \geq 1$. We say that a set $S \subseteq K'$ is *matched* in $G = G(P, N, S)$ if there is a perfect matching in $G(P, B, S)$ for some $B \subseteq N$. To simplify our notation, whenever P, B, S or G are clear from the context we might omit them. We will also assume in this subsection that all α_i's are equal and that all v_is are different; in Section 4.5 we discuss the extension to arbitrary α_i's.

4.3.2 *The Tâtonnement process*

Before describing the Tâtonnement process formally, we provide some useful intuition. The Tâtonnement process begins with a set of prices P_1 such that all first k bidders prefer the first position; i.e., $S_{\rm NE}(P_1) = \{1\}$, $B_1 = \{1, \ldots, k\}$ and $G(P_1, B_1, K')$ is a star graph where each bidder $i \in B_1$ has exactly one edge to the node for position 1. The Tâtonnement process gradually increases prices, increasing the price of only one position during each update. While increasing the prices the algorithm preserves two invariants:

[4]Note that there are equilibria in this case, for example bids $b = (15, 10.5, 10)$ or $b = (11, 10.5, 20)$.

1. At each step, with prices P_t, the set of positions $S_t = S_{\rm NE}(P_t)$ that are the best response for some bidder can only grow; i.e., $S_t \subseteq S_{t+1}$.
2. There is a matching of the positions S_t, such that every position in S_t can be matched to a unique bidder in $G(P_t)$.

Both invariants are preserved by maintaining the conditions of Hall's theorem on every subset of $S_t = S_{\rm NE}(P_t)$, i.e., for every subset $S' \subset S_t$ we require that $|S'| \le |\nu(S')|$ which is a sufficient and necessary condition for a matching by Hall's theorem, which is given here for completeness.

THEOREM 4.4 (HALL'S THEOREM). *A set $S \subset K'$ is matched in $G = G(P, N, S)$ iff for every $S' \subset S$ we have $|S'| \le |\nu_G(S')|$.*

We now present our Tâtonnement process TP:

Tâtonnement process TP
Initialize P_1 such that $p_j = v_{k+1}$ for $j \le k$ and $p_j = 0$ for $j \ge k+1$;
Let $t = 1$ and $S_1 = \{1\}$;
while $\exists \epsilon > 0, j \in S_t$: *MATCH*$(P_t, \epsilon, j) = $ *TRUE* **do**
For each $j \in S_t$ let $\epsilon_j = \max\{\epsilon : \text{MATCH}(P_t, \epsilon, j)\}$;
$s_t = \arg\max_{j \in S_t} \epsilon_j$;
$\epsilon_t = \epsilon_{s_t}$;
$P_{t+1} = (P_t, \epsilon_t, s_t)$;
$S_{t+1} = S_{\rm NE}(P_{t+1})$;
$t = t + 1$;
end
Output the set of price P_t and the allocation is a matching in $G(P_t, N, K)$.

MATCH(P, ϵ, j) = TRUE iff there is a matching for $S' = S_{\rm NE}(P') \cup \{j\}$ in $G(P', N, S')$, where $P' = (P, \epsilon, j)$.

We first show that the Tâtonnement process TP cannot loop indefinitely if all numbers are rational. This is done by showing that there exists ϵ_{min}, which is a function of $v_1, \ldots, v_n$, $R_1, \ldots, R_n$, $\beta_1, \ldots, \beta_k$, and $\alpha_1, \ldots, \alpha_n$, where every increase will be at least ϵ_{min}.

