

Truthful Auctions for Pricing Search Keywords ^{*}

Gagan Aggarwal [†]
Google Inc.

Ashish Goel [‡]
Stanford University

Rajeev Motwani [§]
Stanford University

March 12, 2006

Abstract

We present a truthful auction for pricing advertising slots on a web-page assuming that advertisements for different merchants must be ranked in decreasing order of their (weighted) bids. This captures both the *Overture* model where bidders are ranked in order of the submitted bids, and the *Google* model where bidders are ranked in order of the expected revenue (or utility) that their advertisement generates. Assuming separable click-through rates, we prove revenue-equivalence between our auction and the non-truthful next-price auctions currently in use.

1 Introduction

Keyword auctions are an indispensable part of the business model of modern web search engines and are responsible for a significant share of their revenue. In a keyword auction, a set of merchants submit bids on specific keywords. If a user searches for a keyword, advertisements from the merchants for that keyword are shown to the user along with results that match the keyword. The set of advertisements as well as their order is determined by the bids submitted for that keyword. The search engines (which we will refer to as auctioneers) typically charge a merchant only when a user actually clicks on the advertisement. The price charged depends on the bids submitted. It is clear that merchants will have a preference for paying a smaller price and for obtaining a higher rank in the ordering of advertisements. These conflicting objectives, along with the fact that click-through rates vary depending on position and that the merchants are selfish agents trying to maximize their own utility, makes this an interesting and involved auction. The same framework applies to the

^{*}To appear in ACM conference on electronic commerce, 2006

[†]Gagan Aggarwal was a PhD student at the Computer Science Department of Stanford University when this research was conducted. Email: gagan@cs.stanford.edu.

[‡]Departments of Management Science & Engineering and (by courtesy) Computer Science, Stanford University. Research supported by NSF ITR grant 0428868, an NSF CAREER award 0339262, and an Alfred P. Sloan Faculty Fellowship. Email: ashishg@stanford.edu.

[§]Computer Science Department, Stanford University. Email: rajeev@cs.stanford.edu.

pricing of advertising slots on static web content as opposed to dynamically generated web-pages such as search engine results.

In the auctions currently being used, the search engine first picks the subset of advertisements to be displayed and matches them to slots based on the submitted bids; the matching criteria is referred to as the *ranking function* and is an integral component of the existing keyword auctions. Then, the auctioneer decides on a price for each merchant based on the bids and the allocation. There are two popular ranking methods:

1. The “Overture” method: Merchants are ranked in the decreasing order of the submitted bids. We will call this *direct ranking*.
2. The “Google” method: Merchants are ranked in the decreasing order of the *ranking scores*, where the ranking score of a merchant is defined as the product of the merchant’s bid and estimated click-through rate (or relevance [4]). We refer to this as *revenue ranking*.

These ranking functions are an inherent part¹ of the advertisement methodologies of Overture and Google. Accordingly, we will assume that these ranking functions are fixed. Thus, we impose *rank-equivalence* as a constraint on any auction mechanism that we design, i.e., the truthful mechanism must result in the same ranking as the original mechanism in place. Hence the only degree of freedom in the auction is the price charged per click-through for each merchant. Both Overture and Google currently charge a merchant the minimum amount she would need to bid to retain her current position in the auction². This price can never be larger than the submitted bid, since clearly, the submitted bid was enough to guarantee the merchant her current position. The utility of a merchant is the expected gain per impression (i.e. the expected gain each time her advertisement appears) and is defined in section 2. We will refer to this auction as the *next-price* auction. Despite superficial similarity to the second-price auction [8], the next-price auction is not truthful – in Section 3.1, we present examples where a merchant has an incentive to bid less than her true value under the above auctions

This motivates the following natural question: can we design truthful keyword auctions? In a truthful auction, bidding their true valuations for a keyword is a dominant strategy equilibrium for the merchants. While interesting in its own right, the problem of designing truthful keyword auctions is not merely academic. First, since truth-telling is not a dominant strategy in the current auction, there is no clear prescription for merchants to determine their optimum bid. This optimum bid depends in a complicated dynamic fashion on externalities such as the bids of the other merchants, and it is often necessary for merchants to hire expensive consultants or intermediaries to determine these bids. A truthful mechanism would simplify the bidding process significantly, since it would require a merchant to only determine her valuation for the keyword, a quantity that is intrinsic to the merchant. Secondly, observe that in the current auctions run by Google and Overture, there is an asymmetric incentive for merchants – there may be an incentive for a merchant to bid less than her true value for each click, but there is never an incentive for over-bidding. A truthful

¹One possible reason could be the computational simplicity afforded by this approach.

²Plus a fixed small increment, but we will ignore this minor technical detail.

mechanism would remove this under-bidding incentive³. Finally, in the case of revenue ranking and with an additional separability assumption (defined in section 2), a truthful mechanism is efficient in the sense that it maximizes the total utility obtained by the auctioneer and the merchants together.

