

Designing online auctions with past performance information

De Liu ^{a,*}, Jianqing Chen ^b

^a Gatton College of Business and Economics, University of Kentucky, Lexington, KY 40506, United States

^b Department of Information, Risk, and Operations Management, Center for Research in E-Commerce,
University of Texas at Austin, Austin, TX 78712, United States

Received 25 September 2005; received in revised form 26 October 2005; accepted 1 November 2005

Available online 20 December 2005

Abstract

We investigate the value of past performance information in the context of keyword advertising auctions, where advertisers differ both in valuation-per-click and in the numbers of clicks they can generate (their *performance*). We focus on weighted unit-price-contract (UPC) auctions, in which bidders bid unit prices and pay accordingly if they win, and their bids are weighted by factors based on their own past performance. We characterize the efficient and the revenue-maximizing weighting factors and apply our framework to study Yahoo!'s and Google's auction designs, each of which can be viewed as a special case of weighted UPC auctions.

© 2005 Elsevier B.V. All rights reserved.

Keywords: UPC auctions; Search engine; Google; Keyword advertising; Multi-dimensional values

1. Introduction

The majority of auction design literature has made assumptions that bidders' valuation can be ordered along a single dimension, largely to facilitate solving the design problem. This is not the case, however, in auctions for advertising slots on search engines.¹ Advertisers may differ both on their valuation-per-click on their advertisements and on their abilities to generate clicks. It is no doubt that advertising interme-

diaries (search engines) can still ask advertisers to bid on their total payment so that traditional mechanisms such as first-price sealed-bid or English auctions can be applied. However, one may wonder whether alternative mechanisms could better accommodate the underlying bi-dimensionality of bidders' valuation. This paper explores one Google-like approach that makes use of the information on advertisers' past click-through rates (CTRs). In Google's approach, each advertiser submits a bid on how much they are willing to pay for every click, and the assignment of advertising slots is based on a score rule that weighs advertisers' bid price by their past CTRs.² This auction mechanism, which we shall call weighted *unit-price contract* (UPC) auctions, has two essential features. First, advertisers bid unit

* Corresponding author.

E-mail addresses: de.liu@uky.edu (D. Liu),
chenjq@mail.utexas.edu (J. Chen).

¹ Keyword advertising is an online marketing service provided by search engines (e.g., Google), in which advertisers can specify keywords under which they want their advertisements to appear. These advertisements usually appear along side or above the search results and are labeled as "sponsored links" or "sponsored results". Most search engines use auctions to sell their keyword advertising slots.

² Recently, Google allows the weighting factors to be jointly determined by advertisers' past CTRs, the relevance of their advertisements to keywords, and other relevancy factors.

prices, and winners pay for the actual clicks at unit prices determined by the auction. Hence, advertising intermediaries essentially sell *unit-price contracts* to advertisers. Second, the allocation is based on a score rule that weighs bids by a factor that incorporates the available information on other dimensions of bidders' characteristics.

To our knowledge, weighted UPC auctions as a way of using past information on bidders have not been well studied by researchers. Many questions are pending: how do rational bidders behave in weighted UPC auctions? How do weighted UPC auctions perform, compared to benchmarks such as first-price sealed-bid or second-price auctions? How to choose weighting factors to improve resource allocation efficiency or to maximize the auctioneer's revenue? Is the information about advertisers' past CTRs useful to advertising intermediaries? These are questions to be addressed in this paper.

Answers to the above questions are of importance to the online advertising industry, which is expected to reach US\$13.8 billion in total revenue by 2006 [2]. Google doubled its revenue to US\$3.19 billion, its net income increased fourfold to US\$399 million in 2004, and its market valuation exceeded US\$50 billion. While most search engines have adopted UPC auctions in selling advertising slots, their ways of ranking advertisers differ. As we have mentioned above, Google weighs advertisers' bids, i.e., the price they are willing to pay for each click on their advertisements, by their past CTRs. Yet Yahoo!, another leader in the industry, ranks advertisers solely by their bids. This paper provides a theoretical framework to compare these different ranking mechanisms.

Though we discuss the above questions in the context of keyword advertising, weighted UPC auctions may have wider applications. They may also be used in selling access to other resources, such as publishing rights, rental properties, store fronts in an electronic marketplace, and Internet bandwidth, provided that the outcome from using these resources is verifiable and information on non-price dimensions of bidders' characteristics is available to the auctioneer. Weighted UPC auctions are especially suitable for use on the Internet, which can often provide efficient ways of tracking bidders' past performance.

Our investigation takes place in a setting where a risk-neutral intermediary auctions off one advertising slot to n risk-neutral advertisers. Advertisers differ not only in their valuation-per-click but also in their past CTRs, which serve as signals about their future CTRs. Both valuation-per-click and signals for future CTRs

are symmetrically and independently distributed across advertisers. Advertisers privately learn their valuation-per-click and signals about their future CTRs. In addition, signals about advertisers' future CTRs are also learned by the intermediary. The intermediary's problem is to choose weighting factors for advertisers with different signals for future CTRs to maximize its revenue or resource-allocation efficiency.

We show that when advertisers' bids are weighed by their expected CTRs, the auction is efficient. The efficient UPC auction also generates the same amount of expected revenue as a conventional first-price sealed-bid auction (or other revenue-equivalent formats). Efficient UPC auctions are not necessarily revenue-maximizing, however. We show that under the increasing hazard-rate condition, the revenue-maximizing UPC auctions should favor advertisers with low expected CTRs in the sense that the weighting factor for advertisers with low expected CTRs should be higher (or equivalently, the weighting factor for advertisers with high expected CTRs should be lower) than it is in efficient UPC auctions. By favoring advertisers with low expected CTRs, advertising intermediaries can force advertisers with high expected CTRs to bid more aggressively, which will more than compensate the loss due to allocation inefficiency. The fact that weighted UPC auctions can generate higher expected revenue than standard high-bid auctions suggests that information about advertisers' past is indeed valuable to advertising intermediaries.

By studying weighted UPC auctions under the uniform distribution, we reveal that the optimal weighting factor for advertisers with low expected CTRs decreases with the total number of bidders, but is bounded away from the efficient weighting factor. The intuitive explanation is that when the number of advertisers increase, the increased competition among advertisers with high expected CTRs drives down their information rents, reducing the need of using advertisers with low expected CTRs to increase competition. We also find that the optimal weighting factor for advertisers with low expected CTRs increases with the ratio of low expected CTRs to high expected CTRs.