LEMMA 4.5. *The Tâtonnement process TP always terminates*

PROOF. Since all numbers are rational then we can normalize the number by multiplying them by their lowest common denominator and in the rest of the proof we will assume they are integers. Next we show that this implies that if we update the price by ϵ then we can update the price by $\lceil \epsilon \rceil$. We prove the claim by induction and assume that the current prices are all integers. Let $P' = (P, \epsilon, j)$ and $P'' = (P, \lceil \epsilon \rceil, j)$. We prove that the edges in $G(P'')$ contain the edges in $G(P')$. Assume by contradiction that the edge (i, j) is in $G(P')$ and not in $G(P'')$, thus we have that

$$\begin{aligned} \beta_\ell(v_i - \max\{p'_\ell, R_i\}) &\le \beta_j(v_i - \max\{p'_j, R_i\}) \\ \beta_\ell(v_i - \max\{p''_\ell, R_i\}) &> \beta_j(v_i - \max\{p''_j, R_i\}) \end{aligned}$$

Since the last equation is an integer strict inequality we can rewrite it as

$$\beta_\ell(v_i - \max\{p''_\ell, R_i\}) \ge \beta_j(v_i - \max\{p''_j, R_i\}) + 1.$$

This implies now that

$$\beta_j(v_i - \max\{p''_j, R_i\}) + 1 \le \beta_j(v_i - \max\{p'_j, R_i\}).$$

Rewriting again, using the fact that $\beta_j \le 1$, we get that

$$\max\{p'_j, R_i\} + 1 \le \max\{p''_j, R_i\},$$

which is impossible since $p''_j - p_j = \lceil \epsilon \rceil - \epsilon < 1$. □

The above lemma shows that the Tâtonnement process TP always terminates. Now we would like to prove a few facts on how TP makes progress, until it terminates. Specifically, we would like to show that the set S_t increases monotonically and furthermore, each time it changes it adds the least position which is not in S_t. Therefore, initially we have $S_1 = \{1\}$, and at any time t we will have $S_t = \{1, \ldots, j\}$ for some $j \in K'$. The following two lemmas establish this property.

The following observation shows the effect of a price increase at position j, which is the basic step of the Tâtonnement process TP:

LEMMA 4.6. *Let $P' = (P, \epsilon, j)$, $(N, K', E) = G(P, N, K')$ and $(N, K', E') = G(P', N, K')$. Then every edge in $(b, i) \in E - E'$ is incident to j; i.e., $i = j$. Also, for each edge $(b, s) \in E' - E$ there is an edge $(b, j) \in E$.*

PROOF. Clearly increasing the price of position j can only lower a bidder's utility from position j, while keeping the utility from other position unchanged. Therefore, any bidder whose best response set does not include j is unaffected. A bidder whose best response includes j can now either remove the edge to j or add an edge to another position which became as attractive as j to her. □

The following lemma shows that we preserve the invariant that S_t monotonically grows. Furthermore, since the positions have a strict preference order which is shared by all the bidders, the sets S_t are prefixes of $[1, \ldots, k]$ and can grow by at most one position each time step:

LEMMA 4.7. *Let $S_t = S_{NE}(P_t)$ at time t. Then, for every time $t' > t$ we have $S_t \subseteq S_{t'}$. In addition, if $j \in S_t$ then any $i \le j$ has $i \in S_t$, and if $S_t \ne S_{t+1}$ then $|S_{t+1} - S_t| = 1$.*

PROOF. We prove it for times t and $t+1$ where we increased the price at position s_t. By Lemma 4.6 the only edges that can be in $G(P_t)$ and not $G(P_{t+1})$, are those which are adjacent to s_t, we only need to guarantee that s_t is still a best response for at least one bidder. By the construction of the Tâtonnement process TP we have that MATCH(P_t, ϵ_t, s_t) = TRUE. This implies that there is a matching in $G = G(P_{t+1}, N, S')$ where $s_t \in S'$, which clearly implies that $|\nu_G(s_t)| \ge 1$. For the second part of the lemma, we only need to prove now that if j' is added to S_t then $j' - 1$

is already in S_t. We prove it by induction, when the basis trivially holds with $S_1 = \{1\}$. When we add a new position note that all positions that are not in $S_{\rm NE}$ and have positive multiplier have the same price, now since the multiplier are strictly decreasing, then every bidder in $\{1, \ldots, k\}$ strictly prefers the lowest position (highest ctr) that is currently not in S_t.