One might be tempted to suggest that the famous VCG mechanism [8, 2, 5] or a weighted and biased variant of VCG would yield a solution to this problem. However, we show an example (Section 3.2) where there does not exist any set of weights and biases for which the VCG mechanism always outputs the same merchant ordering as the given ranking function. Hence, the VCG mechanism is not generally applicable to our problem (it is applicable if the click-through rates are separable). We further discuss the applicability of VCG in Section 3.

Our results: We design a simple truthful auction for a general class of ranking functions that includes direct ranking and revenue ranking. More specifically, we study the case where the merchants are assigned arbitrary weights which do not depend on the bids, and then ranked in decreasing order of their weighted bids – we define this formally in section 2. Informally, setting all the weights to 1 results in the direct ranking used by Overture, and setting the weights equal to the estimated click-through rates results in the revenue ranking scheme used by Google.

We call our auction the *laddered auction*, since the price for a merchant builds on the bid of each merchant ranked below it. We show that this auction is truthful. Further, we show that the laddered auction is the unique truthful auction, and hence is trivially revenue-maximal for the auctioneer among all truthful auctions. The auction is presented in section 4 and the analysis is in 5.

We then ask the next natural question: how will the auctioneer’s revenue change as a result of implementing our truthful auction rather than the next-price auction currently in use? Since the next-price auction is not truthful, its revenue should be computed assuming that the bids of the merchants are in a Nash equilibrium. For general weights and click-through rates, we have not been able to answer this question, primarily because we can not obtain a simple characterization of the Nash equilibria imposed by the next-price auction in the general case. However, when the click-through rates are separable (i.e. the click-through rates can be separated into a merchant-specific factor and a position-specific factor; see Section 2 for a formal definition), we prove the following revenue equivalence theorem:

There exists a *pure-strategy* Nash equilibrium for the next-price auction which yields exactly the same revenue for the auctioneer as our laddered auction.

We give an explicit characterization of this Nash equilibrium. These results are presented in section 6. This is arguably the most important special case for the purpose of proving revenue equivalence⁴. Interestingly, we show that there may exist other pure-strategy Nash equilibria under which the next-price auction achieves a smaller revenue than the truthful auction, and yet others

³Admittedly, this is at the expense of decreased prices since our truthful mechanisms require charging the merchant less than the next-price auction given the same bids. However, the revenue equivalence theorem that we prove later under restricted scenarios should alleviate this concern.

⁴Unlike truthfulness, revenue equivalence does not impact the social efficiency or the user experience or the merchant experience. Hence, we believe the correct approach is to guarantee truthfulness in all cases, whereas it is sufficient to prove revenue equivalence for the common case.

under which the next-price auction achieves a higher revenue. These examples are presented at the end of Section 6. In fact, starting from the truthful bids, there may be sequences of self-interested moves (i.e. bid changes) that can lead to a Nash equilibrium for the next-price auction of higher or lower revenue than the truthful auction. This suggests that while the revenue of the current auctions could be better or worse than our ladderred auction depending on which equilibrium the bids settle into, the revenue of the ladderred auction is more predictable.

A subset of our results on revenue equivalence can be obtained by using general techniques such as the envelope theorem [7] as follows: (a) We have proved that VCG is applicable if click-through rates are separable. (b) We have also proved that our ladderred auction is the unique truthful auction, and hence it must coincide with VCG if click-through rates are separable, allowing us to apply general techniques developed for proving revenue equivalence of VCG. However, our first-principles proof of revenue equivalence gives stronger properties, as explained in section 6.

Discussion and Related Work: We assume throughout that the number of slots, K , that the auctioneer is selling for a given keyword does not depend on the submitted bids. Our auction is not truthful if the auctioneer computes the optimum number of slots (in terms of the revenue generated by our ladderred pricing scheme) to be displayed. Extending our auctions to this case appears to be a non-trivial and interesting research direction.

Since the submitted bids are typically used for more than one impression, in addition to a merchant’s valuation, her budget may also be a parameter of relevance [6, 1]. However, before undertaking a combined study of budgets and truthfulness, we need a better understanding and modeling of the role of budgets. For example, one method for a merchant to determine her true value for a keyword may be to just compute the expected immediate profit per click (presumably, merchants want users to go to their web-sites to purchase merchandise which results in immediate profit). If the merchant bids in accordance with the truth-telling strategy that is dominant for our ladderred auction, she would bid an amount equal to her valuation of the keyword. Since the price charged by the auctioneer is never larger than the bid, each click results in an immediate net expected profit. Hence, ignoring budgets may be the right thing to do under this scenario.

Independent of our work, Edelman *et al.* [3] proved the revenue-equivalence of VCG under the assumption that the click-through rate $CTR_{i,j}$ depends only on the position j . This is easily seen to be more restrictive than the separability assumption. Hence, our proof of the revenue equivalence of the ladderred auction (combined with the uniqueness of the ladderred auction as a truthful mechanism) can also be thought of as a generalization of their result.