Finally, the expected revenue from a standard UPC auction (analogous to Yahoo!'s) and from an efficient UPC auction (analogous to Google's) can not be unanimously ranked. Our numerical results suggest the latter design appears to outperform the former when the number of bidders is large (the converse is true when the number of bidders is small).

The paper proceeds as follows. In Section 2 we review related literature. Section 3 lays out our

model, followed by an analysis of bidding functions and the intermediary's expected revenue in Section 4. Section 5 examines the efficient and the revenue-maximizing designs and compares standard, efficient, and optimally weighted UPC auctions. Section 6 concludes the paper.

2. Related literature

Information plays an important role in auction designs. Milgrom and Weber [11] show that auctioneers can gain by revealing their information about the object *ex ante* if bidders' private estimates of the item's value are affiliated. Riley [12] argues that auctioneers can increase their revenues in sealed-bid auctions by making the winner's payment dependant on *ex post* public information about the item's value, such as by introducing positive royalty rates. We focus on investigating UPC auctions as alternative auction mechanisms that can make use of *ex ante* information about bidders in a private-value setting.

Our research relates to the study of asymmetric auctions [10]. In fact, by making signals for expected CTRs public before the auction and limiting the number of bidders to two, we can turn our model into one similar to that in [10]. The key difference between an asymmetric auction model and ours is that in our model an advertiser is not sure whether an opponent has high or low expected CTRs, which implies all advertisers have symmetric beliefs about their opponents. Besides, much of the asymmetric auction literature focuses on comparing the revenues generated by first-price sealed-bid and English auctions, while we study weighted UPC auctions.

The weighted UPC auctions have one feature in common with scoring auctions used in the procurement setting: a single score is computed for each bidder and the allocation is solely determined by scores. A few authors [3,5] study settings where suppliers submit multidimensional bids—e.g., cost and quality—but their private information is one-dimensional. Asker and Cantillon [1] generalize this line of research to allow suppliers to have multidimensional private information. The score rules in the above-mentioned research are quasi-linear in the monetary dimension of bids, and therefore different from this paper. Ewerhard and Fieseler [6] examine UPC auctions in the procurement setting where suppliers bid unit prices for every factor needed to accomplish a task (e.g., highway contracting). In their model, scores are computed as weighted average of factor prices.

Our paper also relates to recent work on search engine pricing. Hu [9] considers whether to incorporate CTRs as a part of incentive contracts between intermediaries and advertisers in a principal-agent setting. Feng et al. [8] use numerical simulations to compare Google's and Yahoo!'s methods of ranking advertisers. Feng [7] studies the optimal allocation mechanism in a multiple-slot setting where advertisers' valuations decrease in the ranks of slots at different speeds. Her work therefore complements ours. Weber and Zheng [13] study a search market model that encompasses both consumers' search problem and advertisers' bidding problem. In their model, advertisers' valuation of an advertising slot is determined by a single parameter, and the intermediary ranks advertisers by the weighted average of their bids and the social surplus they generate. Thus, their model setting is different from ours.

3. Model setup

We consider an online intermediary (hereafter intermediary) who auctions off a single advertising slot to n advertisers at a given time period. The total traffic to the advertising slot is exogenously determined, normalized to 1.

All advertisers are risk-neutral. Each advertiser's total valuation for the advertising slot is its valuation-per-click, v , times the number of clicks it can generate, r , which, in our case, coincides with its click-through rate.³ An advertiser's payoff from winning the slot is additive in its total valuation for the slot and the money it pays, that is, if an advertiser pays P for the slot, its payoff from winning is given by

$$U = rv - P. \quad (1)$$

The difference in advertisers' valuation-per-click may arise from their different abilities to turn a visitor into a buyer and/or from the different profit they make from each purchase. Advertisers may have different CTRs for several reasons. First, the relevance of the advertisement to search traffic varies from one advertiser to another. Second, the presentation of an advertisement may become a differentiation factor. Third, the CTR may differ due to different degrees of brand recognition. It is worth noting that the click-through rate in our setting solely reflects the advertiser's (advertisement's) intrinsic ability to attract clicks, whereas the click-through rate in a general setting may also be

³ Since we have normalized total traffic to size 1, the click-through rate and the number of clicks are interchangeable.

influenced by where (how, when) the slot is placed.⁴ How the slot is presented is common to all potential winners, and therefore can be regarded as the attribute of the advertising slot.

Each advertiser learns its own valuation-per-click before the auction, but not others'.⁵ All advertisers and the intermediary hold a common belief about the distribution of v , denoted as $F(v)$. We assume $F(v)$'s density function, $f(v)$, has a fixed support $[0,1]$, and is positive and differentiable everywhere within the support. We let $F(v)=1$ for all $v>1$.

Each advertiser learns a signal θ before the auction, which allows it to make an inference about its own future click-through rate. The same signal is also learned by the intermediary who will make the same inference about the advertiser's future click-through rate. θ is not observed by other bidders, however.⁶ One such signal is advertisers' past CTRs. To simplify our analysis, we assume $\theta \in \{H, L\}$. Advertisers who receive signal H have high expected CTR, $E_H \equiv E[r|\theta=H]$, and those who receive signal L have low expected CTR, $E_L \equiv E[r|\theta=L]$ ($E_L < E_H$). We also call the former *H-type* advertisers, and the latter *L-type* advertisers. We assume the probabilities for an advertiser to receive a signal H and L are α and $1-\alpha$ ($0 < \alpha < 1$), respectively, and are common knowledge.

The advertising slot is sold through a first-price, sealed-bid UPC auction, i.e., one in which each advertiser places a bid b in a sealed form, and if it wins, will pay for all generated clicks at the unit price b . We assume there is no entry fee or reserve price. The auctioneer assigns each advertiser a score which is a product of the advertiser's bid and a weighting factor based on the type of signal the advertiser receives. The advertiser with the highest score will be the winner of the auction. Since neither the allocation of the slot nor the winner's payment is affected by rescaling the weighting factors, we can normalize the weighting factor for *H-type* advertisers to be 1 without loss of

generality. Let γ ($\gamma > 0$) denote the weighting factor for *L-type* advertisers. The score for an advertiser who receives a signal θ and places a bid b is given by

$$s(b, \theta) = \begin{cases} b, & \text{if } \theta = H \\ \gamma b, & \text{if } \theta = L \end{cases} \quad (2)$$

The intermediary's problem is to choose γ to maximize its expected revenue or the resource-allocation efficiency.

