

Auctions with Revenue Guarantees for Sponsored Search

Zoe Abrams
Yahoo!, Inc.
2821 Mission College Blvd.
Santa Clara, CA , USA
za@yahoo-inc.com

Arpita Ghosh
Yahoo! Research
2821 Mission College Blvd.
Santa Clara, CA , USA
arpita@yahoo-inc.com

ABSTRACT

We consider the problem of designing auctions with worst case revenue guarantees for sponsored search. This problem differs from previous work because of ad dependent clickthrough rates which lead to *two* natural posted-price benchmarks. In one benchmark, the winning advertisers are charged the same price per click, and in the other, the product of the price per click and the advertiser clickability (which can be thought of as the probability an advertisement is clicked if it has been seen) is the same for all winning advertisers. We adapt the random sampling auction from [10] to the sponsored search setting and improve the analysis from [1], to show a high competitive ratio for two truthful auctions, each with respect to one of the two described benchmarks.

However, the two posted price benchmarks (and therefore the revenue guarantees from the corresponding random sampling auctions) can each be larger than the other; further, which is the larger cannot be determined without knowing the private values of the advertisers. We design a new auction, that incorporates these two random sampling auctions, with the following property: the auction has a Nash equilibrium; and *every* equilibrium has revenue at least the larger of the revenues raised by running each of the two auctions individually (assuming bidders bid truthfully when doing so is a utility maximizing strategy). Finally, we perform simulations which indicate that the revenue from our auction outperforms that from the VCG auction in less competitive markets.

1. INTRODUCTION

1.1 Problem and motivation

We address the problem of designing auctions with revenue guarantees in the sponsored search setting. This problem is crucial for search engines, most of which are publicly held companies that rely heavily on sponsored search for revenue and revenue considerations in sponsored search auc-

tions have been receiving recent attention [22, 15, 7]. The efficient auction in this setting, namely VCG [24, 5, 11, 6, 18], and the current mechanism used for sponsored search, namely GSP [6, 23] (of which the VCG outcome is an equilibrium [6]), do not provide revenue guarantees. In fact, the revenue from VCG can be arbitrarily bad compared to the optimal revenue with full knowledge of bidder valuations, as the footnote shows¹. Our simulations indicate that the revenue from the VCG outcome can be particularly small in less competitive yet realistic markets, where there are not many bidders with similar values for the keyword.

Can we ensure revenue guarantees even in markets with less competition? The metric we use to gauge the performance of an auction is to compare its revenue with the revenue that could be raised by an *optimal omniscient posted price auction* (the auction that raises the optimal revenue if the auctioneer knows the true valuations of all the bidders and charges every bidder the same price). The *competitive ratio* is the worst case ratio, over all possible inputs, between the revenue of an optimal omniscient posted price auction and the revenue of the proposed auction.

The main challenge in applying competitive analysis to our setting is to incorporate the existence of multiple posted price benchmarks arising from the structure of clickthrough rates. As in [6, 2], we model clickthrough rates as *separable*, *i.e.*, the probability that a particular ad in a particular slot is clicked can be broken down into two probabilities. First is the slot-clickability, which can be thought of as the probability that the user will look at the displayed advertisement (the higher the slot placement, the more likely it is that the user will see the ad). The second is the ad-clickability, which can be thought of as the bidder dependent probability that the advertisement will be clicked, given that it has been seen. We describe our model formally in §2. While the basic component of our mechanism is a random sampling auction as in [10], there are now *two* natural posted-price benchmarks depending on whether or not advertisers are discounted for their clickabilities: the optimal omniscient single-price revenue where all winning advertisers must be charged the same price per click, and the optimal omniscient 'weighted-price'

¹This is shown by the following (extreme) example. Suppose there are k slots and $k + 1$ bidders, where k of these bidders have value per click 1, and one bidder has a value per click of ϵ . Each of these ads has a clickthrough rate of c in every slot. Under VCG, the k bidders with the highest values are assigned slots, and every bidder pays his negative externality, which in this case is ϵc . So the revenue extracted by VCG is ϵck , which can be arbitrarily small compared to the optimal revenue which is ck .

revenue, where all winning advertisers are charged the same price per sighting, *i.e.*, the product of price per click and ad-clickability is the same for all winning advertisers. These two benchmarks are the natural posted-price analogs of two charging schemes that have been used in practice: charge an advertiser the bid-per-click of the bidder below him, or discount the bidder proportional to his ad-clickability, *i.e.*, divide the bid per click of the advertiser below by the ad-clickability of the bidder being charged.

Although our work is motivated by sponsored search, it is also applicable in other settings where the probability of a successful event is the product of the probability of two separate events: one event based on factors that depend on the allocation, and the other event based on bidder dependent factors. For example, an airport manager may want to auction off a set of vendor sites. One event is that a potential customer walks past a site, and this event depends on the particular location of the site within the airport (*i.e.* the allocation). The other event is that the potential customer walking past the site will actually enter the site, and this event depends on factors related to the bidder occupying the site such as attractiveness of the site and brand familiarity.

1.2 Results

The main contribution of our paper is a mechanism with a Nash equilibrium that raises revenue that is competitive against both the single price and weighted price benchmarks. To do this, we first adapt the random sampling auction to obtain two auctions, each with high revenue guarantees against one of the two benchmarks. The contribution here is improving an existing analysis of the random sampling auction by a factor 2, in addition to an analysis of the random sampling approach in the sponsored setting, accounting for advertiser and slot clickabilities.

The two random sampling auctions are then used as building blocks for a single auction with Nash equilibria that raises revenue at least as large as that raised by each of the two random sampling auctions independently; further, if bidders bid their true value whenever that belongs to the set of utility maximizing strategies, *every* Nash equilibrium of the auction raises this high revenue. This is significant for the following reason. As we show in Section 3, either of these two benchmarks can be larger than the other: the optimal weighted price revenue can be as small as a factor $O(\log k)$ of the optimal single price revenue, and the optimal single price revenue can be as small as a factor $1/k$ of the optimal weighted price revenue (in fact, the cases where these bounds are tight are when the larger revenue is actually the optimal multiprice revenue). In addition, which benchmark is larger cannot be determined *without knowing the private values of the bidders*. One significant challenge in designing an auction competitive against both benchmarks is that the same bidders participate in both auctions and can have higher utility in one or the other. Despite the presence of such bidders, we find that there is a way to *always* raise revenue competitive with both benchmarks. While the auction is no longer truthful, every Nash equilibrium of the auction has this competitive property.

We also perform numerical simulations to compare the performance of our auction against that of the VCG auction. In crowded markets with a large amount of competition, both auctions achieve a large fraction of the optimal

revenue, and the VCG auction obtains more revenue than the competitive auction. However, as the market becomes less competitive and both auctions achieve a smaller fraction of the optimal revenue, the competitive auction overtakes the VCG auction. Our findings that the competitive auction produces more revenue than the VCG alternative in more challenging situations (*i.e.*, less competitive markets) is in keeping with our analytical framework, as competitive auctions are designed to perform well in worst case settings.

1.3 Related work

Incentive compatible auctions for allocation and pricing in the keyword search setting have been considered previously. In [6], the authors show that any equilibrium using generalized second pricing (*i.e.*, where an advertiser is charged the next highest bid), has revenue at least that of the VCG [24, 5, 11] auction. Another approach [2], gives a truthful pricing mechanism when the allocation of slots is externally specified.

There has also been previous work on auctions that maximize revenue. The classical work of Myerson [20, 18] on optimal auction design shows how to design an auction that maximizes the expected revenue of the seller when the bidder values are drawn from a (known) continuous distribution. The expectation of revenue is over this known distribution. In contrast, we are interested in maximizing revenue in the worst case scenario, *i.e.*, for every possible vector of bid values.

There has been recent work concerned with revenue in the context of sponsored search. Roughgarden and Sundararajan show that in the classical Bayesian setting studied by Myerson [20], the VCG mechanism can be made to obtain high revenue by adding enough bidders, *i.e.*, in more competitive markets. Edelman and Schwarz [7] explore setting reserve prices to increase revenues. The work in [15] studies how to use ad-clickabilities in ranking and pricing to improve revenue. All of these differ from our work in that they do not provide worst case revenue guarantees over all realizations of bidder values.