For the same reason, when we add a position it has to be the lowest position which has price v_{k+1}, and therefore if $S_{t+1} \neq S_t$ then we have $|S_{t+1} - S_t| = 1$. □

4.3.3 Conditions when the Price Increases

We have shown that the Tâtonnement process terminates, and maintains key invariants. Since we are maintaining a matching for the set S_t, we essentially just need to show that when the Tâtonnement process TP terminates, we have $S_t = K'$. By Lemma 4.7 it is sufficient to show that at some time $S_t = K'$, since S_t is monotone. Thus the lemmas in this section are steps to show that if $S_t \neq K'$ then some price can be increased, and therefore the Tâtonnement process does not terminate.

The following simple fact regarding bipartite graphs, which will be useful for the lemmas in this section, claims that the matching property is monotone.

LEMMA 4.8. *Let $G = (B, S, E)$ be a bipartite graph in which there is a matching for S, then for any $G' = (B, S, E')$, such that $E \subseteq E'$, there is a matching for S.*

First, consider the case that there is a bidder which has only one position as a best response, we show that in this case we increase the prices and thus the Tâtonnement process TP cannot terminate.

LEMMA 4.9. *Let $G = G(P_t, N, S_t)$ be a best response graph such that $S_t \neq K'$. If for some bidder $b \in N$ we have $\nu_G(b) = \{j\}$, then there exists an $\epsilon > 0$ such that MATCH(P_t, ϵ, j) = TRUE.*

PROOF. Since $S_t \neq K'$ it implies (by Lemma 4.7) that position k is empty, i.e., $\nu_G(k) = \emptyset$. Since $\beta_k > 0$ it implies that bidder b utility from position k is positive, and therefore its utility from position j is positive as well. Let $\delta = \beta_j(v_b - \max\{R_b, p_j\}) - \max_{i \neq j} \beta_i(v_b - \max\{R_b, p_i\})$. Since $\nu_G(b) = \{j\}$ it implies that $\delta > 0$. Let $\epsilon = (1/2)\delta/\beta_j$ and $P' = (P_t, \epsilon, j)$. We need to show that MATCH(P_t, ϵ, j) = TRUE. Note that $\nu_{G(P')}(b) = \{j\}$, by the definition of δ. By Lemma 4.6 we only added edges to other positions, and hence by Lemma 4.8 we still have a matching. □

We say that prices P *induce equal payments* if for any position $s \in S_{\rm NE}(P)$ and any two bidders i, i' for which s is a best response, then $\max\{R_i, p_s\} = \max\{R_{i'}, p_s\}$. For the process it is important to distinguish between prices which induce equal payments and ones which do not. The next lemma claims that in certain subgraphs if P_t do not induce equal payments then they can be increased.

LEMMA 4.10. *Let P_t be prices which do not induce equal payments, and $G = G(P_t, N, S_t)$ be a best response graph, where $S_t \neq K'$. If in G every subset S' of S_t satisfies $|S'| \leq |\nu_G(S')| - 1$, $|S_t| \geq 2$ and G is connected then there exists a position $j \in S_t$ and $\epsilon > 0$ such that MATCH(P_t, ϵ, j) = TRUE.*

PROOF. Since P_t do not induce equal prices, there is a position j whose set of best response bidders I includes two bidders, i and i', who pay different prices for j, i.e., $\max\{R_i, p_j\} \neq \max\{R_{i'}, p_j\}$. Let $\epsilon = \max_{i \in I}\{R_i\} - p_j$ and $P' = (P_t, \epsilon, j)$. Let $G' = G(P', B, S)$ be the graph after changing the prices to P' and let $I'\nu_{G'}(j)$ be the new set of bidders for which j is a best response. If $I = I'$, i.e., best response set did not change, then we still have a matching by Lemma 4.8. Now consider the case where the edge (i, j) was removed from the graph G when creating G'.

Since for every subset S' of S_t satisfies: $|S'| + 1 \leq |\nu_G(S')|$, and $|\eta_{G'}(\{j\})| \geq 1$ then (S', j) will still satisfy the matching constraint in G' and thus we could have increased the price. □

Next we claim that if the prices induce equal payments, then two bidder can have at most one position in the intersection of their best response sets.