2 Model and Notation

There are N merchants bidding for $K < N$ slots on a specific keyword. Let $CTR_{i,j}$ denote the click-through rate of the i -th merchant if placed at slot $j \leq K$. We assume that $CTR_{i,j}$ is arbitrary, but known to the auctioneer. Also, we assume that $CTR_{i,j}$ is non-increasing in j . Set $CTR_{i,j} = 0$ for $j > K$. Let v_i denote the true value of a click-through to merchant i . We assume that v_i is known to merchant i , but not to the auctioneer.

As outlined in the introduction, we will assume that the ranking function is externally specified.

We will consider the class of ranking functions where merchant i is assigned an a priori weight w_i that is independent of her bid. Let b_i denote the bid of the i^{th} merchant for each click-through. The merchants are ranked in the order of decreasing $w_i b_i$. Setting $w_i = 1$ for all i is equivalent to the direct ranking function (the Overture model), while setting $w_i = \text{CTR}_{i,1}$ reduces to the revenue-ranking function (the Google model). Merchant i is charged a price-per-click, $p_i \leq b_i$, which is determined by the auction. We assume the merchants to be risk-neutral. As such, if merchant i is placed at position j , she obtains a utility of $\text{CTR}_{i,j} \cdot (v_i - p_i)$ per impression. Recall that an auction is truthful if bidding her true valuation (i.e. $b_i = v_i$) is a dominant strategy for every agent. Now, we formally define the next-price auctions currently being used.

Definition 2.1 (Next-price Auction) *Given a ranking function, $R = (w_1, w_2, \dots, w_n)$ and a bid vector $\mathbf{b} = (b_1, \dots, b_n)$, the next-price auction ranks the merchants in the decreasing order of $w_i b_i$ and charges the merchant ranked i an amount-per-click equal to the minimum bid she needs to have submitted in order to retain rank i . Let w_a and w_b refer to the weights of the merchants ranked i and $i + 1$ respectively. And let b_b refer to the bid submitted by the merchant ranked $i + 1$. Then the price charged to the merchant ranked i is $\frac{w_b b_b}{w_a}$.*

We will now describe the separability assumption, which we will use (only) for our results on revenue-equivalence. Informally, this assumption states that the click-through rates can be separated into a merchant-specific factor and a position-specific factor. There is evidence to believe that this is a reasonable assumption that holds (approximately) in many real-world cases.

Definition 2.2 (Separable Click-through Rates) *The click-through rates are said to be separable if there exist $\mu_1, \mu_2, \dots, \mu_n > 0$ and $\theta_1 \geq \theta_2 \geq \dots, \theta_K > 0$ such that the click-through rate $\text{CTR}_{i,j}$ of the i^{th} merchant at the j^{th} slot is given by $\mu_i \theta_j$.*

3 Need for a New Auction

In this section, we begin by giving an example to show that the next-price auctions being currently used by Google and Overture are not truthful. In order to construct a truthful auction, the first logical step is to see whether the famous VCG mechanism applies to the problem. However, this is not the case, and we give instances of ranking functions for which there does not exist any set of VCG weights and biases (defined later) for which the ranking output by the VCG mechanism agrees with the given ranking function. We also illustrate why it might be more desirable (in some settings at least) to design a truthful mechanism that agrees with the ranking function as opposed to simply ignoring the ranking function and using VCG.

3.1 Next-price Auction is not Truthful

Consider three merchants A , B and C bidding for two slots. Let all three of them have a click-through rate of 0.5 at the top slot and 0.4 at the bottom slot. Let the true valuations per click of the three merchants be 200, 180, and 100 respectively. Then, if all the merchants bid truthfully,

merchant A ends up paying a price of 180 per click, making an expected profit of $(200 - 180) \times 0.5 = 10$ per impression. In this case, she has an incentive to undercut B by lowering her bid to 110, and make a net profit of $(200 - 100) \times 0.4 = 40$. We note that there is no incentive to bid higher than one's true valuation under the next-price auction. This is because the price-per-click charged is the minimum bid required to retain one's rank; therefore, in cases where bidding higher improves one's rank, the price-per-click charged is higher than one's true valuation.

A similar observation was made by Edelman *et al.* [3].

3.2 Weighted VCG may not Always Apply

In this section, we show by means of a counter-example that even for the simple case of direct ranking, there does not exist any set of (bid-independent) weights and biases for which the VCG solution achieves the same allocation as direct ranking. This will show that, in general, VCG does not apply to our problem. Consider two merchants A and B bidding for two slots on a web page. Let both the merchants have a click-through rate (CTR) of 0.4 at the first slot. For the second slot, merchant A has a CTR of 0.4 while merchant B has a CTR of 0.2. Since any of the merchants can bid the highest and get the top slot in direct ranking, both the merchants must have non-zero weight in order for weighted VCG to achieve the same allocation as direct ranking. Let $\omega_A > 0$ and $\omega_B > 0$ be the weights assigned by the VCG mechanism to merchants A and B respectively. Denote the bias assigned to ranking merchant x followed by y by $H(x, y)$ for $x, y \in \{A, B\}$, $x \neq y$. Then, the VCG mechanism will rank B before A if $\omega_A(0.4b_A) + \omega_B(0.4b_B) + H(B, A) > \omega_A(0.4b_A) + \omega_B(0.2b_B) + H(A, B)$, which is true whenever $b_B > (H(A, B) - H(B, A)) / (0.2\omega_B)$, irrespective of merchant A 's bid. On the other hand, the direct ranking scheme will rank A before B whenever A 's bid is higher than B 's bid. Thus, the VCG mechanism does not apply to this instance. In fact, we show the following general theorem.

