# Keyword Auctions as Weighted Unit-Price Auctions

De Liu
Univ. of Kentucky

Jianqing Chen
Univ. of Texas at Austin

Andrew B. Whinston
Univ. of Texas at Austin

September 7, 2006

In recent years, unit-price auctions in which bidders bid their willingness-to-pay for per-unit realized outcome have been adopted by major keyword advertising providers such as Yahoo!, Google, and MSN, in selling keyword advertising slots on web sites. The majority of keyword auctions are pay-per-click auctions in which advertisers specify their willingness-to-pay per click, and pay by the clicks they actually receive.

In pay-per-click auctions, keyword advertising providers often learn about advertisers' abilities to generate clicks, e.g., by observing the click-through rates of the advertisement in the past. Such information has been gradually integrated into the keyword auction designs. For example, in 2003, Google started ranking advertisers by the product of their bid prices and their historical click-through rates. In 2005, Google adopted a more sophisticated ranking scheme that weighs advertisers' unit-price bids by "quality scores", which are determined by several factors including advertisers' historical click-through rates, the relevance of the advertisement text, and quality of landing pages. Are there economic justifications for auctioneers to weigh bidders' unit-price bids differently? What is the impact of such weighting schemes on bidding behavior, efficiency, and revenue?

The availability of click-through information also enables alternative minimum bid policies. For example, Google has recently abandoned their one-size-fits-all minimum-bid policy in favor of a new policy that imposes higher minimum bids for advertisements with low quality scores. These practices raise new questions: What is the impact of different minimum bid policies on bidding behavior, efficiency, and revenue? The above questions regarding ranking scheme and minimum bid policies are worthy, given the sheer size of the keyword advertising industry. Since 2002, the total revenue of Google has increased almost 15 fold, to $6.1 billion in 2005, most of which comes from keyword advertising.

In order to address the above questions, we study a class of auction mechanisms, weighted unit-price auctions (WUPAs), which capture the key features of keyword advertising auctions. A distinguishing feature of our setting is that the auctioneer can observe a signal about each bidder's future yield (which we term as the bidder's *type*) before the auction. The novelty of the WUPAs lies in that it provides an intuitive way for the auctioneer to take advantage of such interim information, i.e., by assigning different weighting factors and/or imposing different minimum bids for bidders of different types. Our main goals are (1) to study the impact of two design parameters of WUPAs, namely, the weighting scheme and the minimum bid policy, on the equilibrium bidding, allocation efficiency, and revenue, and (2) to compare the efficiency and revenue of WUPAs with those of standard auctions.

Our results have implications more beyond pay-per-click keyword auctions. First, our study can readily apply to other pricing models in keyword auctions. For example, some keyword advertising providers offer pay-per-call and pay-per-action formats, in which advertisers pay only if a customer responds to advertisements by making a phone call or taking a particular action such as signing up or purchasing advertised products. Since the "yield" in our model can also be interpreted as the the number of calls or actions in keyword advertising, the insights gained from our model apply to "pay-per-call" and "pay-per-action" auction formats as well. WUPAs may also be applied to other settings, such as selling computer power on a grid-computing facility or Internet bandwidth, procuring a project with uncertain (but verifiable) outcomes. In each of these settings, the auctioneer are likely to have some information on bidders' future yield, e.g., usage rate and probability of success.

Our research is related to a large literature on the role of ex post information in auction designs ([2], [5], [6], [7]). Unlike these papers, in our model setting, the auctioneer not only uses ex post outcome information but also ex ante information on bidders. Liu and Chen [4] have studied the weighting scheme of the WUPA class in a first-score, single-object setting. This paper advances [4] in several ways: (1) We show that the results obtained in [4] can be generalized to second-score and multi-object settings. (2) We also consider the case with minimum bids and discuss their impact on bidding strategies, allocation efficiency, and revenue. (3) We analyze the case where bidders' unit-valuation is correlated with their yield potential and evaluate the impact of such correlation on the optimal choice of weighting factor. The weighting scheme of keyword auctions is also discussed in Feng and Bhargava (2004) and Lahaie (2006). But they focus on two special cases where bids are either not weighted or are weighted by expected click-through rates.