LEMMA 4.11. *Let P_t be prices which induce equal payments, and let $G = G(P_t, N, S_t)$ be a best response graph. For any bidders $i, i' \in N$ we have $|\nu_G(i) \cap \nu_G(i')| \leq 1$.*

PROOF. Assume towards a contradiction that it is not the case then there are two bidders ℓ and m such that $v_\ell > v_m$, the intersection of their best response sets include positions i and i' and both pay the same for the two positions. Therefore we have that,

$$\begin{aligned} (v_\ell - P_i)\beta_i &= (v_\ell - P_{i'})\beta_{i'} \\ (v_m - P_i)\beta_i &= (v_m - P_{i'})\beta_{i'} \end{aligned}$$

Subtracting the equations we obtain that $(v_\ell - v_m)\beta_i = (v_\ell - v_m)\beta_{i'}$ since we have that $v_\ell \neq v_m$ we can divide by $v_\ell - v_m$ and obtain that $\beta_i = \beta_{i'}$ and arrive at a contradiction. □

Next we present technical lemma which will be crucial in the theorem proof. This lemma shows that if in a subgraph every subset of positions has a slack with respect to Hall's theorem condition, i.e. each subset of positions has strictly more bidders connected to in the graph, then there exists a price we can increase without violating the matching constraint.

LEMMA 4.12. *Let $G = G(P_t, B, S_t)$ be a best response graph. If in G every non-empty subset S' of S_t satisfies $|S'| \leq |\nu_{G(P)}(S')| - 1$, $|S| \geq 2$ and G is connected then there exists some $j \in S_t$ and $\epsilon > 0$ such that MATCH(P_t, ϵ, j) = TRUE.*

PROOF. For contradiction, assume that there is no position $j \in S_t$ for which we can increase the price without violating the matching constraint. I.e., for every $j \in S_t$ and $\epsilon > 0$ we have that S_t is not matched in $G' = G(P', N, S_t)$ where $P' = (P, \epsilon, j)$. By Lemma 4.10 it has to be the case that P_t are prices which do induce equal payments (otherwise we can increase the prices).

By Lemma 4.9 every bidder has at least degree 2. Therefore there is a cycle in the graph G. By Lemma 4.11 two bidders can have only one position in common, therefore the cycle contains at least three positions. Let j_3 be the largest position that appears in the cycle. Let ℓ and m be the two bidders that are the neighbors of j_3 in the cycle, and assume that $v_\ell > v_m$. Let j_1 and j_2 be the positions connected to ℓ and m, respectively, in the cycle.

First, we will show that $j_1 < j_2$. Assume for contradiction that $j_2 < j_1$. For bidder ℓ we have $(v_\ell - p_{j_2})\beta_{j_2} < (v_\ell - p_{j_1})\beta_{j_1}$ and for bidder m we have $(v_m - p_{j_1})\beta_{j_1} < (v_m - p_{j_2})\beta_{j_2}$. Adding the inequalities we get, $v_\ell\beta_{j_2} + v_m\beta_{j_1} < v_\ell\beta_{j_1} + v_m\beta_{j_2}$ which implies that $v_m(\beta_{j_1} - \beta_{j_2}) < v_\ell(\beta_{j_1} - \beta_{j_2})$. Since $v_m < v_\ell$ we need to require that $\beta_{j_2} < \beta_{j_1}$ which is equivalent to $j_2 > j_1$. Therefore we reached a contradiction, and the order has to be $j_1 < j_2 < j_3$.