By allowing γ to take different values, we can accommodate different auction formats. When $\gamma=1$, the winners are determined solely by bids, like in the auction format adopted by Yahoo!. We call the case $\gamma=1$ *standard* UPC auctions. When $\gamma < 1$, bids from advertisers with lower expected CTRs are weighed less than those with higher expected CTRs. One version of $\gamma < 1$ is implemented by Google. To our knowledge, Google weighs advertisers' bids on a particular keyword by their CTRs over their entire history on this keyword, which, if we believe high past CTRs lead to high future CTRs, translates to the case where bids from advertisers with high expected CTRs are weighted more.

To summarize, every advertiser gets a signal about its expected CTR and learns its valuation-per-click before the auction. Each advertiser then places a bid b as its pay per click. The intermediary selects the winner based on the pre-announced score rule specified in (2). The winner's advertisement is displayed in the slot during the period. At the end of the period, the winner pays the intermediary according to the actual number of clicks at unit price b .

4. The bidding functions and expected revenue

We consider a symmetric, pure-strategy Bayesian-Nash equilibrium for the above problem. By "symmetric", we mean bidders with the same valuation-per-click and the same expected CTR will bid the same in equilibrium. Let $\beta_H(v)$ ($\beta_L(v)$) denote the mapping from an *H-type* (*L-type*) bidder's valuation-per-click to its equilibrium bid. We term $\beta_H(v)$ and $\beta_L(v)$ as the *bidding functions* for *H-type* and *L-type*, respectively.

The lowest-valuation bidder ($v=0$) will always bid zero in equilibrium. Obviously, a bidder with zero valuation-per-click will never bid more than zero since it would incur a loss if it were to win the auction. Thus, assuming negative bids are not allowed, we must have $\beta_L(0)=\beta_H(0)=0$.

We conjecture that both *H-type*'s and *L-type*'s bidding functions are strictly increasing (we verify this in Appendix 2). Let $\bar{b}_H \equiv \beta_H(1)$ and $\bar{b}_L \equiv \beta_L(1)$ denote the upper bounds of bids from *H-type* and *L-type* bidders.

⁴ For example, an advertising slot at the top of the page can normally attract more clicks (therefore generate higher click-through rates) than a same-sized slot located at the bottom of the page. This latter factor is common to all advertisers and therefore can be attributed to the advertising slot.

⁵ Advertisers may not be informed about opponents' valuations for several reasons: first, it may be impossible to infer an advertiser's valuation-per-click from their ranking because intermediaries rarely publish bids. Second, even if one could learn opponents' past bids, it is still difficult to infer their valuation because their bids may also be affected by their past CTRs. Third, advertisers may drop out of an auction or join one at any time—the dynamic pool of bidders makes it difficult to infer valuation-per-click.

⁶ Neither Google nor Yahoo! publishes advertisers' CTR information.

Inverse bidding functions exist, denoted as $\phi_H(b)$, $b \in [0, \bar{b}_H]$ and $\phi_L(b)$, $b \in [0, \bar{b}_L]$, respectively.

An L-type bidder who bids b wins the auction if and only if all other L-type bidders bid less than b , and all H-type bidders bid less than γb . The probability for an L-type bidder to bid less than b is $F(\phi_L(b))$ if $b \leq \bar{b}_L$ and is 1 if $b > \bar{b}_L$. To simplify presentation, we define

$$\phi_i(b) = 1, \text{ for } b > \bar{b}_i, \quad i = \{H, L\}. \quad (3)$$

By definition (3) the probability for an L-type bidder to bid less than b can be uniformly written as $F(\phi_L(b))$. The probability for an H-type bidder to bid less than γb is $F(\phi_H(\gamma b))$. An L-type bidder's expected payoff is

$$U_L(v, b) = E_L(v - b)[\alpha F(\phi_H(\gamma b)) + (1 - \alpha)F(\phi_L(b))]^{n-1}. \quad (4)$$

Similarly, an H-type bidder's expected payoff is

$$U_H(v, b) = E_H(v - b)[\alpha F(\phi_H(b)) + (1 - \alpha)F(\phi_L(b/\gamma))]^{n-1}. \quad (5)$$

The optimal bidding functions $\beta_L(v)$ and $\beta_H(v)$ should necessarily satisfy the first order conditions

$$\begin{aligned} \frac{\partial U_L(v, b)}{\partial b} \Big|_{b=\beta_L(v)} &= 0 \text{ and } \frac{\partial U_H(v, b)}{\partial b} \Big|_{b=\beta_H(v)} \\ &= 0, \text{ for all } v \in [0, 1]. \end{aligned} \quad (6)$$

Conditions (6) yield two differential equations involving ϕ_L and ϕ_H . Lemma 1 is critical for us to explicitly solve (6).

Lemma 1.

$$\text{If } \gamma \leq 1, \beta_H(\gamma v) = \gamma \beta_L(v), \quad \forall v \in [0, 1] \quad (7)$$

$$\text{If } \gamma > 1, \beta_H(v) = \gamma \beta_L(v/\gamma), \quad \forall v \in [0, 1] \quad (8)$$

Proof. See Appendix for all proofs. \square

Lemma 1 says in equilibrium an H-type bidder with valuation-per-click γv and an L-type bidder with valuation-per-click v will bid the same score. The intuition for Lemma 1 is as follows. Consider an H-type with valuation-per-click γv who bids γb and an L-type with valuation-per-click v who bids b . By the score rule, the former has the same winning probability as the latter. The payoff, conditional on winning, of the former differs from that of the latter only by a scalar. Since they also have the same probability of winning, their expected payoff functions differ only by a scalar too. The fact that multiplying a payoff function by a scalar does not alter the solution to an optimization problem

implies that b maximizes the L-type's expected payoff if and only if γb maximizes H-type's, hence Lemma 1.