In terms of competitive analysis for auctions, the random sampling approach was first proposed in [10], and has since been used in several problems and contexts, see for example [16, 4, 12]. Finally, there are several papers that combine multiple auctions into a single auction [3, 19, 1]. In [3], the generalized auction uses two successive auctions to create an auction that is truthful while maintaining the competitive ratio. Unfortunately, this composition does not apply in our context.

2. MODEL

Our model is the following. There are n bidders competing for k slots. Each bidder has a private valuation for a click, v_i . We order bidders by value, *i.e.*, $v_1 \geq \dots \geq v_n$. Every slot-bidder pair has a clickthrough rate c_{ij} associated with it, which is the probability with which the advertisement of bidder i in slot j is clicked. We assume that this clickthrough rate is *separable*, *i.e.*,

$$c_{ij} = \mu_i \theta_j,$$

where we refer to μ_i as the ad-clickability of bidder i , and θ_j as the slot-clickability of slot j . The separability assumption is equivalent to saying that the events of clicking on a particular ad (regardless of which slot it is displayed in) and

a particular slot (regardless of which ad is displayed in it) are independent. Although this assumption is not always entirely accurate, analysis shows it is often reasonable [25], and it has been widely adopted in the literature [2, 6, 17, 21, 14]. We assume that the ad-clickabilities μ_i and slot-clickabilities θ_i are public knowledge. For our results in §4, we only need μ_i and θ_i to be known to the seller (in fact, this is true for all auctions where truthfulness is a dominant strategy), which is a realistic assumption.

We assume that the clickabilities of the slots decrease with position, *i.e.*, $\theta_1 \geq \theta_2 \geq \dots \geq \theta_k$. We define

$$\Theta_i = \sum_{j=1}^i \theta_j, \quad (1)$$

i.e., Θ_i is the sum of the clickabilities of the top i slots.

We denote by b_i the bid of bidder i , and the price charged to bidder i in an allocation by p_i . The auction mechanism takes the bids b_i , and computes an allocation x and pricing p , where $x_i = j$ if the bidder is assigned to slot j , and is 0 if bidder i is not assigned a slot, and p_i is the price that bidder i pays per click he receives in his slot.

For a bidder i , we define

$$w_i = v_i \mu_i,$$

which is the expected value to the bidder from a slot with clickability $\theta_j = 1$. By the separability assumption, the expected value to bidder i in a slot with clickability θ_j is $w_i \theta_j$.

3. OPTIMUM PRICING SOLUTIONS

The previous work on digital goods auctions uses as a benchmark the optimal multi-price and optimal single price revenues [12, 10, 1]. In this section, we extend these concepts to our problem, introducing a new benchmark, optimal weighted price revenue, and bound these benchmarks against each other. While current auctions do not sell clicks in different slots at the same price, a single price (or single weighted price) per click is still meaningful when interpreted as a common reserve price for all slots (note that an advertiser's net payment still depends on the clickthrough rate in the assigned slot even if the price-per-click is the same in all slots).

DEFINITION 1. *Multi-price optimal (OPT_{MP}):* The multi-price optimal revenue, OPT_{MP} , is the maximum possible revenue that can be extracted with k slots, when the true values of all bidders are known. Let $w_{i(j)}$ denote the j th largest value in w , then

$$OPT_{MP} = \sum_{j=1}^{\min(n,k)} w_{i(j)} \theta_j. \quad (2)$$

We denote by O_M the set of bidders that are assigned slots in this allocation.

DEFINITION 2. *Single price optimal (OPT_{SP}):* The single price optimal revenue OPT_{SP} is the maximum revenue that can be extracted with k slots, when the true values of all bidders are known, and every bidder assigned to a slot must be charged the same price per click. Here $p \leq k$ items are sold at a single price v_p , where the single price is chosen to

maximize revenue. Let $\mu_{i(j)}^p$ be the j th largest μ_i of bidders with values $v_i \geq v_p$. Then, OPT_{SP} is computed as

$$OPT_{SP} = \max_{p=1, \dots, \min(n,k)} v_p \sum_{j=1}^p \mu_{i(j)}^p \theta_j. \quad (3)$$

We denote the set of bidders contributing positive revenue to OPT_{SP} as O_S .

Unlike in settings without ad-clickabilities, the optimal single price here is not necessarily limited to one of the values v_1, \dots, v_k – the optimal single price can be any of the values v_1, \dots, v_n . (If $v_i \geq v_j$ implies $\mu_i \geq \mu_j$, however, v_p is clearly greater equal v_k).

DEFINITION 3. *Weighted price optimal (OPT_{WP}):* The weighted price optimal revenue OPT_{WP} is the maximum revenue that can be extracted with k slots, when the true values of all bidders are known, and every bidder assigned to a slot is charged a price inversely proportional to his clickability, *i.e.*, such that $p_i \mu_i$ is constant. OPT_{WP} is computed as follows: sort the w in decreasing order, and choose an index $r \leq k$ that maximizes the revenue when every bidder with $w_i \geq w_r$ contributes w_r to the revenue, *i.e.*,

$$OPT_{WP} = w_r \Theta_r = \max_{j=1, \dots, \min(k,n)} w_{i(j)} \Theta_j. \quad (4)$$

Every bidder who is allocated a slot pays a price

$$p_i = \frac{w_r}{\mu_i} \leq \frac{w_i}{\mu_i} = v_i.$$

We denote the set of bidders contributing positive revenue to OPT_{WP} as O_W .

Note that when all ad-clickabilities μ_i are equal, the weighted price and single price revenues are exactly the same.

We will sometimes use $OPT_{WP}(S)$ and $OPT_{SP}(S)$ to denote the optimal weighted price and single price revenues for a set of bidders S .

The OPT_{WP} benchmark, that weights prices proportional to ad-clickabilities, is attractive for several reasons. It seems natural to give a discount to bidders that bring the auction most value; this is the prominent framework in both theory (VCG, GSP, and the ladder auction) and in practice (Google and soon Yahoo! charge bidders proportional to ad-clickabilities). In addition, Theorems 4 and 5 show that OPT_{MP} is at most H_k times as large as OPT_{WP} , as opposed to k times as large as OPT_{SP} . We also point out that when $|O_S| = |O_W|$, then the competitive ratio against OPT_{SP} is worse than the competitive ratio against OPT_{WP} .

But a further examination of Theorem 2 indicates that, in fact, weighted prices are not clearly superior to charging a single price. As we would anticipate, in practice it is often the case that value and ad-clickability are correlated, since the ultimate goal is to match the searcher with a relevant advertisement. We can think of the ad-clickability and the value as both being increasing functions of the quality of the searcher-advertisement match. Since in this case we always have $OPT_{SP} \geq OPT_{WP}$, it is quite common for single prices to provide better revenue than weighted prices.

3.1 Relating OPT_{SP} and OPT_{WP}

Either OPT_{SP} or OPT_{WP} can be larger, depending on the values of (v, μ) and θ , as the following example shows. Suppose $\theta_i = 1$ for all slots, and bidders clickabilities are

$\mu_1 = 12, \mu_2 = 6, \mu_3 = 4, \mu_4 = 3$. If the bidders valuations are $v = (1, 1, 1, 1)$, then $OPT_{SP} = 25$, and $OPT_{WP} = 12$. However if the values are $v = (1/12, 1/6, 1/4, 1/3)$, then $OPT_{SP} = 13/6$ which is less than $OPT_{WP} = 4$. Notice that which of OPT_{SP} and OPT_{WP} has larger revenue cannot be determined without knowing the true valuations of the bidders.

We now show some theoretical results about how OPT_{SP} and OPT_{WP} are related.

THEOREM 1. *The optimal single price and weighted price revenue are related as follows:*

$$\frac{1}{k} OPT_{WP} \leq OPT_{SP} \leq H_k OPT_{WP}.$$

PROOF. The first inequality is easy:

$$OPT_{SP} \geq v_r \mu_r \theta_1 \geq \frac{1}{k} OPT_{WP},$$

where r is the index chosen by OPT_{WP} as before. The same example that shows that OPT_{MP} can be as large as k times OPT_{SP} also shows that this inequality is tight, since $OPT_{WP} = OPT_{MP}$ for that example.

To show the second inequality, consider the set of bidders in O_S each of whom pays the optimal single price v_p . Consider a modified set of bidders \tilde{O}_S obtained by changing the values of bidders in O_S to $\tilde{v}_i = v_p$. The value of OPT_{SP} for this set of bidders is unchanged. Now consider the optimal weighted price revenue that we can obtain from \tilde{O}_W , which is certainly less than or equal to OPT_{WP} : first, since $\tilde{v}_i \leq v_i$, \tilde{w}_i is less equal w_i , so the optimal weighted price revenue for the bidders in \tilde{O}_W is less equal that for the bidders in O_S . Next, we are considering a subset of the set of all bidders used to compute OPT_{WP} , so the revenue cannot increase.