Consider the j_1, j_2, j_3. First note that $p_{j_1} > p_{j_2} > p_{j_3}$ otherwise bidder m will strictly prefer j_1 over j_2. Since the prices P_t induces equal payments, bidder m pays at position j_2 price p_{j_2} and bidder ℓ pays price p_{j_1}, i.e., they are not using their minimum prices. Recall that both pay the same price at position j_3. The following must hold due to the fact that they are both indifferent

$$\begin{aligned}(v_\ell - p_{j_1})\beta_{j_1} &= (v_\ell - p_{j_3})\beta_{j_3}\\ (v_m - p_{i_2})\beta_{j_2} &= (v_m - p_{j_3})\beta_{j_3}\end{aligned}$$

Now rewriting the last equation and using the fact that $v_\ell > v_m$ we have that

$$v_\ell(\beta_{j_2} - \beta_{j_3}) > v_m(\beta_{j_2} - \beta_{j_3}) = p_{j_2}\beta_{j_2} - p_{j_3}\beta_{j_3}$$

rewriting again $(v_\ell - p_{j_2})\beta_{j_2} > (v_\ell - p_{j_3})\beta_{j_3}$ which contradicts the fact that positions j_1 and j_3 are part of bidder ℓ best response set. We reached a contradiction to the assumption that the prices can not be increased. □

4.3.4 *Proof of Theorem 4.3*

Now all the pieces are in place to prove Theorem 4.3.

PROOF OF THEOREM 4.3. Suppose toward a contradiction that the algorithm halts at price P_t and there is no envy free mapping. This implies that there are empty positions, and by Lemma 4.7 $k \notin S_t$. Since there are empty positions, i.e., $S_t \neq K'$, there exists a position which is in more than one bidder's best response set, and its connected component, $C = (B, S, E)$ in $G(P_t)$, satisfies $|S| + 1 \leq |\nu_G(S)|$. We know that we cannot increase the price of any position in S without violating the matching constraint, otherwise the Tâtonnement process TP would not have terminated. We consider two possible scenarios: one is when S contains a single position and multiple bidders and the second is when S contains several positions.

We start with S being a singleton, i.e. $S = \{j\}$ and show that the price of j can be increased a contradiction to the fact that the process terminated. Since $S = \{j\}$, every bidder $i \in \nu_G(j)$ must satisfy $|\nu_G(i)| = 1$. Therefore, by Lemma 4.9 we can increase the prices.

Now we know that S is not singleton. Instead of considering $C = (B, S, E)$ we would consider its subset $\bar{C} = (\bar{B}, \bar{S}, \bar{E})$. We obtain $\bar{C}$ by removing bidders $B' \subset B$ and positions $S' \subset S$ from C. Specifically, we iteratively remove every subset S' of positions in the remaining subgraph if $|S'| = |\nu(S')|$ in the subgraph.[5] At the end we define $\bar{C}$ as the largest connected component of the resulting subgraph. From now on on we abuse notation by considering only best response inside $\bar{C}$. Note that for every subset of positions S' in $\bar{C}$ we have $|S'| + 1 \leq |\nu_G(S')|$. In addition we have that by changing prices inside $\bar{C}$ the only way to violate the matching constraint is within a subset inside $\bar{C}$ since by adding any subset of positions outside $\bar{C}$ we always increase the amount of bidders by at least the same amount. By Lemma 4.12 this cannot happen if $|\bar{S}| \geq 2$ and if $|\bar{S}| = 1$ then it's impossible from similar reasoning to the singleton analysis. Therefore, we cannot stop when there are empty positions. Clearly, if there are no empty positions by Hall's theorem there exists a perfect matching and thus an envy free mapping in prices P_t.

It still remains to show that these prices can be produced by bids. Now if the bids are $b_i = p_{i-1}$, then they do indeed produce the prices, and since $p_{i-1} > p_i > R_i$, the proof follows. □

Note that in the Tâtonnement process TP we do not necessarily terminate with a price vector P_t which induce equal payments. We might terminate with prices P_t which do not induce equal payments, since we already reached a state in which $S_t = K'$. Our proof technique only shows that as long as $S_t \neq K'$ we can increase the price of some position.