Lemma 1 implies that

$$\text{If } \gamma \leq 1, \phi_H(\gamma b) = \gamma \phi_L(b), \quad \forall b \in [0, \bar{b}_L] \quad (9)$$

$$\text{If } \gamma > 1, \phi_H(b) = \gamma \phi_L(b/\gamma), \quad \forall b \in [0, \bar{b}_H]. \quad (10)$$

Based on Lemma 1 and (9) and (10), we can solve the two differential equations in (6) separately.

Proposition 1. Given $\gamma > 0$, the equilibrium bidding functions are given by

$$\begin{cases} \beta_L(v) = v - \frac{\int_0^v [\alpha F(\gamma t) + (1 - \alpha)F(t)]^{n-1} dt}{[\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1}} \\ \beta_H(v) = v - \frac{\int_0^v [\alpha F(t) + (1 - \alpha)F(t/\gamma)]^{n-1} dt}{[\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1}} \end{cases}, \quad (11)$$

for all $v \in [0, 1]$.

Both L-type and H-type advertisers bid less than their true valuation (except for zero-valuation bidders), which is common among first-price auctions. In the Appendix we verify that both $\beta_L(v)$ and $\beta_H(v)$ are monotonically increasing.

The kinks in equilibrium bidding functions are worth noting. When $\gamma < 1$, $F(t/\gamma) = 1$ for $t \geq \gamma$, so the H-type's bidding function has a kink at $v = \gamma$. Specifically,

$$\beta_H(v) = \begin{cases} v - \frac{\int_0^v [\alpha F(t) + (1 - \alpha)F(t/\gamma)]^{n-1} dt}{[\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1}}, & v \in [0, \gamma] \\ v - \frac{\int_0^\gamma [\alpha F(t) + (1 - \alpha)F(t/\gamma)]^{n-1} dx + \int_\gamma^v [\alpha F(t) + 1 - \alpha]^{n-1} dt}{[\alpha F(v) + 1 - \alpha]^{n-1}}, & v \in [\gamma, 1] \end{cases}$$

The intuition is as follows. We consider the case of $\gamma \leq 1$ first. It is straightforward to see that an H-type bidder with valuation-per-click $v < \gamma$ can possibly lose the auction to either H-types or L-types (by Lemma 1). In other words, those H-type bidders have to face competition from both H-types and L-types in equilibrium. On the other hand, an H-type bidder with high valuation ($v \geq \gamma$) will only face competition from other H-types, since, according to Lemma 1, L-type bidders with the highest valuation ($v = 1$) will bid the same score as H-type bidders with valuation-per-click $v = \gamma$, and thus no L-type bidders can match the score bid by an H-type bidder with even higher valuation. Thus, H-type bidders who have valuation-per-click lower than γ and who have valuation-per-click higher than γ face different numbers of competitors. As a result, the H-type's bidding function has a kink at $v = \gamma$. Similarly, when $\gamma > 1$, the highest-valuation H-type bidders will bid the




Fig. 1. Bidding functions under two different weighting factors.

same score as an L-type bidder with valuation $v=1/\gamma$. Therefore the L-type bidder's bidding function has a kink at $v=1/\gamma$.

The following example illustrates the impact of the weighting factor on bidding functions using two stylized UPC designs, $\gamma=1$ (the standard UPC auction) and $\gamma=E_L/E_H=0.5$ (the weighted UPC auction). Notice that if $\gamma=1$, H-type and L-type bidders will follow the exactly same bidding strategy. This is because when $\gamma=1$ their payoff functions differ only by a scalar. When $\gamma=0.5$ there is a kink in H-type's bidding function. H-types with valuation-per-click exceeding the kink bid less aggressively than in the case $\gamma=1$, since they can beat any L-type and face the only competition from other H-types. It is worth noting that although the L-type bidder with value 1 bids much more than the H-type bidder with value γ , they have the same score and therefore the same winning probability in this weighted UPC auction.

Example 1. Let $F(v)=1-(1-v)^2$, $n=10$, $E_L=0.05$, $E_H=0.1$, and $\alpha=0.3$. Fig. 1 illustrates the bidding functions under the standard UPC auction ($\gamma=1$), and the UPC auction weighted by expected CTRs ($\gamma=E_L/E_H=0.5$), respectively. In the latter case, the H-type's bidding function has a kink at $v=0.5$.

We can explicitly evaluate the expected revenue of the intermediary.

Proposition 2. *The expected revenue of the intermediary is*

$$n(1-\alpha)E_L \int_0^1 [\alpha F(\gamma v) + (1-\alpha)F(v)]^{n-1} \times \left(v - \frac{1-F(v)}{f(v)} \right) f(v) dv + n\alpha E_H \int_0^1 [\alpha F(v) + (1-\alpha)F(v/\gamma)]^{n-1} \times \left(v - \frac{1-F(v)}{f(v)} \right) f(v) dv. \quad (12)$$

The first term above is the expected revenue from L-type bidders and the second is expected revenue from H-type bidders. $[\alpha F(\gamma v) + (1-\alpha)F(v)]^{n-1}$ is L-type advertisers' equilibrium probability of winning and $[\alpha F(v) + (1-\alpha)F(v/\gamma)]^{n-1}$ is H-type advertisers'. $E_L(v - \frac{1-F(v)}{f(v)})$ and $E_H(v - \frac{1-F(v)}{f(v)})$ can be viewed as the "marginal revenues" generated by L-type and H-type advertisers, respectively, in Bulow and Robert's [4] terms. The intermediary's problem is reduced to maximize its expected "marginal revenue" from L-types and H-types by choosing two winning probability functions: $[\alpha F(\gamma v) + (1-\alpha)F(v)]^{n-1}$ for L-types and $[\alpha F(v) + (1-\alpha)F(v/\gamma)]^{n-1}$ for H-types. Though the expected revenue in weighted UPC auctions (12) looks similar to those in asymmetric auctions [10], there are a few notable differences. First, unlike asymmetric auctions, weighted UPC auctions allocate resources based on classes (in this case, expected CTRs) instead of identities. This feature is useful when discrimination based on identities is not possible. Second, the value of the weighted UPC auction lies in the distinctive set of allocation plans it offers. To our knowledge, this set of allocation plans has not been studied in previous literature. Though the weighted UPC auctions are a subset of all possible mechanisms, they have the advantage of intuitive implementation.