Let \tilde{r} be the number of bidders in the optimal weighted price solution for this modified subset of bidders, and let $\widetilde{OPT}_{WP} = w_{\tilde{r}} \theta_{\tilde{r}}$ denote this revenue. Then, for all bidders in \tilde{O}_P ,

$$\Theta_{\tilde{r}} w_{\tilde{r}} \geq \Theta_i \tilde{w}_i,$$

that is,

$$\mu_i \leq \frac{\Theta_{\tilde{r}}}{\Theta_i} \mu_{\tilde{r}}.$$

So

$$\begin{aligned} OPT_{SP} &= v_p \sum_{i(j) \in O_S} \mu_{i(j)} \theta_j \\ &\leq v_p \sum_{i(j) \in O_S} \mu_{\tilde{r}} \frac{\Theta_{\tilde{r}}}{\Theta_j} \theta_j \\ &\leq H_p \widetilde{OPT}_{WP} \\ &\leq H_k OPT_{WP}, \end{aligned}$$

where the third line uses the same argument as in Theorem 5, and H_k is the harmonic sum as before.

This bound is tight, as shown by the following example. Let there be $n = k$ bidders each with value $v_i = 1$, clickability $\mu_i = k\mu/i$ for $1 \leq i \leq k-1$, and $\mu_k = \mu + \epsilon$. Then $OPT_{MP} \approx k\mu \log k = OPT_{SP}$. However, $OPT_{WP} = \max i w_{j(i)} = k(\mu + \epsilon)$. \square

This theorem showed that either of OPT_{SP} or OPT_{WP} can be larger than the other; OPT_{SP} can be smaller by a

factor k , while it can only be larger by a factor $O(\log k)$ than OPT_{WP} . While it might still be tempting to choose the weighted price revenue as our benchmark, the next result shows that in the important case when bidders' clickabilities decrease with their values, OPT_{SP} is always greater equal OPT_{WP} .

THEOREM 2. *Suppose clickabilities decrease with values, i.e., $v_i \geq v_j$ implies $\mu_i \geq \mu_j$. Then, the optimal single price revenue is greater equal the optimal weighted price revenue.*

PROOF. Let p be the optimal index in the single price auction; then, since the μ are decreasing with v (i.e., $\mu_{i(j)} = \mu_j$), the revenue is

$$OPT_{SP} = v_p \left(\sum_{i=1}^p \mu_i \theta_i \right).$$

In the weighted price auction, we order the bidders by $w_i = v_i \mu_i$, which, by assumption, is the same as the ordering of the v 's. Let r be the optimal index picked by the weighted price auction. Then, we extract revenue $w_r \theta_j = v_r \mu_r \theta_j$ from the bidder assigned to slot j . So the revenue is

$$\begin{aligned} OPT_{WP} &= v_r \mu_r \left(\sum_{j=1}^r \theta_j \right) \\ &\leq v_r \left(\sum_{j=1}^r \mu_j \theta_j \right) \\ &\leq OPT_{SP}, \end{aligned}$$

where the first inequality follows since $\mu_i \geq \mu_r$ for $i \leq r$, and the second follows from the definition of single price optimum. \square

Note that it cannot be argued that if $v_1 \geq \dots \geq v_n$ and $\mu_1 \leq \dots \leq \mu_n$, then OPT_{WP} is always larger than OPT_{SP} , since the ordering of bidders according to w and the ordering according to v can be unrelated.

Finally we show that OPT_{SP} and OPT_{WP} are close to each other when the clickabilities of winning bidders are not very different.

THEOREM 3. *Let μ_{max} and μ_{min} be the largest and smallest clickabilities of bidders in $O_S \cup O_W$. Then*

$$\frac{\mu_{min}}{\mu_{max}} OPT_{WP} \leq OPT_{SP} \leq \frac{\mu_{max}}{\mu_{min}} OPT_{WP}.$$

PROOF. To show the first inequality, consider the set of bidders in O_W , i.e., the bidders who contribute positive revenue to

$$OPT_{WP} = w_r \sum_{i=1}^r \theta_i.$$

The smallest value of bidders in O_W is at least $\frac{w_r}{\mu_{max}}$. Therefore, by definition of OPT_{SP} ,

$$\begin{aligned} OPT_{SP} &\geq \frac{w_r}{\mu_{max}} \sum_{i \in O_W} \mu_{i(j)} \theta_i \\ &\geq \frac{w_r}{\mu_{max}} \sum_{i \in O_W} \mu_{min} \theta_i \\ &= \frac{\mu_{min}}{\mu_{max}} \left(\sum_{i \in O_W} w_r \theta_i \right), \\ &\geq \frac{\mu_{min}}{\mu_{max}} OPT_{WP}. \end{aligned}$$

Next we show the second inequality.

$$\begin{aligned} OPT_{SP} = v_p \sum_{i \in O_S} \mu_{i(j)} \theta_j &\leq \frac{\mu_{\max}}{\mu_{\min}} (v_p \mu_{\min} \sum_{i=1}^p \theta_i) \\ &\leq \frac{\mu_{\max}}{\mu_{\min}} OPT_{WP}, \end{aligned}$$

where the last inequality uses the fact that for every bidder in O_S ,

$$w_i = v_i \mu_i \geq v_p \mu_{\min},$$

since by definition, v_p is the smallest value of bidders in O_S , and μ_{\min} is less equal the smallest clickability of these bidders. \square

3.2 Bounding Against OPT_{MP}

We now relate OPT_{WP} and OPT_{SP} to OPT_{MP} . Note that while the worst case bounds for both benchmarks are large, the results in Theorem 3 and 4 show that when the top k bidders values for slots is not very widely different, these benchmarks are quite close to OPT_{MP} .

THEOREM 4. $OPT_{MP} \leq k OPT_{SP}$, and this bound is tight.

PROOF. From (2) and (3),

$$OPT_{SP} \geq w_{i(1)} \theta_1 \geq \frac{1}{k} \sum_{j=1}^k w_{i(j)} \theta_j,$$

since the θ_j s are decreasing. To show that this bound is tight, consider the following example. Suppose there are $n = k$ bidders, with $v_j = 1/c^{j-1}$, and $\mu_j = 1/v_j = c^{j-1}$, where c is a large positive constant. All slots have equal clickability $\theta_j = 1$. Then $OPT_{MP} = k$.

For any choice v_i of single price, the revenue is

$$OPT_{SP} = \max_i \frac{1}{c^{i-1}} \sum_{j=1}^i c^{j-1} = \frac{c^i - 1}{c^{i-1}(c-1)},$$

which approaches 1 for large c . \square

However, when clickthrough rates are bidder independent (*i.e.*, $\mu_i = 1$), the optimal single-price revenue can be no smaller than a factor $O(\log k)$ of the optimal multi-price revenue. This follows directly from the next result since in this case $OPT_{WP} = OPT_{SP}$.

THEOREM 5. $OPT_{MP} \leq H_k OPT_{WP}$, where $H_k = 1 + \frac{1}{2} + \dots + \frac{1}{k}$. This bound is tight.

PROOF. Let $r = |O_W|$ be the number of slots sold by OPT_{WP} . From (4) and (1), for $j = 1, \dots, k$,

$$w_{i(j)} \leq w_r \frac{\Theta_r}{\Theta_j}.$$

So the optimal multi-price revenue, OPT_{MP} , is

$$\begin{aligned} \sum_{j=1}^k w_{i(j)} \theta_j &\leq \sum_{j=1}^k \theta_j w_r \frac{\Theta_r}{\Theta_j} \\ &= w_r \Theta_r \sum_{j=1}^k \frac{\theta_j}{\Theta_j} \\ &\leq OPT_{WP} \sum_{j=1}^k \frac{1}{j}, \end{aligned}$$

where the last inequality follows from the fact that $j\theta_j \leq \sum_{i=1}^j \theta_i$, since the θ s are decreasing. When all θ and all μ are equal to 1 and $v_i = \frac{1}{i}$, all inequalities are tight, so this bound is tight as well. \square

While these theorems show that OPT_{SP} and OPT_{WP} can be quite small compared to the multiprice optimal, when bidders' valuations are more consistent, OPT_{SP} and OPT_{WP} are quite close to OPT_{MP} , as shown in the following theorems.