Theorem 3.1 *Let the number of merchants with non-zero click-through rates be $n > K$. If the click-through rates are not separable, then there exists a ranking function $R = (w_1, w_2, \dots, w_n)$ for which there does not exist any set of weights for which unbiased, weighted VCG always yields the same ranking as the ranking function R .*

The proof of this theorem is in appendix A. Although we have presented this theorem for unbiased VCG, a similar statement holds for biased, weighted VCG as well. We can see this as follows. In each of the constraints 8, 9, 10 and 11 in the proof in appendix A, the two sides of the equation represent the rating given by the VCG mechanism to two different outcomes. If the biased VCG mechanism adds an unequal bias to the two outcomes, it can be viewed as adding a non-zero bid-independent constant term to the right-hand side. The key idea is that we can scale up the bids uniformly without changing the ordering output by the ranking function R . Thus, the chosen VCG weights need to satisfy the constraints both for the scaled and the unscaled bid vector, which is impossible in the presence of a non-zero bid-independent constant term. Thus, the VCG mechanism must have added the same bias to both sides of each of the constraints, thereby leaving the constraints unchanged. We state the following theorem without proof.

Theorem 3.2 *Let the number of merchants with non-zero click-through rates be $n > K$. If the click-through rates are not separable, then there exists a ranking function $R = (w_1, w_2, \dots, w_n)$ for which there does not exist any set of weights for which biased, weighted VCG always yields the same ranking as the R .*

Of course one might argue that VCG is inapplicable only because of the *rank-equivalence* constraint: we want the auction mechanism to produce the same ranking as the original next-price auction. So retaining VCG at the expense of rank-equivalence might superficially seem to be a viable alternative. However, consider the simple two-merchant scenario where the VCG weights are 1 for both merchants (a natural choice). Merchant A has a CTR of 0.8 at both positions 1 and 2, and merchant B has a CTR of 0.15 at position 1 and of 0.1 at position 2. Regardless of what A and B bid, the VCG mechanism (since it maximizes efficiency) will place bidder A in 2nd place. This is certainly against the spirit of keyword auctions and akin to the loss of “voter-sovereignty” in social choice theory. It is unlikely that merchant A would find this ranking acceptable if her bid is higher than that of merchant B.

To summarize, VCG does not appear to be an appropriate choice if the click-through rates are not separable. Interestingly, VCG is applicable under the separability assumption, with appropriately chosen weights. The following theorem is easy to verify:

Theorem 3.3 *Let the click-through rates be separable. Then the VCG mechanism having merchant i 's VCG weight set to $w_i/CTR_{i,1}$ always produces the same ordering as the ranking function (w_1, \dots, w_n) .*

4 The Laddered Auction

In this section, we will assume without loss of generality that the i^{th} merchant also has the i^{th} rank in the auction. The truthful auction is quite simple: For $1 \leq i \leq K$, set the price-per-click p_i charged to merchant i according to the equation:

$$p_i = \sum_{j=i}^K \left(\frac{CTR_{i,j} - CTR_{i,j+1}}{CTR_{i,i}} \right) \frac{w_{j+1} b_{j+1}}{w_i}. \quad (1)$$

Equivalently,

$$CTR_{i,i} p_i = \sum_{j=i}^K (CTR_{i,j} - CTR_{i,j+1}) \frac{w_{j+1} b_{j+1}}{w_i} \quad (2)$$

which lends itself to an intuitive explanation:

1. For those clicks which merchant i would have received at position $i + 1$, she pays the same price as she would have paid at position $i + 1$.
2. For the additional clicks, merchant i pays an amount equal to the minimum bid required to retain position i .

The above intuition is inherently recursive (or laddered), since step 1 requires computation of the price that the merchant would have paid at position $i + 1$. Hence the sum over all j from i to K in the formula for computing p_i naturally corresponds to the above intuitive explanation.

Since $w_i b_i \geq w_j b_j$ for $j > i$, it follows that $p_i \leq b_i$. Hence the price charged per click-through can be no larger than the submitted bid. We will refer to this auction as *Laddered Auction*(w_1, \dots, w_n) or simply as the *laddered auction* when the w_i s are clear from the context.

As an example, consider four merchants A, B, C and D bidding for three slots. The click-through rates are given in the table below.

	A	B	C	D
Slot 1	0.4	0.35	0.3	0.2
Slot 2	0.30	0.25	0.25	0.18
Slot 3	0.18	0.2	0.2	0.15

Let their weights w_i and bids b_i be as shown in the first two columns of the table below. Then, the last three columns show the scores, ranks and prices assigned by the laddered auction.