**Model Setup.** Let there be $n$ bidders, indexed by $i = 1, 2, ..., n$ and $m$ $(m \leq n)$ assets, indexed by $j = 1, ..., m$. The bidder $i$'s valuation for asset $j$ is $v_i q_i \delta_j$, where $1 = \delta_1 > \delta_2 > ... > \delta_m$. We interpret $\delta_j$ as asset $j$'s size (or quality). Upon exploiting the asset $\delta_j$, each bidder $i$ generates a stochastic but verifiable yield $\delta_j q_i$, for which the bidder has a unit-valuation of $v_i$. We assume $q_i$ is not influenced by bidder $i$'s actions during the auction; rather, it is determined by bidder $i$'s intrinsic attributes. In the keyword advertising setting, $q_i$ is the number of clicks generated.

Unit valuation $v_i$ is bidder $i$'s private information. The bidder also observes a signal $\theta_i$ about his future yield $q_i$. The signal can be either $H$ or $L$. Let $F_H(v)$ and $F_L(v)$ denote the cumulative distributions of $v_i$ conditional on $i$ receiving signal $H$ and $L$ respectively. We assume $F_H(v)$ and $F_L(v)$ have a fixed support $[0, 1]$ and the density functions, $f_H(v)$ and $f_L(v)$, are positive and differentiable everywhere within the support.

The auctioneer can also observe each bidder's yield signal. This is possible because e-commerce technologies allow auctioneers to keep a record of bidders' yield rates in the past, which are used by both the bidder and the auctioneer to infer the bidder's future yield. We assume $E(q|\theta, v) = E(q|\theta)$ so that a bidder's private information $v_i$ has no implications on $q_i$. Denote $Q_H \equiv E(q|H)$ and $Q_L \equiv E(q|L)$ as the expected yield for a bidder who receives signal $H$ and signal $L$, respectively. We assume $Q_H > Q_L$ and call bidders who receive $H$ $(L)$ as high-yield (low-yield) bidders. Other bidders do not observe bidder $i$'s yield signal, but they hold a common belief about the probabilities of receiving $H$ and $L$, denoted as $\alpha$ and $1 - \alpha$, respectively.

**Auction Format.** The assets are sold through a unit-price auction in which each bidder $i$ is asked to submit his/her willingness-to-pay for per unit yield, $b_i$ (hereafter *unit price*), and the winners of the auction pays his/her realized yield at a unit-price $p$ determined by the auction rule. The auctioneer allocates the assets based on a weighted score rule. In particular, let the weight for low-yield bidders' bid prices be $w$ and normalize the weight for high-yield bidders' to 1. Then the score for bidder $i$ is

$$s_i = \begin{cases} b_i, \text{ if } i \text{ is high-yield} \\ wb_i, \text{ if } i \text{ is low-yield} \end{cases} \tag{1}$$

The bidder with the highest score wins the first asset, the bidder with the second-highest score wins the second asset, and so on. We call such an auction format as *weighted unit-price auctions*

(WUPAs). By allowing $w$ to take different values, we can accommodate the following stylized auction formats. When $w = 1$, the winners are determined solely by bid prices. One example is the auction format adopted by Yahoo!. When $w < 1$, bid prices from low-yield bidders are weighted less than those from high-yield bidders. Google's auction is among this category.

A WUPA may be *first-score* or *second-score* depending on whether the winning bidders are required to "fulfill" their own scores or the next highest score. We can show that the revenue equivalence holds for the two formats. In light of this, we will derive our main results using the first-score variation, which is simpler to work with.

All bidders are risk-neutral, and their payoff functions are additive in their total valuation and the payment. In particular, in first-score WUPAs, if the probability of winning asset $\delta_j$ with a score $s$ is $P_j(s)$, bidder $i$'s expected payoff is given by

$$u = E\left[q|\theta_i\right](v_i - b_i)\sum_{j=1}^{m}\delta_j P_j(s_i) \tag{2}$$

**Synopsis of Findings.** We focus on symmetric, pure-strategy Bayesian-Nash equilibria. By "symmetric", we mean that bidders with the same unit-valuation and yield-signal will bid the same. We denote their equilibrium probabilities of winning the first asset as $\rho_H^1(v)$ and $\rho_L^1(v)$ respectively.[1] Let minimum bids for high-yield and low-yield bidders be $\underline{b}_L$ and $\underline{b}_H$ respectively. Assuming $\underline{b}_L \le \underline{b}_H/w \le 1$ and $w \le 1$,[2] we have

$$\rho_L^1(v) = \begin{cases} \left[\alpha F_H(\underline{b}_H) + (1-\alpha)F_L(v)\right]^{n-1} & \text{for } v \in [\underline{b}_L, v_0] \\ \left[\alpha F_H(wv) + (1-\alpha)F_L(v)\right]^{n-1} & \text{for } v \in [v_0, 1] \end{cases} \tag{3}$$