It is worth pointing out that changing γ not only affects H-type's and L-type's bidding functions, but also affects their equilibrium winning probabilities. In particular, decreasing γ tends to (but not always) cause H-type bidders' to bid less, which has a negative effect on total revenues, and L-type bidders to bid more, which has a positive effect (see Fig. 1 for an example). Meanwhile, decreasing γ will also increase H-type bidders' winning probabilities ($[\alpha F(v) + (1-\alpha)F(v/\gamma)]^{n-1}$) and decrease L-type bidders' ($[\alpha F(\gamma v) + (1-\alpha)F(v)]^{n-1}$), which has a positive effect on total revenues. The overall effect of the change γ depends on the balances of the above effects. In Example 1, we can calculate that as γ decreases from 1 to 0.5, the total expected revenue increases from 0.0387 to 0.0394. In the next, we examine the efficient and the optimal weighting factors.

5. The efficient weighting factor and the optimal weighting factor

When an advertiser with valuation-per-click v and expected click-through rate E_θ , $\theta \in \{H, L\}$ wins the advertising slot, it creates a social surplus of vE_θ . Given the probabilities of winning for H-type and L-type advertisers are $[\alpha F(\gamma v) + (1-\alpha)F(v)]^{n-1}$ and

$[\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1}$, respectively, the total *social surplus* generated by a weighted UPC auction is given by

$$W = n(1 - \alpha)E_L \int_0^1 [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1} v f(v) dv \\ + n\alpha E_H \int_0^1 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1} v f(v) dv. \quad (13)$$

We define the *efficient weighting factor*, γ^* , as the factor that maximizes W . We call the UPC auction with an efficient weighting factor an *efficient UPC auction*. Meanwhile, we denote W^{FB} as the *first best social surplus* that a social planner can achieve under a complete-information setting. The first best efficiency is reached when an advertiser is assigned the slot if and only if it has the highest total valuation (vE_θ).

Proposition 3. *The efficient weighting factor is given by $\gamma^* = E_L/E_H$. When $\gamma = \gamma^*$, $W = W^{FB}$. Moreover, the expected revenue of the efficient UPC auction equals that of a standard first-price auction where advertisers bid their total payment.*

According to Lemma 1, when the weighting factor is E_L/E_H , an H-type advertiser whose valuation-per-click is E_L/E_H times of an L-type advertiser's will bid E_L/E_H times of the L-type advertiser's bid—so they will tie. Because they also have the same total expected valuation for the advertising slot, the ranking of bids is in fact consistent with their total valuation, thereby assuring allocation efficiency. It follows that efficient UPC auctions allocate the same way as standard first-price auctions and generate the same expected revenue to auctioneers. Bidders also pay the same expected amount, except that in efficient UPC auctions they pay a fee that varies according to the ex post outcome, whereas in standard auctions, an upfront lump sum. In this regard, bidders in efficient UPC auctions assume less risk than those in standard auctions.

We define the *optimal weighting factor*, γ^{**} , as the factor that maximizes the total expected revenue of the intermediary. We call the UPC auction with an optimal weighting factor an *optimally weighted UPC auction*. The optimal weighting factor may depend on a number of factors, including the number of bidders, the distribution of valuation, the signals (E_L and E_H), and the distribution of the signals (α). The explicit formula for γ^{**} is not generally attainable, except for some special distributions (see Corollary 1). So we turn to characterize the boundaries for the optimal weighting factor.

Proposition 4.

(a) *The optimal weighting factor $\gamma^{**} > E_L/E_H$, if the distribution function F satisfies the property of increasing hazard rate (IHR), i.e.,*

$$\frac{d}{dv} \left(\frac{f(v)}{1 - F(v)} \right) \geq 0 \text{ for any } v \quad (14)$$

(b) *The optimal weighting factor $\gamma^{**} \leq 1$, if F satisfies IHR and the negative impact of raising γ on H-type advertisers' winning probability increases in v , i.e.,*

$$\partial^2 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1} / \partial \gamma \partial v \leq 0 \text{ for any } \gamma \geq 1 \\ \text{and any } v. \quad (15)$$

In our setting, the IHR property is interpreted as that an advertiser's valuation-per-click above a certain threshold value is more likely to fall into the low end of its range as the threshold value increases. The IHR property is a frequently made assumption in games of incomplete information. The property is known to be satisfied by a wide range of distributions, including uniform, normal, and exponential.

Proposition 4 suggests that when (14) and (15) are satisfied, the optimal weighting factor lies between the efficient weighting factor γ^* and 1. The intuition is as follows. When $\gamma = E_L/E_H$, an L-type bidder with valuation-per-click v bids the same score as an H-type bidder with valuation-per-click vE_L/E_H in equilibrium. When the IHR condition holds, the “marginal revenue” from the former $\left(E_L \left(v - \frac{1-F(v)}{f(v)} \right) \right)$ is higher than the “marginal revenue” from the latter $\left(E_H \left(\gamma v - \frac{1-F(\gamma v)}{f(\gamma v)} \right) \right)$. We can show that marginal effect of increasing γ is to redistribute part of H-type bidders' winning probability to L-type bidders who bid the same score as the former in the equilibrium (see Appendix 7). Because the intermediary's total expected total revenue is the total expected “marginal revenue” from all winning bidders (Proposition 2), increasing γ will increase the expected total revenue for the intermediary. Thus, the optimal weighting factor should generally be higher than the efficient weighting factor γ^* . It is worth nothing that this result does not depend on proportion of H-type bidders (α)—although the optimal weighting factor (γ^{**}) may do.

When γ keeps increasing, the difference between an L-type's marginal revenue and that of an H-type with the same score in equilibrium reduces or even reverses for some v 's. The optimal γ^{**} is generally determined endogenously by all model parameters. Proposition 2

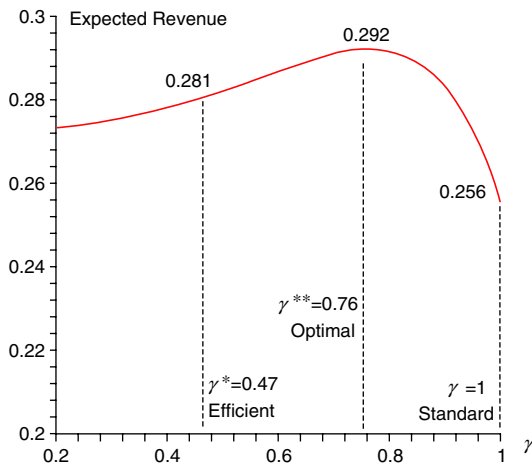


Fig. 2. Expected revenues under different weighting factors.

shows that when regularity condition (15)⁷ holds, the optimal weighting factor γ^{**} is bounded above by 1.