THEOREM 6. Let v_{\max} be the largest, and v_{\min} be the smallest value of the bidders contributing to OPT_{MP} . Then $OPT_{MP} \leq (v_{\max}/v_{\min}) OPT_{SP}$.

PROOF. We have, with $w_{i(j)} = v_{i(j)} \mu_{i(j)}$,

$$\begin{aligned} OPT_{MP} &= \sum_{j=1}^k w_{i(j)} \theta_j \\ &= \frac{1}{v_{\min}} \sum_{j=1}^k v_{\min} v_{i(j)} \mu_{i(j)} \theta_j \\ &\leq \frac{v_{\max}}{v_{\min}} \left(v_{\min} \sum_{j=1}^k \mu_{i(j)} \theta_j \right) \\ &\leq \frac{v_{\max}}{v_{\min}} OPT_{SP}, \end{aligned}$$

where the last inequality follows from the definition of OPT_{SP} , since every bidder in O_M has value greater equal v_{\min} . \square

Note here that v_{\max} and v_{\min} are values from OPT_{MP} , and need not be the largest and smallest values from the entire set of bidders (*i.e.*, not necessarily v_1 and v_n).

A nearly identical argument can be used to show

THEOREM 7. Let w_{\max} be the largest, and w_{\min} be the smallest revenues of the bidders contributing to OPT_{MP} . Then $OPT_{MP} \leq (w_{\max}/w_{\min}) OPT_{WP}$.

4. AUCTIONS COMPETITIVE AGAINST A SINGLE BENCHMARK

In this section we describe two truthful auctions that are competitive against the optimal single price and weighted price revenues. The auctions in this section are based on the random sampling auction from [10]. However, extending previous analyses gives us a competitive ratio that approaches 2 against the optimum weighted price revenue and 4 against the optimum single price revenue. First we improve upon the analysis in [1] by a factor 2, to obtain a competitive ratio that also approaches 2 against the optimum single price revenue. Next, we incorporate decreasing slot-clickabilities into our analysis to further improve our guarantees to approach near optimal, as the steepness in clickthrough rates increases.

The two competitive auctions use versions of *ProfitExtract* from [10] that are described in the Appendix. Given a set of bidders S and a revenue R , $ProfitExtract_{WP}^R$ is an incentive compatible auction that extracts revenue R using weighted pricing, if $OPT_{WP}(S) \geq R$. Given a set of bidders S and a revenue R , $ProfitExtract_{SP}^R$ is an incentive compatible auction that extracts revenue R using *single* pricing, when possible. Unlike *ProfitExtract_{WP}*, this auction assigns higher slots to bidders whose ads have higher clickabilities.

4.1 Mechanism competitive with OPT_{WP}

Now we give an auction mechanism M_{WP} which has high competitive ratio (less equal 4 and asymptotically optimal as a function of bidder dominance and slot clickabilities) with respect to OPT_{WP} .

Mechanism M_{WP}

1. Partition bidders independently and uniformly at random into two subsets S_1 and S_2 .
 2. Compute $R_1 = OPT_{WP}(S_1) - \epsilon$, and $R_2 = OPT_{WP}(S_2) + \epsilon$.
 3. Run $ProfitExtract_{WP}^{R_1}$ on the bidders in S_2 , and $ProfitExtract_{WP}^{R_2}$ with the bidders in S_1 .
-

We assume that revenues are calculated to some finite precision, and we choose $\epsilon > 0$ to be small compared with this precision.

A straightforward application of the analysis from [10] provides at most a guarantee of two, because the revenue extracted is the lesser of the random division of contributions to the optimum. Our setting has a unique structure which allows us to improve upon this guarantee: clickthrough rates are decreasing with respect to rank. The performance of M_{WP} depends on the bidder dominance with respect to participants (*i.e.*, the inverse of the number of participants), and the drop-off rate of the slot-clickabilities. We show that the revenue from M_{WP} is at least a factor 1/4 of OPT_{WP} , and approaches optimal as the bidder dominance decreases and the drop-off in slot-clickabilities becomes steep:

THEOREM 8. M_{WP} is truthful, and has competitive ratio

$$\beta_{WP} = \frac{\bar{\theta}_r}{g(\alpha_{WP})\bar{\theta}_{\lfloor r/2 \rfloor}}$$

with respect to OPT_{WP}^2 (the optimal weighted price auction selling at least two items), where $g(\alpha_{WP}) \geq 1/4$, and $g(\alpha_{WP}) \rightarrow 1/2$ as $\alpha_{WP} \rightarrow 0$.

Here $\bar{\theta}_m = \frac{\Theta_m}{m}$ is the average clickability for the top m slots. (Since the θ s are decreasing, $\bar{\theta}_m$ decreases as m increases, *i.e.*, as we average over more slots.) The bidder dominance, α_{WP} , is defined as

$$\alpha_{WP} = \frac{1}{r},$$

where $r = |O_W|$ is the number of slots sold in OPT_{WP} . The function $g(x)$ is defined in (5) from the Appendix, and lies between 1/4 and 1/2 for $x \leq 1/2$.

The value of β_{WP} is roughly the product of two values: one value starts at 2 and tends to 1 as the number of bidders in the optimum solution increases, the other value is the sum of all slot clickabilities, divided by the sum of the largest half of the slot clickabilities, and always lies between 1 and 2.

4.2 Mechanism competitive with OPT_{SP}

Next we describe and analyze a mechanism which is competitive with respect to OPT_{SP} . An application of previous results [1, 10] gives an auction that approaches a competitive ratio of 4 as the bidder dominance decreases. We give a new proof that tightens previous analysis and allows us to achieve a competitive ratio of 2 (this also improves on

the results in [1]). We define bidder dominance in the context of single price, to be the largest advertiser clickability in the optimum solution divided by the sum of advertiser clickabilities in the optimum solution. Then, we provide an analysis showing that as the CTRs become more steep, and the bidder dominance approaches 0, the competitive ratio approaches 1.

Recall that O_S is the set of bidders contributing positive revenue to OPT_{SP} , $p = |O_S|$ and the optimal single price is v_p .

Define the average clickability of bidders in O_S as

$$\bar{\mu} = \frac{\sum_{i \in O_S} \mu_i}{p},$$

and the bidder dominance

$$\alpha_{SP} = \frac{\mu_{\max}}{\sum_{i \in O_S} \mu_i},$$

where μ_{\max} is the largest clickability of bidders in O_S . The smallest value of α_{SP} with p bidders in the optimal single price solution is $1/p$, when all bidders have the same clickability. (Note that this bidder dominance depends both on bidders' values (which are implicitly present in α_{SP} through p), and the clickabilities of the bidders in O_S .)

Define a second bidder dominance parameter

$$\alpha'_{SP} = \frac{\theta_1 \mu_{\max}}{\sum_{i \in O_S} \theta_j \mu_{i(j)}}.$$

Observe that since the θ are decreasing, $\alpha_{SP} \leq \alpha'_{SP}$, with equality when all the θ_i are equal.

We prove that the mechanism below achieves near optimal revenue as $\alpha_{SP} \rightarrow 0$, and the slot clickabilities decrease steeply enough. The competitive ratio also shows that the revenue is always greater than $\frac{1}{4}$ when at least two items are sold.

Mechanism M_{SP}

1. Partition bidders independently and uniformly at random into two subsets S_1 and S_2 .
 2. Compute $R_1 = OPT_{SP}(S_1) - \epsilon$ and $R_2 = OPT_{SP}(S_2) + \epsilon$.
 3. Run $ProfitExtract_{SP}^{R_1}$ on the bidders in S_2 , and $ProfitExtract_{SP}^{R_2}$ with the bidders in S_1 .
-

We prove the following theorem about this mechanism (the proof is included in the Appendix):

THEOREM 9. M_{SP} is truthful, and has competitive ratio

$$\beta_{SP} = \max \left(\frac{p \bar{\theta}_p \alpha_{SP}}{g(\alpha_{SP}) \bar{\theta}_{p - \frac{1}{2\alpha_{SP}}}}, \frac{1}{g(\alpha'_{SP})} \right),$$

against OPT_{SP} when $\alpha_{SP} \leq 1/2$, where $\frac{1}{2} \leq \frac{1}{2\alpha_{SP}} \leq \frac{p}{2}$, and $g(x)$ is defined in (5) from the Appendix.