$$\rho_H^1(v) = \begin{cases} \left[\alpha F_H(v) + (1-\alpha)F_L(v_0)\right]^{n-1} & \text{for } v \in [\underline{b}_H, wv_0] \\ \left[\alpha F_H(v) + (1-\alpha)F_L\left(\frac{v}{w}\right)\right]^{n-1} & \text{for } v \in [wv_0, w] \\ \left[\alpha F_H(v) + (1-\alpha)\right]^{n-1} & \text{for } v \in [w, 1] \end{cases} \tag{4}$$

where $v_0$ is the point at which lower-yield bidders jump their bids and is given by

$$w\int_{\underline{b}_L}^{v_0}\sum_{j=1}^{m}\rho_L^j(t)\delta_j dt = \int_{\underline{b}_H}^{wv_0}\sum_{j=1}^{m}\rho_H^j(t)\delta_j dt. \tag{5}$$

Given these equilibrium winning probabilities, we can easily evaluate the equilibrium bidding functions, efficiency, and the auctioneer's revenue. Our analysis generates the following novel findings:

- Unlike in standard auction settings, there may be kinks in bidding functions in WUPAs. The kinks occur when one type's bids rise (or drop) to a point the other type of bidders are no longer able to match. Figure 1 shows that the bidding function for high-yield bidders has a kink at $v = w$ when $w < 1$.

- With minimum bids, there may be a jump in one type's bidding function. The reason is that when two types of bidders face different minimum bids, the type who faces the more constraining minimum bid will have some of its bidders "clustering" near the minimum bid (zone-I high-yield bidders). As a result, it is rational for the other type to either leap-frog the

[1] The equilibrium winning probabilities of winning other assets, $\rho_L^j$ and $\rho_H^j$, $j = 2, ..., m$, can be derived analogously using different order statistics.

[2] The case with $w > 1$ and/or $\underline{b}_L > \underline{b}_H/w$ can be analyzed in a similar way.

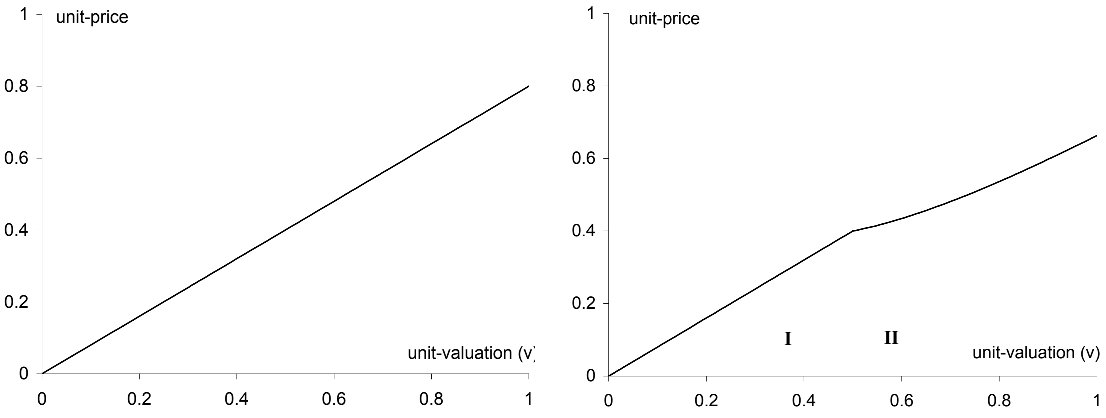


(a) low-yield bidder (b) high-yield bidder

Figure 1: equilibrium bidding functions without minimum bids

entire "cluster" or to bid lower than them. The jump does not occur only when minimum bids are equally constraining, i.e., when the minimum score is the same across bidder types. For example, in Figure 2, all zone-I low-yield bidders score lower than zone-I high-yield bidders (the cluster). There is a jump in low-yield bidders' bidding function at the border of zone-I and zone-II. All Zone-II low-yield bidders score higher than zone-I high-yield bidders. No