Proposition 4 implies that the optimally weighted UPC auction is generally sub-efficient: it may award the advertising slot to an L-type advertiser even though an H-type advertiser may value it more. Because the efficient UPC auction generates the same expected revenue as standard auctions (Proposition 3), Proposition 4 also implies that optimally weighted UPC auctions can generate more revenue than standard first-price auctions that do not take advantage of past performance information on bidders.

To facilitate understanding the optimal weighting factor, we analyze the case in which valuation-per-click is uniformly distributed on $[0,1]$. The following corollary establishes the optimal weighting factor under the uniform distribution and some comparative statics about it.

Corollary 1. *Assuming v is uniformly distributed on $[0,1]$, the optimal weighting factor is given by*

$$\gamma^{**} = \frac{(n-1)E_L + (n+1)E_H}{2nE_H} \quad (16)$$

$E_L/E_H < \gamma^{**} < 1$, γ^{**} decreases in E_H/E_L and in n .

Corollary 1 shows that under the uniform distribution the optimal weighting factor is always between E_L/E_H and 1. When n is large, γ^{**} is approximately $\frac{E_L + E_H}{2E_H}$. This implies the optimal weighting factor is bounded away from the efficient weighting factor even in a limit case. When the ratio of an H-type's expected CTR to an L-type's expected CTR increases, it is more profitable to

let H-types win. As a result, the intermediary should decrease the weighting factor for L-types. γ^{**} decreases in n because the competition within H-type advertisers will increase when n increases, reducing the need to induce competition by L-type advertisers. In such a case, the intermediary should decrease L-type's weighting factor to reduce allocation efficiency distortion.

In the following example, we numerically compare the expected revenues of the standard, efficient, and optimally weighted UPC auctions.

Example 2. Assume the valuation-per-click is uniformly distributed on $[0,1]$, $n=10$, $E_L=0.2$, $E_H=0.425$, and $\alpha=0.5$. We plot the expected revenues under different weighting factors in Fig. 2. We highlight the expected revenues from standard, efficient, and optimally weighted UPC auctions.

As Fig. 2 illustrates, neither the standard UPC auction (analogous to Yahoo!'s) nor the efficient UPC auction (analogous to Google's) is optimal. The efficient UPC auction assigns a lower weighting factor to L-types than the optimal one. The standard UPC auction assigns a higher weighting factor than the optimal one. In this particular case, the standard UPC auction and the efficient UPC auction endure eleven percent and four percent revenue loss, respectively. Given the size of keyword search markets (for Google, the revenue is US\$1.38 billion of the second quarter, 2005), the financial gains for Yahoo! and Google to shift to optimally weighted UPC auction can be significant.

Given the popularity of standard and efficient UPC auctions in practice, it is also tempting for us to rank the two formats in terms of revenues. However, the following numerical example shows that the expected revenues from a standard UPC auction and an efficient UPC auction cannot be unanimously ranked. This and other numerical examples we compute also reveal an interesting pattern: the efficient UPC auction tends to per-

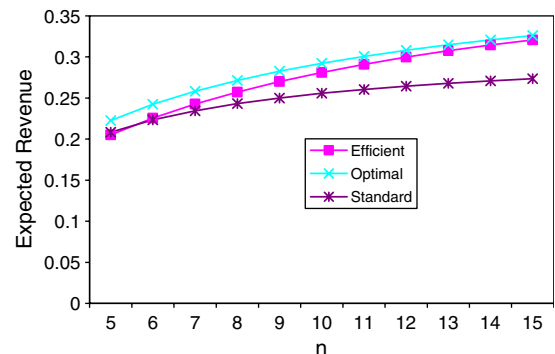


Fig. 3. Expected revenues under different n .

⁷ This condition requires that as γ increases, high-valuation H-type advertisers lose their probability of winning faster than low-valuation H-type advertisers.

form better than the standard UPC auction when the number of bidders is large while the opposite is true when the number of bidders is small (Fig. 3). It also appears that revenue generated by the efficient UPC auction approaches that of the optimally weighted auction as the number of bidders increases while the standard UPC auction does not. Besides the number of bidders, the composition of L-type and H-type bidders can also affect the revenue ranking of the two designs.

6. Conclusion

We studied the issue of exploiting past information on one dimension of bidders' valuation in the keyword advertising auction context. Our study was motivated by increasingly available information on bidders in recurring online auctions. In doing so, we studied a class of weighted UPC auctions that encompass popular auction formats including those adopted by Yahoo! and Google. Although UPC auctions may not be the theoretically optimal form, they are still interesting due to their ease of implementation and practical relevance.

One of our main findings was that efficient UPC auctions, in which unit-price bids are weighted by expected CTRs, can achieve the first-best (*ex ante*) efficient allocation. This makes weighted UPC auctions an attractive mechanism since as we have mentioned in the introduction, weighted UPC auctions can reduce bidders' risks. We also showed that auctioneers can achieve higher revenues by using appropriate weighting factors based on past performance information. The revenue-maximizing weighting factor assigned to disadvantageous bidders should be higher than is suggested by the efficient resource-allocation criteria. The intuition for this result is familiar in economics literature: by favoring disadvantaged players, intermediaries can reduce the economic rent of the advantageous players, the benefit of which can more than compensate the loss caused by misallocation. Although the above results are derived in the keyword advertising setting, they can be generalized to other online settings where auctioneers can observe bidders' past performances.

We applied our model framework to study the keyword auction designs of Yahoo! and Google, leaders of two main camps in keyword advertising. In terms of resource allocation efficiency, Google's approach (weighted UPC auctions) dominates Yahoo!'s. However, in terms of the ability to generate expected revenues, neither company's design can dominate another. Numerical results suggest that Google's ap-

proach may generate higher revenue when the number of bidders is large.

Our analysis generates two implications for the keyword advertising industry. First, firms that are concerned about assigning advertisement slots to those who value them the most should weigh advertisers' unit-price bids by estimates of their future click-through rates. Weighted UPC auctions provide an effective framework to achieve such resource allocation efficiency. Of course, the more information on advertisers' past performance and the better estimation procedure, the higher allocation efficiency advertising intermediaries can achieve.