	Weight	Bid	Score	Rank	Price
Merchant A	60	25	1500	1	13.5
Merchant B	40	30	1200	2	16
Merchant C	50	16	800	3	12
Merchant D	40	15	600	4	0

5 Truthfulness of the Laddered Auction

Theorem 5.1 *Given fixed w_1, \dots, w_n , the laddered auction is truthful. Further, it is the unique truthful auction that ranks according to decreasing $w_i b_i$.*

Proof: Consider a merchant M . Fix the bids of all the other merchants arbitrarily. With these bids, let $p(j)$ be the price charged by the laddered auction to merchant M if her rank is j , with $p(K + 1) = 0$. Note that the price charged depends only on merchant M 's rank and is independent of her exact bid value. Let v_M be the true valuation of a single click for merchant M . If merchant M bids v_M , let her be ranked x . Also, without loss of generality, assume that all the merchants are indexed such that merchant j would be ranked j if merchant M bids v_M . Then, $w_j b_j \geq w_x v_x$ for all $j < x$ and $w_x v_x \geq w_j b_j$ for all $j > x$. To show that the auction is truthful, we will show that merchant M cannot benefit by lying about her valuation. Among all ranks that give the merchant the highest profit per impression (i.e., utility – cost), let r be the rank closest to x , i.e. the one with the least $|r - x|$. Now suppose that the merchant can benefit by lying, i.e., $r \neq x$. For a contradiction, we will show that there is a rank closer to x which gives at least the same profit. For this, observe that if $r > x$, then the change in profit by moving to rank $r - 1$ is $(\text{CTR}_{x,r-1} - \text{CTR}_{x,r})(v_x - \frac{w_r}{w_x} b_r)$, which is non-negative. On the other hand, if $r < x$, the change in profit in moving to rank $r + 1$ is $(\text{CTR}_{x,r+1} - \text{CTR}_{x,r})(v_x - \frac{w_r}{w_x} b_r)$, which is again non-negative.

To show uniqueness, consider any truthful auction \mathcal{A} that ranks the merchants in the decreasing order of $w_i b_i$. Consider any merchant M and fix the bids of all the other merchants arbitrarily. With these bids, let $p_{\mathcal{A}}(j)$ be the price charged by auction \mathcal{A} to merchant M if she is ranked j , with $p_{\mathcal{A}}(K+1) = 0$. Note that in a truthful auction, $p_{\mathcal{A}}(j)$ can depend on the bids of other merchants, but is independent of M 's bid. Assume, without loss of generality, that the other merchants are indexed such that merchant i would be ranked i if merchant M bids ∞ . To prove uniqueness, it suffices to show that for any truthful auction,

$$p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) = (\text{CTR}_{x,j} - \text{CTR}_{x,j+1}) \frac{w_{j+1}}{w_x} b_{j+1} \quad (3)$$

First suppose merchant M has valuation $v_M = \frac{w_{j+1}}{w_x} b_{j+1} + \epsilon$. Then, if she bids truthfully, for sufficiently small $\epsilon > 0$, she is ranked j . The additional valuation per impression of being ranked j instead of $j+1$ is given by $(\text{CTR}_{x,j} - \text{CTR}_{x,j+1})v_x$. Thus, this is the maximum amount that can be charged by a truthful auction for this additional valuation (otherwise, the merchant can benefit by bidding lower to get rank $j+1$). Since ϵ can be made arbitrarily small, this proves that

$$p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) \leq (\text{CTR}_{x,j} - \text{CTR}_{x,j+1}) \frac{w_{j+1}}{w_x} b_{j+1} \quad (4)$$

Next, suppose merchant M has valuation $v_M = \frac{w_{j+1}}{w_x} b_{j+1} - \epsilon$. Then, if she bids truthfully, for sufficiently small $\epsilon > 0$, she is ranked $j+1$. The additional valuation per impression of being ranked j instead of $j+1$ is given by $(\text{CTR}_{x,j} - \text{CTR}_{x,j+1})v_x$. Thus, this is the minimum amount that can be charged by a truthful auction for this additional valuation (otherwise, the merchant can benefit by bidding higher to get the j^{th} rank). Since ϵ can be made arbitrarily small, this proves that

$$p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) \geq (\text{CTR}_{x,j} - \text{CTR}_{x,j+1}) \frac{w_{j+1}}{w_x} b_{j+1}. \quad (5)$$

Putting together 4 and 5, we get 3, thereby completing the proof. \blacksquare

Corollary 5.2 *For any fixed w_1, \dots, w_n , the ladder auction is the profit-maximizing truthful auction that ranks merchants by decreasing $w_i b_i$.*

6 Revenue Equivalence with Separable Click-Through Rates

In this section, we compare the revenue of the ladder auction to the revenue achieved by the next-price auctions currently being used. As mentioned earlier, truth-telling is not a dominant strategy for the existing auctions. Thus, we consider the revenue of the existing auctions under equilibrium conditions, i.e. a setting of bids for which no merchant can increase her profit by a unilateral change in his bid.