The first term in the max, roughly in words, is the product of three values. The first is the largest ad clickability divided by the average ad clickability. The second is the average slot clickability, divided by the average slot clickability of a portion of the largest slot clickabilities (at least the half largest). Finally, the last value is at least 1/4 and approaches 1/2 as

the bidder dominance decreases (here, bidder dominance is measured by ad-clickabilities and is assumed to be at most $1/2$).

To understand why decreasing clickabilities is advantageous, consider a weighted price solution with two bidders. Each is capable of contributing the same amount to the optimum solution. We could place them in arbitrary positions and still obtain the same optimal revenue. However, in the optimum solution the one placed in the highest position contributes more. Now suppose they have been divided into two bins, (a.k.a. the first step of the random sampling auction). Each bidder can now potentially contribute as much as the highest contributor to revenue, even though its true contribution in the optimum is actually much less. This is the intuition behind our improved analysis.

5. AN AUCTION COMPETITIVE AGAINST MULTIPLE BENCHMARKS

In this section, we describe a mechanism with high revenue guarantees against both the single price and weighted price benchmarks. To do this, we use the two random-sampling auctions from §4 that have high competitive ratio against OPT_{SP} and OPT_{WP} respectively. We combine these two auctions to derive a single auction with a Nash equilibrium that raises revenue at least that raised by each of the individual random-sampling auctions.

As we saw in §3, for a particular set of values and clickabilities (v_i, μ_i) , either the optimum weighted price revenue OPT_{WP} or optimum single price revenue OPT_{SP} could be larger. However, which of the two is actually larger cannot be determined without knowing the true values of the bidders.

Here, we describe a new mechanism that builds on the two auctions in §4 to raise a larger amount of revenue. Of course, we can combine the two auctions using randomization into a single truthful auction that raises expected revenue $\frac{1}{2}(OPT_{SP}/\beta_{SP} + OPT_{WP}/\beta_{WP})$. To achieve a revenue that is the better of the two auctions, we break from truthful mechanism design and instead design an auction with equilibria (which we show always exist) such that the revenue raised is at least the larger of the revenues that would be raised by the auctions M_{WP} and M_{SP} . The resulting equilibrium analysis framework for the random sampling approach is more robust and malleable. Our hope is that this additional flexibility will have implications for other contexts and applications as well.

We point out that the revelation principle does not apply in our setting; also, bidding truthfully is not necessarily an equilibrium strategy (we will say later for which bidders it is not an equilibrium strategy).

Mechanism M_C

1. Partition the bidders randomly into two sets A and B , announce the partition, and collect bids from all bidders.
2. Compute $R^A = \max(OPT_{SP}^A, OPT_{WP}^A)$, and $R^B = \max(OPT_{SP}^B, OPT_{WP}^B)$ using the reported bids.
3. Run $ProfitExtract_{SP}^{R^B}$ on the bidders in A ; if the auction fails to raise revenue R^B , run $ProfitExtract_{WP}^{R^B}$. Do the same for the bidders in B .

4. If $R^A = R^B$, then items are only assigned to bidders in partition A .

In what follows, we will use R^{A^*} to denote the value of R^A when every bidder bids his true value (similarly for R^B , OPT_{SP}^A , OPT_{WP}^A , OPT_{SP}^B , and OPT_{WP}^B).

We show the following result for the combined auction for every instance of the random partition of bidders:

THEOREM 10. *There always exists an equilibrium solution with revenue at least*

$$R = \min(\max(OPT_{SP}^{A^*}, OPT_{WP}^{A^*}), \max(OPT_{SP}^{B^*}, OPT_{WP}^{B^*})).$$

Further, if bidders bid their true value whenever bidding truthfully belongs to the set of utility maximizing strategies, every Nash equilibrium of M_C has this property.

PROOF. Assume wlog that $R^{A^*} \geq R^{B^*}$. First we will show existence. We consider the following cases:

- Case I: $R^{B^*} > \min(OPT_{SP}^{A^*}, OPT_{WP}^{A^*})$, *i.e.*, only one of the two auctions can raise the revenue R^{B^*} from bidders in A . Then $b_i = v_i$ is a Nash equilibrium for the combined auction: every bidder who does not win an item has no incentive to deviate from $b_i = v_i$, since his utility is 0 for all $b_i \leq v_i$, and can only be non-positive if he reports a bid $b_i > v_i$. Every bidder who wins an item has no incentive to deviate: if he reports $b_i \leq v_i$, his utility cannot increase, since he either fails to win an item, or wins an item but still pays a price independent of his bid. This Nash equilibrium raises revenue R^{B^*} , since every bidder in B reports his true value.
- Case II: $R^{B^*} \leq \min(OPT_{SP}^{A^*}, OPT_{WP}^{A^*})$, *i.e.*, the revenue R^{B^*} can be extracted using both single price and weighted price mechanisms from bidders in A . We will show that there is a Nash equilibrium in which B is the losing partition, and bids are as specified below.

First, note that for all bidders (in both partitions) who do not win an item in either solution, there is no incentive to deviate from $b_i = v_i$, using the same reasoning as above. Since the bidders in B lose, the mechanism tries to extract revenue R^{B^*} from the bidders in A .

For the same reason, every bidder who can win an item in only one of OPT_{SP} or OPT_{WP} has no incentive to deviate from $b_i = v_i$. This leaves us with bidders who might win an item in both OPT_{SP} and OPT_{WP} . We consider two sub-cases for bidders with such values, based on the following condition:

Condition C: There is no bidder with higher utility in OPT_{WP} who can unilaterally decrease his bid enough to ensure that $ProfitExtract_{SP}$ fails to extract R^{B^*} , while still winning an item in OPT_{WP} .

- Condition C holds: In this case, $b_i = v_i$ is an equilibrium vector of bids. A bidder winning an item in both OPT_{WP} and OPT_{SP} has no incentive to bid $b_i > v_i$; if he reports $b_i < v_i$, he might fail to win an item in OPT_{SP} , which still extracts revenue R^{B^*} by assumption.

- Condition C does not hold (*i.e.*, there is at least one bidder with higher utility in OPT_{WP} who can unilaterally decrease his bid enough to ensure that $ProfitExtract_{SP}$ fails to extract R^{B^*} while still winning an item in OPT_{WP} .)

Let w^* be the single weighted price at which $ProfitExtract_{WP}$ extracts revenue R^{B^*} from the bidders in A . Let i be a bidder satisfying the condition above. Then the vector of bids with $b_i = w^*/\mu_i$ for any one bidder satisfying this condition, and $b_i = v_i$ for all other bidders is a Nash equilibrium: there is no incentive for i to change his bid because b_i is the lowest bid at which i still can win an item in OPT_{WP} ; by assumption this bid is low enough to ensure that $ProfitExtract_{SP}$ fails to raise R^{B^*} . Further, bidder i cannot increase his utility by deviating from this value, nor can any other bidder improve its utility by deviation. Note that bidder i can be any single bidder that causes the condition to be violated.

Therefore, in either subcase, there is a Nash equilibrium in which B is the losing partition, and that extracts the specified revenue.

We now prove the second part of the theorem. If bidders bid their true value whenever bidding truthfully belongs to the set of utility maximizing strategies, bidders in the losing partition always bid their true value. Therefore, the only Nash equilibria are those where B is the losing partition, in which case a revenue of R^{B^*} is extracted. So *every* Nash equilibrium of M_C extracts the specified revenue. \square

For a particular partition of the bidders into A and B , the revenue extracted by M_{SP} is

$$R_{SP} = \min(OPT_{SP}^{A^*}, OPT_{SP}^{B^*}),$$

and the revenue extracted by M_{WP} is

$$R_{WP} = \min(OPT_{WP}^{A^*}, OPT_{WP}^{B^*}).$$

From Theorem 10, the revenue extracted by the auction M_C is $\min(\max(OPT_{SP}^{A^*}, OPT_{WP}^{A^*}), \max(OPT_{SP}^{B^*}, OPT_{WP}^{B^*}))$, which is greater equal $\max(R_{WP}, R_{SP})$. Taking the expectation over random partitions, we see that the expected revenue from M_C is $\max(\beta_p OPT_{SP}, \beta_r OPT_{WP})$. (Note that M_C is actually stronger, since we obtain the larger revenue of M_{WP} and M_{SP} for *every* partition, not just in expectation over partitions.)