Second, firms that are concerned about total revenue should bias more toward low-CTR advertisers than suggested by efficient UPC auctions. This can be especially useful when the playing field is uneven and there are only a few bidders. In practice, an optimal UPC auction can be implemented by specifying the allocation rule (weighting factors), and can also be approximated, for example, by employing a CTR estimating system that biases toward low CTRs. However, one potential drawback of tilting the field toward low-CTR advertisers is that, in the long term, it may select advertisers with low CTRs, causing the pool of advertisers to deteriorate.

Our model may be expanded in a few directions in the future. First, it will be interesting to look at cases where bidders are able to manipulate their signals. Second, we have assumed that advertisers do not know others' past performance, which permits a symmetric equilibrium bidding strategy. It would be interesting to ask whether our results can be carried over to the case where bidders know others' past performance. When there are only two bidders, the question reduces to an asymmetric game such as analyzed in [10]. However, the generalization from the two-bidder case to the many-bidder case is nontrivial. For instance, we can anticipate that, unlike the two-bidder case, kinks may emerge in the many-bidder case where a bidder faces unequal numbers of competitors when their valuation is low and when their valuation is high. We speculate that our equilibrium may be considered as the "average" of all equilibriums, each of which corresponds to a realized asymmetric setting. Third, in this paper we have assumed all bidders have past performance records. We may also relax this assumption by allowing new entrants who have not established their performance records. This brings out a practical issue of how to level the field between those with past performance information and those without.

Our model framework may offer a starting point for studying *click-through spam*, which has been an increasing threat to the keyword advertising industry. Click-through spam occurs when individuals or parties maliciously click on advertising links in which they have no interest. Click-through spam can quickly drain advertisers' budgets without generating any returns to them. Both standard and weighted UPC auctions are vulnerable to click-through spam.⁸ This novel problem posts some interesting challenges: What is the impact of click-through spam on various designs of pay-per-click keyword auctions? Can we reduce or eliminate the negative impact of click-through spam through appropriate auction designs? If yes, how should intermediaries choose their auction design while taking into consideration click-through spam? These issues need to be addressed in future research.

Acknowledgements

We thank Thomas Wiseman at the University of Texas at Austin and Juan Feng at the University of Florida for their insightful comments on this paper. We also thank participants in the University of Texas at Austin research seminars, in the Mark C. Burger Applied Microeconomic Workshop at the University of Kentucky, and in the 2004 Workshop on Information Systems and Economics (WISE) for their helpful feedback.

Appendix A. Proofs

A.1. Proof of Lemma 1

We first consider the case $\gamma \leq 1$. By (4) and (5),

$$\begin{aligned} U_H(\gamma v, \gamma b) &= E_H \gamma (v - b) [\alpha F(\phi_H(\gamma b)) \\ &\quad + (1 - \alpha) F(\phi_L(b))]^{n-1} \\ &= \frac{\gamma E_H}{E_L} U_L(v, b), \forall v \in [0, 1] \end{aligned} \quad (A1)$$

$$\begin{aligned} \beta_L(v) &= \arg\max_b \{U_L(v, b)\} = \arg\max_b \left\{ \frac{E_L}{\gamma E_H} U_H(\gamma v, \gamma b) \right\} \\ &= \frac{1}{\gamma} \arg\max_{b'} \{U_H(\gamma v, b')\} = \frac{1}{\gamma} \beta_H(\gamma v), \forall v \in [0, 1] \end{aligned}$$

⁸ Weighted UPC auctions may have limited resistance to click-through spam targeted on particular advertisers: the spam causes one's CTR to increase, which may bring the advertiser some advantage in the next period.

where the first and the last step is by definition of bidding functions. When $\gamma > 1$, we can similarly have

$$U_H(v, b) = \gamma \frac{E_H}{E_L} U_L(v/\gamma, b/\gamma), \forall v \in [0, 1]. \quad (A2)$$

The rest is analogous. \square

A.2. Proof of Proposition 1

Denote $V_L(v) \equiv U_L(v, \beta_L(v))$ and $V_H(v) \equiv U_H(v, \beta_H(v))$ as equilibrium payoffs of L-type and H-type advertisers with valuation-per-click v .

$$\begin{aligned} V_L(v) &= U_L(v, \beta_L(v)) \\ &= E_L(v - \beta_L(v)) [\alpha F(\phi_H(\gamma \beta_L(v))) \\ &\quad + (1 - \alpha) F(\phi_L(\beta_L(v)))]^{n-1} \\ &= E_L(v - \beta_L(v)) [\alpha F(\gamma \phi_L(\beta_L(v))) \\ &\quad + (1 - \alpha) F(\phi_L(\beta_L(v)))]^{n-1} \\ &= E_L(v - \beta_L(v)) [\alpha F(\gamma v) + (1 - \alpha) F(v)]^{n-1} \end{aligned} \quad (A3)$$

Where the third step follows from (9), if $\gamma \leq 1$, and from (10), if $\gamma > 1$ and $\gamma \beta_L(v) \leq \bar{b}_H$. If $\gamma > 1$ and $\gamma \beta_L(v) > \bar{b}_H$ (which imply $v > 1/\gamma$, according to (8)), the third step results from $F(\phi_H(\gamma \beta_L(v))) = F(\phi_H(\bar{b}_H)) = F(1) = F(\gamma v) = F(\gamma \phi_L(\beta_L(v)))$.

$$\frac{dV_L(v)}{dv} = \frac{\partial U_L(v, \beta_L(v))}{\partial v} + \frac{\partial U_L(v, \beta_L(v))}{\partial b} \frac{d\beta_L(v)}{dv} \quad (A4)$$

According to the first order condition, $\partial U_L(v, \beta_L(v)) / \partial b = 0$. So,

$$\begin{aligned} \frac{dV_L(v)}{dv} &= \frac{\partial U_L(v, \beta_L(v))}{\partial v} \\ &= E_L [\alpha F(\phi_H(\gamma \beta_L(v))) + (1 - \alpha) F(\phi_L(\beta_L(v)))]^{n-1} \\ &= E_L [\alpha F(\gamma v) + (1 - \alpha) F(v)]^{n-1}. \end{aligned} \quad (A5)$$