4.4 GSP: revenue

It is known that in the setting of position auctions without minimum prices, the bidder-optimal equilibrium of GSP is equivalent to the truthful outcome of VCG. So it is natural to ask how this carries over when we enforce minimum prices. This is largely still open, but we can prove that under certain conditions, enforcing minimum prices can improve revenue:

THEOREM 4.13. *Let P^{EF} be envy-free prices for GSP with minimum prices and let P^{VCG} be the VCG prices without minimum price. Then $P^{EF}_j \geq P^{VCG}_j$, assuming that for every bidder i, $v_i > R_i$.*

PROOF. We prove it by backward induction on the prices of P^{VCG} and P^{EF}. For the basis we have that the player in the kth position pays at least the maximum between her minimum price and the valuation of the $k+1$ bidder (otherwise bidder $k+1$ can outbid her and get a positive utility).

Assume the induction holds for all positions $\ell' > \ell$, i.e., $P^{EF}_{\ell'} \geq P^{\text{VCG}}_{\ell'}$, and prove for ℓ. Let j be the bidder with

[5] Recall that since we have a matching for S_t in $G(P_t)$ then by Hall's Theorem we have that $|S'| \leq \nu_G(S')$.

highest valuation in positions $\{\ell+1,\dots,k\}$. Clearly, $v_j \geq v_{\ell+1}$, and let the position of bidder j be m. Since the prices P^{EF} are envy free, we have that

$$\beta_m(v_j - P_m^{EF}) \geq \beta_\ell(v_j - P_\ell^{EF}),$$

rearranging we obtain that

$$\begin{aligned} P_\ell^{EF} &\geq \frac{\beta_m P_m + v_j(\beta_\ell - \beta_m)}{\beta_\ell} \geq \frac{\beta_m P_m^{\mathrm{VCG}} + v_j(\beta_\ell - \beta_m)}{\beta_\ell} \\ &\geq P_\ell^{\mathrm{VCG}}, \end{aligned}$$

where we used in the last inequality that

$$\begin{aligned} P_\ell^{\mathrm{VCG}} &= \sum_{i=\ell+1}^{k} (\beta_{i-1} - \beta_i) v_i = P_m^{\mathrm{VCG}} + \sum_{i=\ell+1}^{m} (\beta_{i-1} - \beta_i) v_i \\ &\leq P_m^{\mathrm{VCG}} + (\beta_\ell - \beta_m) v_j \end{aligned}$$

□

4.5 Extension to unequal α

While we assumed that all α_is are identical this can be fixed easily by running the same algorithm only this time charging bidder j for position i $\max\{p_i/\alpha_j, R_j\}$. Now one can verify that all the algorithm properties still hold. To verify that bids can produce these prices note that we need that $b_{i+1}\alpha_{i+1} = p_i$, rewriting it we get that $b_{i+1} = p_i/\alpha_{i+1} \geq p_{i+1}/\alpha_{i+1} = R_{i+1}$ and thus the algorithm produces the envy free prices when different players have different α's.

5. CONCLUDING REMARKS

Position auctions are important because search engine companies use them to sell advertisements in sponsored search. Recently, [25, 3] and others developed mathematical notions to understand the equilibrium properties of such auctions, in particular, that of GSP which is widely used. Our work here adds to the understanding of position auctions started in [25, 3] by incorporating a practical feature that search engines widely use, namely, bidder-specific minimum prices. This feature is not innocent since if not handled properly in VCG pricing, it leads to loss of truthfulness; in case of GSP, it results in key technical difficulties, but, as we prove in this paper, it still leads to envy-free equilibrium.

There are a number of open technical and conceptual open issues. While in this paper we considered the bidder-specific minimum prices as are currently used by Google, there are other natural ways to enforce quality:

- *Position-specific minimum prices (only).* In this case, there is simply a reserve price for each position. One is then tempted to apply the multi-item auction of Demange et al [8] which works with reserve prices to compute an envy-free price vector. However the problem is that some positions can go unsold; indeed it sometimes impossible to sell all positions in an envy-free equilibrium.[6] While this is allowable in the context of an abstract multi-item auction, for sponsored search it is not desirable to have empty slots. It would be interesting to see if there are other alternatives that can give a more suitable solution to this case.
- *Minimum ctr.* Here each bidder has a minimum click-through-rate (ctr) required to participate in the auction. If the ctr is exogenous, then this simply applies filter to the set of bidders and this case reduces to the others. However, if the bidder has a choice of which ad to use, with varying ctrs (and values), then this becomes potentially interesting.