A natural question is whether such a situation, where a bidder prefers one mechanism to the other, and can indeed make the other mechanism fail to raise revenue can indeed exist. Also, can the natural opposite mechanism not be tried then- is it possible to have a situation where there are both bidders who prefer WP and who prefer SP , and both can make the other mechanism fail?

The following example will show that this can indeed happen, and help understand why the mechanism is not truthful. We give an example with two bidders where, given the revenue to be raised by the weighted price and single price auctions, each bidder prefers a different auction, and further, is able to make the other mechanism fail, *i.e.*, unable to raise the required revenue.

Suppose $v_1 = 1, v_2 = 2, \mu_1 = 2, \mu_2 = 1$, and $\theta_i = 1$. For these values, $OPT_{WP} = 4, OPT_{SP} = 3$, so both auctions can raise a revenue of 3. To raise this revenue using the weighted price mechanism, bidder 1's price is $1.5/2 = 3/4$, and his utility is $2(1 - 3/4) = 1/2$. Bidder 2 is charged a price of 1.5, and his utility is $1(2 - 1.5) = 1/2$. When the single price mechanism is used, both bidders are charged a price of 1, and their utilities are $2 * (1 - 1) = 0$, and $1 * (2 - 1) = 1$ respectively. Thus bidder 1 has higher utility when the revenue of 3 is raised using weighted price, and bidder 2 has higher utility under single price. Further, if bidder 1 bid $3/4$ instead of her true utility of 1, the single price mechanism can no longer raise the revenue of 3, while weighted price still can; similarly, if bidder 2 bids 1 instead of her true utility of 2, the weighted price mechanism can no longer raise a revenue of 3, while single price still raises the required revenue. Thus each bidder has a bid different from her true value that can improve her utility given that the other bidder bids her true value.

6. SIMULATION RESULTS

In this section we discuss our simulation results. We draw bidder valuations from a lognormal distribution with increasing variance and unit mean. This distribution has been used previously [8] and also fits the distribution observed in practice. For our simulations, we used $n = 50$ bidders, $k = 12$ slots, and ad-clickabilities μ_i proportional to v_i . Each point plotted in a figure is obtained by averaging over 800 draws of bidder valuations from a lognormal distribution of the corresponding variance and unit mean. We use two sets of vectors for the slot clickabilities θ . We call slot clickabilities with $\theta_i = 0.7^i$ *Geometric Slot-clickabilities*. This distribution for slot clickabilities is in keeping with [9]. When several advertisements are shown at the top of the page and others shown along the right hand side, the slot clickabilities tend to be significantly larger for advertisements shown along the top. To model this situation, we use a set of *Sharp Geometric Slot-clickabilities*, where the first four slots (presumably shown along the top), decrease by a factor of .85, starting from .85, and the remaining slots along the east, starting from .4, decrease by a factor of .4. We also point out that because ad-clickabilities have the same ordering as the bid values, due to Theorem 2, the revenue of a Nash equilibrium using Mechanism M_C equals the revenue extracted using Mechanism M_{SP} .




Figure 1: Geometric Slot-clickabilities: Revenue versus Variance of Bidder Valuations Drawn from a Log-normal Distribution

The general shape of Figures 2 and 1 follow a similar pattern. For $\sigma = 0$, there is no variance in the bids and both




Figure 2: Sharp Geometric Slot-clickabilities: Revenue versus Variance Including OPT_{mp}

algorithms achieve the revenue of the optimal multi-price solution. Initially, the variance of the bids is small, and the VCG auction outperforms the combined auction. As the variance in the bid values begin to diverge more sharply, the combined mechanism outperforms VCG.

VCG revenue decreases dramatically because as the bid values become more varied and every individual's bid value more distinctive, the externalities a bidder imposes on others decreases (because externalities measure, to some degree, how 'replaceable' a bidder is). We can also consider highly varied bid values as a less competitive market. If a single bidder's value lies far away from others, it does not have to fight other contenders off for his position: it is clear who the winners should be and there is not much competition for the clicks.

It is often difficult to design incentive compatible auctions for markets with little competition. Truthful auctions rely on bids other than b_i to set values for bidder i . When there is a lot of variance in the bids, choosing a reasonable price is more challenging. This can be seen by observing Figure 3. The multiprice optimum shoots up, relative to both algorithms, as the bidder variance increases. This suggests that both algorithms have difficulty obtaining revenue in these situations. The simulations corroborate the findings in Theorem 6, which prove analytically that the tighter the range of bidder values, the higher the performance guarantee.

Since the combined mechanism is designed to do well in a worst case setting, it is not surprising that its performance improves relative to VCG exactly when maintaining a minimal amount of revenue in the face of a challenging situation (*i.e.* non-competitive market) is encountered.

Figures 1 and 2 highlight how the steepness of slot-clickabilities impacts the algorithms' revenues. There is very little difference in the curve for the VCG mechanism when the slot clickabilities are steeper. However, the improvement for the combined mechanism is more noticeable, outperforming VCG earlier and by a larger margin. This is consistent with our analysis, which indicates that the auction will perform better as the steepness in slot clickabilities increases.

Our simulations use the algorithms described in §4 and §5, but the auctioneer could alternatively implement a variation of the combined auction where the partition splits into two sets of equal size, chosen uniformly at random. In practice, this algorithm maintains an equilibrium (and truthfulness where appropriate). Although more cumbersome to analyze, it is a more appropriate algorithm in to use in practice and leads to a slight increase in performance.




Figure 3: Revenue versus Variance Including OPT_{MP}

7. FUTURE WORK

There are a number of interesting questions that remain open. First, is whether it is possible to design truthful auctions that achieve better guarantees (*i.e.*, better competitive ratios, an impossibility result, or the larger revenue of the M_{SP} and M_{WP}). Another question is whether it is possible to perform competitive analysis using random sampling optimal price, along the lines of [10, 4].

One possible direction for future work is to compare the performance of these auctions against other benchmarks. Perhaps we can theoretically bound the revenue in our auction against VCG revenue, or against the best VCG revenue obtained by artificially limiting the supply as in [13]. Another possible benchmark would be to compare against the optimal revenue auction from Myerson[20] given noisy information about bidder valuations.

A considerable obstacle in achieving good bounds for keyword search problems is that the performance relies on having a large scale problem where no individual bidder has too much influence on the optimum solution. If there are many auctions with similar properties, it is possible that they could be used either to merge markets together so that the competitive ratio approaches optimal more quickly, or to use advertisers and bidders for one set of keywords to determine solutions for other sets of keywords.

Finally, it would be interesting to set reserve prices using the auctions presented here.

Acknowledgements: We are very grateful to Andrei Broder, Ravi Kumar, Ofer Mendelevitch and Michael Schwarz for helpful discussions.

8. REFERENCES

- [1] Z. Abrams. Revenue maximization when bidders have budgets. In *Proc. Symposium on Discrete Algorithms*, pages 1074–1082, 2006.
- [2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proc. 7th ACM conference on Electronic Commerce*, pages 1–7, 2006.
- [3] G. Aggarwal and J. Hartline. Knapsack auctions. *Symposium on Discrete Algorithms*, 2006.
- [4] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proc. 6th ACM*

Conference on Electronic Commerce, pages 44–51, 2005.

- [5] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17-33, 1971.
- [6] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalizaed second price auction: Selling billions of dollars worth of keywords. *Second Workshop on Sponsored Search Auctions, Ann Arbor, MI. June, 2006*.
- [7] B. Edelman and M. Schwarz. Optimal auction design in a multi-unit environment: The case of sponsored search auctions. *ACM Conference on Electronic Commerce*, 2007.
- [8] B. Edelman and M. Schwarz. Optimal auction design in a multi-unit environment: The case of sponsored search auctions. *Manuscript in preparation*, 2007.
- [9] J. Feng, H. Bhargava, and D. M. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 2006.
- [10] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. *Proc. Symposium on Discrete Algorithms*, 2001.
- [11] T. Groves. Incentives in teams. *Econometrica*, 41:617-631, 1973.
- [12] J. Hartline. Optimization in the private value model: Competitive analysis applied to auction design. *PhD Thesis*, 2003.
- [13] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. *Proc. of 6th ACM conference on Electronic Commerce*, 2005.
- [14] N. Immorlica, K. Jain, M. Mahdian, and K. Talwar. Click fraud resistant methods for learning click-through rates. *Workshop on Internet and Network Economics*, 2005.
- [15] S. Lahaie and D. M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. *ACM Conference on Electronic Commerce*, 2007.
- [16] M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. *ACM Conference on Electronic Commerce*, 2006.
- [17] S. Sathiya Keerthi and J. A. Tomlin. Constructing an optimal slate of advertisements. *Manuscript in preparation*, 2006.
- [18] P. Milgrom. *Putting Auction Theory to Work*. something, 9999.
- [19] A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *AAAI*, 2002.
- [20] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [21] S. Pandey and C. Olston. Handling advertisements of unknown quality in search advertising. *Neural Information Processing Systems Conference*, 2006.
- [22] T. Roughgarden and M. Sundararajan. Is efficiency expensive? *Third Workshop on Sponsored Search Auctions*, 2007.
- [23] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2005.
- [24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*,

16:8-37, 1961.