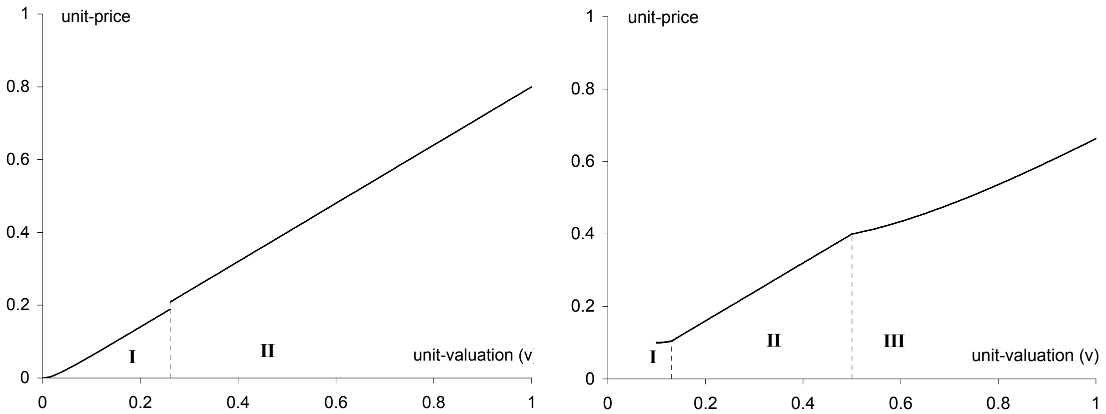


(a) low-yield bidder (b) high-yield bidder

Figure 2: equilibrium bidding functions with minimum bids

- Remarkably, although bidders generally do not bid their true valuation (except in the case of single-object WUPA with the second-score rule), the efficient weighting scheme is simple and the same in both the first-score and the second-score WUPAs: each unit-price bid is weighted by the respective expected yield, *as if bidders are bidding their true unit-valuation.* This insight is due to a result central to our findings: in equilibrium, the ratio of bid prices from the high-yield and the low-yield bidders *with the same score* is the same as the ratio of their true unit-valuation.

- It is weakly efficient (i.e., the allocation is efficient among participating bidders) to impose the same minimum score across bidder types, provided the auctioneer also uses the efficient weighting factor. As we have suggested, setting different minimum scores for different bidder types causes clustering among the more-constrained type and jump-bidding among the less-constrained the type, which distorts allocation efficiency locally.

- The revenue-maximizing weighting scheme is generally different from the efficient one. The revenue-maximizing weighting scheme may favor high-yield or low-yield bidders in comparison with the efficient weighting scheme. When high-yield bidders and low-yield bidders differ only in their expected yield, the optimal weighting scheme always favor low-yield bidders. But as the unit-valuation of low-yield bidders increases by the hazard-rate order, the revenue-maximizing weighting scheme will favor low-yield bidders less and may even favor the high-yield bidders. This finding is due to that when bidders differ both in the expected yield and in the unit-valuation distribution, the high-yield bidders are not necessarily "stronger".

- Combining the previous results with the finding that efficient WUPAs are revenue-equivalent to standard auctions in which bidders bid total willingness-to-pay, we conclude that WUPAs may generate higher revenue than standard auctions. The "abnormal" revenue arises precisely because the interim information on bidders' future-yield allows auctioneers to further discriminate among bidders, often at the cost of allocation efficiency.

- An optimal minimum bids policy is generally different from what we see in the optimal mechanism design literature. The reason is that besides the effect of excluding bidders of certain valuation, the minimum bids also affect the resource allocation among the *participating* bidders, particularly through its impact on the "jump" point and the size of the "cluster". Furthermore, the optimal minimum-bids are generally not a uniform minimum score either, thus inefficient.

# References


[1] Feng, Juan, H. Bhargava, and D. Pennock (2005). Implementing paid placement in web search engines: Computational evaluation of alternative mechanisms. INFORMS Journal on Computing, forthcoming.

[2] Laffont, J. J. and J. Tirole (1988). "The Dynamics of Incentive Contracts." Econometrica 56(5): 1153-1175.

[3] Lahaie, S. (2006). An Analysis of Alternative Slot Auction Designs for Sponsored Search. Proceedings of the 7th ACM Conference on Electronic Commerce, Ann Arbor, Michigan, ACM Press.

[4] Liu, D. and J. Chen (2005). "Designing online auctions with performance information." Decision Support Systems, forthcoming.

[5] McAfee, R. P. and J. McMillan (1986). "Bidding for Contracts - a Principal-Agent Analysis." Rand Journal of Economics 17(3): 326-338.

[6] Riley, J. G. (1988). "Ex-Post Information in Auctions." Review of Economic Studies 55(3): 409-429.

[7] Samuelson, W. (1986). "Bidding for Contracts." Management Science 32(12): 1533-1550.