These quality-control techniques (or combinations thereof) require more research and may have different structural properties than the bidder-specific minimum prices case we studied and solved here.

Acknowledgments. We thank Hal Varian and Martin Pál for useful discussions.

6. REFERENCES

[1] Z. Abrams and A. Ghosh. Auctions with revenue guarantees for sponsored search. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.

[2] Z. Abrams, O. Mendelevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. In *ACM Conference on Electronic Commerce*, pages 272–278, 2007.

[3] M. Ostrovsky B. Edelman and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In Second workshop on sponsored search auctions, 2006.

[4] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proc. WWW*, 2007.

[5] C. Borgs, J. T. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*, 2005.

[6] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. R. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. In *ACM Conference on Electronic Commerce*, pages 262–271, 2007.

[7] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[8] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *The Journal of Political Economy*, 94:863–872, 1986.

[9] E. Even-Dar, M. Kearns, and J. Wortman. Sponsored search with contexts. In *WINE*, pages 312–317, 2007.

[10] J. Feldman, S. Muthukrishnan, M. Pal, and C. Stein. Budget optimization in search-based advertising auctions. In *ACM Conference on Electronic Commerce*, pages 40–49, 2007.

[11] J. Feldman, E. Nikolova, S. Muthukrishnan, and M. Pál. A truthful mechanism for offline ad slot scheduling, 2007.

[6]For instance, consider the case of 3 bidders with valuations $(100, 20, 10.5)$, minimum position prices of $(30, 19, 10)$, and position multipliers of $(1, 1/2, 1/4)$. Any matching where everyone pays at most their valuation and at least their reserve price puts the third bidder in the third position, and therefore prices this position at at most 10.5; however then the second bidder will prefer this position to either of the other two.

[12] A. Goel G. Aggarwal and R. Motwani. Truthful auctions for pricing search keywords. In *ACM Conference on Electronic Commerce*, 2006.

[13] J. Feldman G. Aggarwal and S. Muthukrishnan. Bidding to the top: Vcg and equilibria of position-based auctions. In Proc. WAOA, 2006.

[14] T. Groves. Incentives in teams. Econometrica, 41(4):617–631, 1973.

[15] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.

[16] S. Lahie, D. Pennock, A. Saberi, and R. Vohra. *Sponsored Search Auctions, in: Algorithmic Game Theory*, pages 699–716. Cambirdge University Press, 2007.

[17] L. Liang and Q. Qi. Cooperative or vindictive: Bidding strategies in sponsored search auction. In *WINE*, pages 167–178, 2007.

[18] A. Likhodedov and T. Sandholm. Methods for boosting revenue in combinatorial auctions. In *National Conference on Artificial Intelligence (AAAI)*, 2004.

[19] M. Mahdian, H. Nazerzadeh, and A. Saberi. Allocating online advertisement space with unreliable estimates. In *ACM Conference on Electronic Commerce*, pages 288–294, 2007.

[20] M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 243–249, New York, NY, USA, 2006. ACM.

[21] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. In *FOCS*, 2005.

[22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[23] P. Rusmevichientong and D. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proc. 7th ACM conference on Electronic commerce*, 2006.

[24] M. Pál S. Muthukrishnan and Z. Svitkina. Stochastic models for budget optimization in search-based advertising. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.

[25] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.

[26] W. Vickrey. Counterspeculation, auctions and competitive-sealed tenders. *Finance*, 16(1):8–37, 1961.

[27] Y. Vorobeychik and D. M. Reeves. Equilibrium analysis of dynamic bidding in sponsored search auctions. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.

[28] J. Wortman, Y. Vorobeychik, L. Li, and J. Langford. Maintaining equilibria during exploration in sponsored search auctions. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.