- [25] T. Zhang. Clickthrough-rates positional effect and calibration model. *Manuscript in preparation*, 2006.

APPENDIX

First we give the two extensions of *ProfitExtract* used in M_{WP} and M_{SP} .

*ProfitExtract*_{WP}^R - A Weighted Price Auction

Set $K = \min(k, |S|)$. While $K > 0$

1. Set $w = \frac{R}{\sum_{i=1}^K \theta_i}$.
2. If there are at least K bidders with bid $b_i \geq w/\mu_i$, assign slot j to the bidder with clickability $\mu_{i(j)}$ for $j = 1, \dots, k$, and return this allocation and w .
3. Set $K = K - 1$.

If $K = 0$, all bidders lose.

LEMMA 1. *ProfitExtract*_{WP}^R is truthful, and extracts revenue R if $R \leq OPT_{WP}(S)$, and 0 otherwise.

We note that arbitrarily allocating winning bidders to slots can also be used for the same results; we use this for consistency with the *ProfitExtract*_{SP}.

*ProfitExtract*_{SP}^R - A Single Price Auction

Set $\bar{S} = S$, the set of all bidders. While $\bar{S} \neq \emptyset$

1. Set $K = \min(k, |\bar{S}|)$.
2. Set $p = \frac{R}{\sum_{j=1}^K \mu_{i(j)} \theta_j}$, where $\mu_{i(j)}$ denotes the j th largest clickability of bidders in \bar{S} .
3. If each of the k bidders contributing to the denominator has bid $b_{i(j)} \geq p$, assign slot j to the bidder with clickability $\mu_{i(j)}$ and return this allocation and a single price of p .
4. Remove all bidders with $b_i < p$ from \bar{S} .

If $\bar{S} = \emptyset$, all bidders lose.

LEMMA 2. *ProfitExtract*_{SP}^R is truthful, and extracts revenue R if $R \leq OPT_{SP}(S)$, and 0 otherwise.

A. COMPETITIVE RATIO OF M_{WP}

THEOREM 11. M_{WP} is truthful, and has competitive ratio

$$\beta_{WP} = \frac{\bar{\theta}_r}{g(\alpha_{WP})\bar{\theta}_{\lfloor r/2 \rfloor}}$$

with respect to OPT_{WP}^2 , where $g(\alpha_{WP}) \geq 1/4$, and $g(\alpha_{WP}) \rightarrow 1/2$ as $\alpha_{WP} \rightarrow 0$.

PROOF. The revenues R_1 (resp. R_2) to be extracted and the number of slots k are independent of the bids of bidders in S_2 (resp. S_1). Since *ProfitExtract*_{WP} with independent parameters is truthful, M_{WP} is truthful in this case. The addition and subtraction of ϵ ensures $R_1 \neq R_2$. Combined with Lemma 1, the revenue from this auction is $R_{WP} =$

$\min(R_1, R_2)$, exactly one side of the partition wins, and we do not oversell advertisement slots. Since ϵ is chosen to be very small compared to the precision of the revenue, we ignore it in the analysis that follows.

Observe that R_1 is greater equal the optimal weighted price revenue from bidders in $S_1 \cap O_W$. So we need only consider partitioning the bidders in O_W to bound the revenue. If $|S_1 \cap O_W| = i$, and $|S_2 \cap O_W| = r - i$, then $R_1 \geq w_r i \bar{\theta}_i$, and $R_2 \geq w_r (r - i) \bar{\theta}_{r-i}$, where $w_r = v_r \mu_r$ is the contribution of each bidder in OPT_{WP} . So we have

$$\begin{aligned} \frac{E[R_{WP}]}{OPT_{WP}^2} &\geq \frac{1}{r\bar{\theta}_r} \sum_{i=1}^{r-1} \min(\Theta_i, \Theta_{r-i}) \binom{r}{i} 2^{-r} \\ &\geq \frac{\bar{\theta}_{\lfloor r/2 \rfloor}}{r\bar{\theta}_r} \sum_{i=1}^{\lfloor \frac{r}{2} \rfloor} i \binom{r}{i} 2^{-r} \\ &\geq \frac{\bar{\theta}_{\lfloor r/2 \rfloor}}{\bar{\theta}_r} \left(\frac{1}{2} - \binom{r-1}{\lfloor \frac{r}{2} \rfloor} 2^{-r} \right), \end{aligned}$$

where the second line follows since $i\bar{\theta}_{\lfloor r/2 \rfloor} \leq i\bar{\theta}_i = \Theta_i$ for all $i \leq \lfloor r/2 \rfloor$. Define, for $x \leq 1/2$,

$$g(x) = x \lfloor \frac{1}{x} \rfloor \left(\frac{1}{2} - \binom{\lfloor \frac{1}{x} \rfloor - 1}{\lfloor \frac{1}{2x} \rfloor} 2^{-\lfloor \frac{1}{x} \rfloor} \right). \quad (5)$$

Thus, the competitive ratio is $\frac{\bar{\theta}_r}{g(\alpha_{WP})\bar{\theta}_{\lfloor r/2 \rfloor}}$ as stated. \square

B. COMPETITIVE RATIO OF M_{SP}

THEOREM 12. *M_{SP} is truthful, and has competitive ratio*

$$\beta_{SP} = \max \left(\frac{p\bar{\theta}_p \alpha_{SP}}{g(\alpha_{SP})\bar{\theta}_{p-\frac{1}{2\alpha_{SP}}}}, \frac{1}{g(\alpha'_{SP})} \right),$$

against OPT_{SP} when $\alpha_{SP} \leq 1/2$, where $\frac{1}{2} \leq \frac{1}{2\alpha_{SP}} \leq \frac{p}{2}$, and $g(x)$ is as in (5).

PROOF. Following the same reasoning as in the proof of Theorem 11, M_{SP} is truthful, and the revenue extracted is $\min(R_1, R_2)$ (and exactly one side of the partition wins). Again, we ignore ϵ in the analysis since it is negligibly small.

Let $\mathbf{r} = (r_1, \dots, r_n)$, where r_i is the revenue contributed by bidder i to OPT_{SP} . Observe that $R_1 \geq \sum_{i \in S_1 \cap O_S} r_i$: every bidder in $S_1 \cap O_S$ has value greater equal v_p , and is assigned to a slot with θ_j greater equal his assignment in OPT_{SP} (the same argument holds for R_2). So it is enough to consider bidders in O_S , and bound

$$R(r) = E[\min(\sum_{i \in S_1 \cap O_S} r_i, \sum_{i \in S_2 \cap O_S} r_i)]$$

against $\sum_{i \in O_S} r_i$.

To do this we will apply Lemma 3 to a vector \mathbf{r}' with $m = \lfloor 1/\alpha'_{SP} \rfloor$ non-zero entries of value $r_{\max} = \theta_1 \mu_{\max}$ each, where \mathbf{r}' is obtained by repeatedly applying REDISTRIBUTE (see Lemma 3) to \mathbf{r} . From the Lemma, bounding

$$R(\mathbf{r}') = E[\min(\sum_{i \in S_1 \cap O_S} r'_i, \sum_{i \in S_2 \cap O_S} r'_i)]$$

gives us a bound on revenue. But this is easy since each of the non-zero entries in \mathbf{r}' have the same value r_{\max} :

$$R(\mathbf{r}') = r_{\max} \sum_{i=1}^{m-1} \min(i, m-i) \binom{m}{i} 2^{-m} = m r_{\max} g(\alpha'_{SP}).$$

Therefore,

$$\frac{R(\mathbf{r})}{OPT_{SP}} \geq g(\alpha'_{SP}). \quad (6)$$

However, this analysis does not account for the fact while computing the optimum single price revenues on each side, the winning bidders are associated with clickthrough rates greater equal those in OPT_{SP} . Next we obtain another bound accounting for this; the final competitive ratio is the better of the two bounds.