For separable click-through rates (see Definition 2.2), we show that there exists a *pure-strategy* Nash equilibrium under the next-price auction that yields the same revenue as the ladder auction. Let the weights used by the next-price auction be (w_1, w_2, \dots, w_n) . Re-index the merchants in the decreasing order of $w_i v_i$ so that

$$w_i v_i \geq w_{i+1} v_{i+1} \quad \text{for } i = 1, \dots, n-1 \quad (6)$$

Let the click-through rates be separable with the click-through rate of merchant i at position j given by $\mu_i \theta_j$. Also let $\theta_{K+1} = 0$. Then, the bids b_i for this Nash equilibrium are recursively defined for $i = K, \dots, 1$ as follows:

$$w_i b_i = \left(\frac{\theta_i}{\theta_{i-1}} \right) w_{i+1} b_{i+1} + \left(1 - \frac{\theta_i}{\theta_{i-1}} \right) w_i v_i \quad (7)$$

with the initialization $b_{K+1} = v_{K+1}$.

Theorem 6.1 (Revenue-Equivalence Theorem) *The bids defined by the recursive formula given in Equation 7 are in equilibrium. Moreover, the ranking induced by these bids is the same as the ranking induced by truthful bidding.*

Proof: To prove this, we unroll the recursion to get:

$$w_i b_i = \frac{1}{\theta_{i-1}} \sum_{j=i-1}^K (\theta_j - \theta_{j+1}) w_{j+1} v_{j+1}$$

Thus, $w_{i+1} b_{i+1}$ is a convex linear combination of $w_j v_j$ for $j = i + 1, \dots, K + 1$. Since, $w_i v_i \geq w_j v_j$ for $j = i + 1, \dots, K + 1$, we get $w_i v_i \geq w_{i+1} b_{i+1}$. We also know that $w_i b_i$ is a convex linear combination of $w_{i+1} b_{i+1}$ and $w_i v_i$. Hence, $w_i b_i \geq w_{i+1} b_{i+1}$. This shows that the ranking induced by these bids is the same as that induced by truth-telling, i.e., merchant i is ranked i by the ranking function of the next-price auction.

Next, we show that under the next-price auction, no merchant can gain by changing her bid unilaterally. Consider the merchant ranked (and indexed) x . With the above bids, she is making a profit of:

$$\begin{aligned} U(x) &= \mu_x \theta_x (v_x - b_{x+1}) \\ &= \mu_x \sum_{j=x}^K (\theta_j - \theta_{j+1}) \left(v_x - \frac{w_{j+1} v_{j+1}}{w_x} \right) \end{aligned}$$

If the merchant changes her bid in order to be ranked y , her profit becomes

$$\begin{aligned} U(x) &= \mu_x \theta_y (v_x - b_{y+1}) \\ &= \mu_x \sum_{j=y}^K (\theta_j - \theta_{j+1}) \left(v_x - \frac{w_{j+1} v_{j+1}}{w_x} \right) \end{aligned}$$

If the merchant decreases her bid in order to be ranked y , i.e. $y > x$, then the net change in profit is:

$$-\mu_x \sum_{j=x}^{y+1} (\theta_j - \theta_{j+1}) \left(v_x - \frac{w_{j+1} v_{j+1}}{w_x} \right)$$

By equation 6, $w_x v_x \geq w_{j+1} v_{j+1}$ for $j = x, \dots, y + 1$, which in turn implies that the above change is non-positive. Similarly, if the merchant increases his bid in order to be ranked y , i.e. $y < x$, the net change in profit is:

$$\mu_x \sum_{j=y}^{x+1} (\theta_j - \theta_{j+1}) \left(v_x - \frac{w_{j+1} v_{j+1}}{w_x} \right)$$

Again, equation 6 implies that this change is non-positive. Thus, none of the merchants has any incentive to change her bid unilaterally and the bids are in equilibrium. ■

Note that the merchants can achieve this equilibrium by solely using the knowledge of their true valuation and the current price being charged to them. To do this, the merchants start by bidding their true valuations, after which the k^{th} -ranked merchant changes her bid to the one indicated in the formula above in order to prevent anybody from under-cutting her, followed by merchant $k - 1$ changing her bid to the b_{k-1} value defined above and so on. For example, consider four merchants A , B , C and D bidding for three slots. Let all four of them have a click-through rate of 0.5 at the top slot, 0.4 at the middle slot and 0.2 at the bottom slot. Let the true valuations per click of the three merchants be 200, 150, 100 and 40 respectively. Let the ranking function be Google's revenue-ranking function. The merchants start off by bidding their true valuations, and A , B and C get the top, middle and bottom slot respectively, and make a profit of 25, 20 and 12 respectively. At this point, merchant B has an incentive to undercut C by bidding 80, which will result in B making a profit of 22, while C makes a reduced profit of 8. In order to remove any incentive for B to undercut her, C can change her bid to 70 as prescribed by the above formula. At this point B is making a profit of 32. Now, B faces the problem of A trying to undercut her by bidding 100 (say) in order to make a profit of 52, reducing B 's profit to 25. To remove A 's incentive to undercut her, B can change her bid to 86 as prescribed by the above formula.

With suitable assumptions, including separability of click-through rates, one can also use standard techniques such as the envelope theorem [7] to prove revenue equivalence. We omit the details since our first-principles analysis gives a stronger result in the form of a pure strategy Nash equilibrium under which the next-price auction is revenue-equivalent to the laddered auction, while the envelope theorem would only guarantee a mixed-strategy Nash equilibrium. Further, we obtain a simple and explicit characterization of the revenue-equivalent Nash equilibrium. Admittedly, these results show revenue equivalence to the next-price auction only. However, since the next-price auction is the auction currently in deployment, it is arguably the most interesting auction to consider in terms of showing revenue-equivalence.

Existence of Multiple Nash Equilibria: The foregoing discussion shows that there exists an equilibrium for the next-price auction which achieves the same ranking and the same revenue as the laddered auction. It should be pointed out that not all equilibria of the next-price auction have these properties. We next give an example that shows that there may exist other pure-strategy Nash equilibria under which the next-price auction achieves a smaller revenue than the truthful auction, and yet others under which the next-price auction achieves a higher revenue. Consider three merchants A , B and C having valuation-per-click of 500, 480 and 100 respectively bidding for two slots. Assume that the click-through rates are separable, and that all the merchants have the same μ_i of 1, and that the position-specific factors are given by $\theta_1 = 0.2$ and $\theta_2 = 0.15$. Let the ranking function be revenue-ranking. Assuming that everyone follows the dominant strategy of truth-telling, the laddered auction earns a revenue of $15 + (15 + 24) = 54$. Moreover, if everyone bids truthfully, then

the next-price auction would earn a revenue of $15 + 96 = 111$, more than twice the revenue of the laddered auction. However, truthful bidding is not an equilibrium for the next-price auction. One way to achieve equilibrium is for merchant A to change his bid to 110 in which case the revenue earned is $15 + 22 = 37$. On the other hand, if equilibrium is achieved by merchant B changing her bid from 480 to 200 before merchant A changes his bid, a different equilibrium is reached. In this case, the revenue earned is $15 + 40 = 55$. In this particular example, unless merchant B bids 200 or lower, merchant A will have an incentive to undercut her. This indicates that among all possible equilibria for this instance (excluding the ones where merchant C bids more than her true valuation, as there is no incentive for merchant C to do so), the highest revenue earned is 55. In order to achieve an equilibrium that achieves the same revenue as the laddered auction, merchant B could change her bid to 195, again preventing under-cutting by merchant A .

References

- [1] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 44–51, 2005.
- [2] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [3] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Working paper*, <http://rwj.berkeley.edu/schwarz/>, 2005.
- [4] J. Feng, H. Bhargava, , and D. Pennock. Implementing sponsored search in web search engines: computational evaluation of alternative mechanisms. *To appear in the Inform Journal on Computing*, 2006.
- [5] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [6] A. Mehta, A. Saberi, U. Vazirani, , and V. Vazirani. Adwords and generalized online matching. *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [7] C. Simon and L. Blume. *Mathematics for Economics*. W.W. Norton & Company, Inc., New York, 1994.
- [8] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

A Proof of theorem 3.1

Proof: Let $CTR_{i,j}$ be the click-through rate of merchant with index i at the j^{th} position. First note that if $CTR_{i,j}/CTR_{i,j+1} = CTR_{i',j}/CTR_{i',j+1}$ for all values of $i, i' \leq n$ and $j \leq K - 1$, then the click-through rates are separable: just set $\mu_i = CTR_{i,K}$ and $\theta_j = CTR_{1,j}/CTR_{1,K}$.

We show that if for every rank function $R = (w_1, w_2, \dots, w_n)$, there exists a set of VCG weights which always yield the same ranking as R , then

$$\text{CTR}_{i,j}/\text{CTR}_{i,j+1} = \text{CTR}_{i',j}/\text{CTR}_{i',j+1}$$

for all values of $i, i' \leq n$ and $j \leq K - 1$. We will prove this by downward induction on j .