For any partition of the bidders, assume wlog that the sum of clickabilities of bidders from O_S is smaller in the partition S_1 , and let

$$\sum_{i \in S_1 \cap O_S} \mu_i = \delta \left(\sum_{i \in O_S} \mu_i \right) = \delta p \bar{\mu},$$

where $0 \leq \delta \leq 1/2$. Let $X = |O_S \cap S_1|$. The optimal single price revenue from this subset of the bidders is

$$\begin{aligned} R_1 &\geq v_p \sum_{i \in O_S \cap S_1} \mu_{(j)} \theta_j \\ &\geq v_p \frac{\delta p \bar{\mu}}{X} \sum_{j=1}^X \theta_j \\ &= v_p \delta p \bar{\mu} \bar{\theta}_X, \end{aligned}$$

Recall that the mechanism assigns bidders with highest clickabilities to the top slots. Similarly,

$$R_2 \geq v_p (1 - \delta) p \bar{\mu} \bar{\theta}_{p-X}.$$

So the smaller of the two revenues is bounded by

$$\begin{aligned} \min(R_1, R_2) &\geq (p v_p \bar{\mu}) \min(\delta, 1 - \delta) \min(\bar{\theta}_X, \bar{\theta}_{p-X}) \\ &= p v_p \bar{\mu} \delta \bar{\theta}_{\max(X, p-X)}, \end{aligned}$$

since we assumed $0 \leq \delta \leq 1/2$ and $\bar{\theta}_m$ decreases with increasing m since the θ s are decreasing.

Define $\gamma = \frac{\mu_{\max}}{\bar{\mu}} = \alpha p$. We upper bound X as follows: the number of bidders in the partition with the larger fraction of $\bar{\mu} p$ must be at least

$$\begin{aligned} (p - X) &\geq \frac{(1 - \delta) p \bar{\mu}}{\mu_{\max}}, \\ \Rightarrow X &\leq \frac{(\gamma - 1 + \delta) p}{\gamma} \leq \frac{p(\gamma - \frac{1}{2})}{\gamma}, \end{aligned}$$

since $\delta \leq 1/2$. Since $\gamma \geq 1$, $(\gamma - 1/2)/\gamma \geq 1/2$, and so

$$\max(X, p - X) \leq \frac{p(\gamma - \frac{1}{2})}{\gamma}$$

as well.

So for a particular partition with ratio δ , we have

$$\min(R_1, R_2) \geq \delta p v_p \bar{\mu} \bar{\theta}_{\frac{p(\gamma - \frac{1}{2})}{\gamma}}, \quad (7)$$

where now the only term that depends on the random partition is δ .

The single price optimal revenue is bounded as

$$OPT_{SP} \leq v_p \mu_{\max} p \bar{\theta}_p = \gamma p v_p \bar{\mu} \bar{\theta}_p.$$

So the expected revenue from this mechanism is

$$\frac{\min(R_1, R_2)}{OPT_{SP}} \geq \frac{E[\delta] \bar{\theta}^{\frac{p(\gamma-\frac{1}{2})}{\gamma}}}{\gamma \theta_p}, \quad (8)$$

where

$$E[\delta] = \left(\sum_{i \in O_S} \mu_i \right) E[\min(\sum_{i \in S_1 \cap O_S} \mu_i, \sum_{i \in S_2 \cap O_S} \mu_i)].$$

We bound $E[\delta]$ using Lemma 3 as we did above, to obtain

$$\frac{\min(R_1, R_2)}{OPT_{SP}} \geq g(\alpha_p) \frac{\bar{\theta}^{\frac{p(\gamma-\frac{1}{2})}{\gamma}}}{\gamma \theta_p}. \quad (9)$$

Combining the two results in (6) and (9), and using $\gamma = \alpha p$, we have the theorem. \square

Now we state and prove Lemma 3. Let $b = (b_1, \dots, b_n)$ be a vector of nonnegative numbers. For i, j with $b_i \geq b_j$, and any Δ with $0 \leq \Delta \leq b_j$, define $b' = \text{REDISTRIBUTE}(b, i, j, \Delta)$ to be the vector with $b'_i = b_i + \Delta$, $b'_j = b_j - \Delta$, and $b'_m = b_m$ for $m \neq i, j$. Define $R(b) = E(\min(\sum_{i \in S_1} b_i, \sum_{i \in S_2} b_i))$, where each b_i is independently thrown into S_1 or S_2 with probability $1/2$ (i.e., $R(b)$ is the expected value over random partitions of the sum of entries in the smaller partition).

LEMMA 3. *For any nonnegative vector b , $R(b) \geq R(b')$, where $b' = \text{REDISTRIBUTE}(b, i, j, \Delta)$.*

PROOF. Let $S_0 = \{1, \dots, n\}$. Consider the set S_{\min} of all subsets with the lesser sum for the given vector b , i.e., $S_{\min} = \{S \subset S_0 \mid \sum_{j \in S} b_j \leq \sum_{j \in S_0 - S} b_j\}$. Given i and j , the indices of the bids in the REDISTRIBUTE operation, partition the sets in S_{\min} into four sets as $S_{b_i b_j} = \{S \in S_{\min} \mid b_i, b_j \in S\}$, $S_{\bar{b}_i \bar{b}_j} = \{S \in S_{\min} \mid b_i, b_j \notin S\}$, $S_{b_i \bar{b}_j} = \{S \in S_{\min} \mid b_i \in S, b_j \notin S\}$; $S_{\bar{b}_i b_j} = \{S \in S_{\min} \mid b_i \notin S, b_j \in S\}$.

Let p_S denote the probability of a particular set $S \in S_{\min}$ being the subset in the random partition with the smaller value (note that choosing S is the same as choosing the partition of the bids b_i). Let us write $|S|_b = \sum_{i \in S} b_i$, and $|b| = \sum_{i=1}^n b_i$. Then,

$$\begin{aligned} R(b) &= \sum_{S \in S_{\bar{b}_i \bar{b}_j}} p_S |S|_b + \sum_{S \in S_{b_i b_j}} p_S |S|_b \\ &\quad + \sum_{S \in S_{b_i \bar{b}_j}} p_S |S|_b + \sum_{S \in S_{\bar{b}_i b_j}} p_S |S|_b \end{aligned} \quad (10)$$

and

$$\begin{aligned} R(b') &= \sum_{S \in S_{\bar{b}_i b_j}} p_S (|S|_b - \Delta) \\ &\quad + \sum_{S \in S_{b_i \bar{b}_j}, |S|_b + \Delta \leq \frac{|b|}{2}} p_S (|S|_b + \Delta) \\ &\quad + \sum_{S \in S_{b_i \bar{b}_j}, |S|_b + \Delta > \frac{|b|}{2}} p_S (|b| - |S|_b - \Delta) \\ &\quad + \sum_{S \in S_{b_i b_j}} p_S |S|_b + \sum_{S \in S_{\bar{b}_i \bar{b}_j}} p_S |S|_b. \end{aligned} \quad (11)$$

Note that for sets S with $|S|_b + \Delta > |b|/2$, $|S|_b - (|b| - |S|_b - \Delta) = 2|S|_b - |b| + \Delta > -\Delta$. Subtracting (11) from

(10) and using this, we see that

$$\begin{aligned} R(b) - R(b') &> \sum_{S \in S_{\bar{b}_i b_j}} \Delta p_S + \sum_{S \in b_i \bar{b}_j, |S|_b + \Delta \leq |b|/2} p_S (-\Delta) \\ &\quad + \sum_{S \in S_{b_i \bar{b}_j}, |S|_b + \Delta > |b|/2} p_S (-\Delta) \\ &= \Delta \left(\sum_{S \in S_{\bar{b}_i b_j}} p_S - \sum_{S \in S_{b_i \bar{b}_j}} p_S \right). \end{aligned}$$

But this difference is clearly positive: since $b_j \leq b_i$, for every set $S \in S_{b_i \bar{b}_j}$, there is a set $S' \in S_{\bar{b}_i b_j}$ obtained by swapping b_i with b_j ; also $p_{S'} = p_S$. So $\sum_{S \in S_{\bar{b}_i b_j}} p_S > \sum_{S \in S_{b_i \bar{b}_j}} p_S$, and the lemma is proved. \square