First consider the base case of $j = K - 1$. Consider any pair of merchants. Re-index the merchants such that the two merchants are indexed j and $j + 1$. Let α denote the quantity $\text{CTR}_{j+1,j+1}/\text{CTR}_{j,j+1}$ and let ϕ denote $\text{CTR}_{j,j}/\text{CTR}_{j,j+1}$. Now, consider the ranking function R with $w_j = 1, w_{j+1} = \alpha$. All the other merchants are assigned a weight of 1. Suppose there exists a weighted VCG mechanism that always results in the same ranking as this ranking function. Let ω_i be the weight assigned by the VCG mechanism to merchant i , normalized such that $\omega_j = 1$. Then, the VCG mechanism chooses that ranking scheme that maximizes $\sum_{i=1}^K \omega_{m(i)} \text{CTR}(m(i), i) b_{m(i)}$, where $m(i)$ is the index of the merchant placed i in the ranking scheme. Let ρ be the ratio of the maximum click-through rate to the minimum click-through rate over all merchants and positions, and let ν be the ratio of the maximum VCG weight to the minimum VCG weight. Also, let $\beta_{max} = \max\{1, 1/\alpha\}$ and let $\beta_{min} = \min\{1, 1/\alpha\}$. Consider the following set of bids: $b_i = (2n)^{j-i} \rho \nu \beta_{max}$ for $i = 1, \dots, j - 1$, $b_j = 1, b_{j+1} = 1/\alpha$, and $b_i = \beta_{min}/((2n)^{i-(j+1)} \rho \nu)$ for the rest. Then, it is easy to verify that merchant i is placed at position i for $i = 1, 2, \dots, j - 1$ and merchants j and $j + 1$ share the remaining two positions (i.e., positions j and $j + 1$) under both the ranking function R as well as the VCG mechanism. The ranking score under R of merchant j and $j + 1$ is exactly the same, namely 1. Therefore, the ranking function R can be forced to place them in any chosen order by an infinitesimal change in the bids. In order for the VCG mechanism to produce the same ranking as R after the change, VCG must rate both possible orderings of j and $j + 1$ equally as well, i.e. the weighted sum of utilities (with the above bids) must be the same for both possible orderings.

$$\phi + \omega_{j+1} = 1 + \omega_{j+1} \frac{\text{CTR}_{j+1,j}}{\text{CTR}_{j+1,j+1}} \quad (8)$$

We could also have set the bids of the other merchants such that they get ranks $1, \dots, j$, leaving merchants j and $j + 1$ to compete for rank $j + 1 = K$. Then, a reasoning similar to above would show that

$$1 = \omega_{j+1} \quad (9)$$

Putting equations 8 and 9 together, we get

$$\frac{\text{CTR}_{j+1,j}}{\text{CTR}_{j+1,j+1}} = \phi.$$

This completes the proof of the base case.

By the induction hypothesis, we have $\text{CTR}_{i,j}/\text{CTR}_{i,j+1} = \text{CTR}_{i',j}/\text{CTR}_{i',j+1}$ for all values of $i, i' \leq n$ and $\hat{j} < j \leq K - 1$. Next consider $j = \hat{j}$. Consider the ranking function R with $w_i = 1$ for all merchants i . Let ω_i be the weight assigned by the corresponding VCG mechanism to merchant i . Again consider a pair of merchants and re-index merchants such that the pair is indexed j and

$j + 1$. Let $b_j = b_{j+1} = 1$. As before, we can set the bids of other merchants such that merchant i is ranked i for $i = 1, \dots, j - 1, j + 2, \dots, K$, while j and $j + 1$ share ranks j and $j + 1$. Since the ranking score given by R is the same for both j and $j + 1$, the VCG mechanism must also be ambivalent towards their order, i.e.

$$\begin{aligned}\omega_j \text{CTR}_{j,j} b_j + \omega_{j+1} \text{CTR}_{j+1,j+1} b_{j+1} \\ = \omega_j \text{CTR}_{j,j+1} b_j + \omega_{j+1} \text{CTR}_{j+1,j} b_{j+1}\end{aligned}$$

which implies that

$$\frac{\omega_{j+1}}{\omega_j} = \frac{\text{CTR}_{j,j} - \text{CTR}_{j,j+1}}{\text{CTR}_{j+1,j} - \text{CTR}_{j+1,j+1}}. \quad (10)$$

We can also set the bids of the other merchants such that they get ranks $1, \dots, j, j + 3, \dots, K$, leaving ranks $j + 1$ and $j + 2$ for merchants j and $j + 1$. Then, a reasoning similar to above would show that

$$\frac{\omega_{j+1}}{\omega_j} = \frac{\text{CTR}_{j,j+1} - \text{CTR}_{j,j+2}}{\text{CTR}_{j+1,j+1} - \text{CTR}_{j+1,j+2}} \quad (11)$$

From Equations 10 and 11, we get

$$\frac{\text{CTR}_{j,j} - \text{CTR}_{j,j+1}}{\text{CTR}_{j+1,j} - \text{CTR}_{j+1,j+1}} = \frac{\text{CTR}_{j,j+1} - \text{CTR}_{j,j+2}}{\text{CTR}_{j+1,j+1} - \text{CTR}_{j+1,j+2}}$$

Also, by induction hypothesis, we have

$$\frac{\text{CTR}_{j,j+1}}{\text{CTR}_{j,j+2}} = \frac{\text{CTR}_{j+1,j+1}}{\text{CTR}_{j+1,j+2}}$$

Using the above two equations and some elementary algebra, we get,

$$\frac{\text{CTR}_{j,j}}{\text{CTR}_{j+1,j}} = \frac{\text{CTR}_{j,j+1}}{\text{CTR}_{j+1,j+1}}$$

This completes the proof by induction. ■