The differential equation (A5) can be solved explicitly. Moving dv to the right hand side, integrating both sides from 0 to v , and applying the boundary condition $V_L(0) = 0$, we get

$$V_L(v) = E_L \int_0^v [\alpha F(\gamma t) + (1 - \alpha) F(t)]^{n-1} dt, \text{ for } v \in [0, 1]. \quad (A6)$$

Combining (A6) and (A3), we can solve L-type's equilibrium bidding function as

$$\beta_L(v) = v - \frac{\int_0^v [\alpha F(\gamma t) + (1 - \alpha)F(t)]^{n-1} dt}{[\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1}}, \text{ for } v \in [0, 1]. \quad (\text{A7})$$

Following similar steps, we can solve the equilibrium payoff function and the bidding function for H-types:

$$V_H(v) = E_H \int_0^v [\alpha F(t) + (1 - \alpha)F(t/\gamma)]^{n-1} dt, \text{ for } v \in [0, 1], \quad (\text{A8})$$

$$\beta_H(v) = v - \frac{\int_0^v [\alpha F(t) + (1 - \alpha)F(t/\gamma)]^{n-1} dt}{[\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1}}, \quad (\text{A9})$$

for $v \in [0, 1]$.

In the following, we denote

$$\rho_L(v) \equiv [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1} \text{ and} \quad (\text{A10})$$

$$\rho_H(v) \equiv [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1} \quad (\text{A11})$$

as the equilibrium winning probabilities for L-types and H-types, respectively. It is clear that both $\rho_L(v)$ and $\rho_H(v)$ are strictly increasing in v for $v \in [0, 1]$.

Now we show that $\beta_L(v)$ is indeed monotonically increasing. By integration by parts,

$$\beta_L(v) = v - \frac{\int_0^v \rho_L(t) dt}{\rho_L(v)} = \frac{\int_0^v \rho'_L(t) t dt}{\rho_L(v)}. \quad (\text{A12})$$

$$\begin{aligned} \frac{d\beta_L(v)}{dv} &= \frac{\rho'_L(v) \left[\rho_L(v)v - \int_0^v \rho'_L(t) t dt \right]}{\rho_L(v)^2} \\ &= \frac{\rho'_L(v) \int_0^v \rho'_L(t) (v - t) dt}{\rho_L(v)^2} > 0. \end{aligned} \quad (\text{A13})$$

The proof for $\beta_H(v)$ is analogous.

By now we have showed that $\beta_L(v)$ satisfies the first order necessary condition. In the following, we show $\beta_L(v)$ is indeed optimal. We examine the payoff of an L-type bidder with valuation v who bids $b' \neq \beta_L(v)$, when every other L-type bids according to $\beta_L(v)$ and every other H-type bids according to $\beta_H(v)$. Since it is never optimal for an L-type bidder to bid more than $\{\bar{b}_L, \bar{b}_H/\gamma\}$, we limit ourselves to the case $b' \leq \max\{\bar{b}_L, \bar{b}_H/\gamma\}$.

If $\gamma > 1$, we know from Lemma 1 that $\bar{b}_H = \beta_H(1) = \gamma\beta_L(1/\gamma) < \gamma\beta_L(1) = \gamma\bar{b}_L$. Thus there exists $v' \in [0, 1]$, such that $b' = \beta_L(v')$.

$$U_L(v, \beta_L(v)) - U_L(v, \beta_L(v'))$$

$$\begin{aligned} &= E_L \int_0^v \rho_L(t) dt - E_L \left[v - \left(v' - \frac{\int_0^{v'} \rho_L(t) dt}{\rho_L(v')} \right) \right] \rho_L(v') \\ &= E_L \int_0^v \rho_L(t) dt - E_L \left[v \rho_L(v') - v' \rho_L(v') + \int_0^{v'} \rho_L(t) dt \right] \\ &= E_L \left[\rho_L(v')(v' - v) + \int_{v'}^v \rho_L(t) dt \right] \\ &= E_L \int_{v'}^v [\rho_L(t) - \rho_L(v')] dt. \end{aligned}$$

Since $\rho_L(t)$ is strictly increasing, $U_L(v, \beta_L(v)) - U_L(v, \beta_L(v')) > 0$ for both $v' < v$ and $v' > v$.

If $\gamma < 1$, we know from Lemma 1 that $\bar{b}_L = \beta_L(1) = \beta_H(\gamma)/\gamma < \beta_H(1)/\gamma = \bar{b}_H/\gamma$. Thus there exists $v' \in [0, 1]$ such that $b' = \beta_H(v')/\gamma$. In addition, the L-type advertiser wins with a probability of $\rho_H(v')$.

$$U_L(v, \beta_L(v)) - U_L(v, \beta_H(v')/\gamma)$$

$$\begin{aligned} &= E_L \int_0^v \rho_L(t) dt - E_L \left[v - \frac{1}{\gamma} \left(v' - \frac{\int_0^{v'} \rho_H(t) dt}{\rho_H(v')} \right) \right] \rho_H(v') \\ &= \frac{E_L}{\gamma} \left[\int_0^{v\gamma} \rho_H(t) dt + \rho_H(v')(v' - v\gamma) - \int_0^{v'} \rho_H(t) dt \right] \\ &= \frac{E_L}{\gamma} \int_{v'}^{v\gamma} [\rho_H(t) - \rho_H(v')] dt \end{aligned}$$

where step 2 is due to $\rho_H(\gamma v) = \rho_L(v)$, $v \in [0, 1]$. Since $\rho_H(v)$ is strictly increasing, $U_L(v, \beta_L(v)) - U_L(v, \beta_H(v')/\gamma) > 0$ holds for both $v' < \gamma v$ and $v' > \gamma v$. By the same logic, we can show $\beta_H(v)$ is indeed optimal. \square

A.3. Proof of Proposition 2

Let $M_L(v)$ and $M_H(v)$ denote the expected payments from L-type and H-type bidders, respectively, with valuation v . Because the expected payment from a bidder is equal to its total expected valuation upon winning minus its expected payoff,

$$M_L(v) = E_L v \rho_L(v) - U_L(v) = E_L \left[v \rho_L(v) - \int_0^v \rho_L(t) dt \right], \quad (\text{A14})$$

$$M_H(v) = E_H v \rho_H(v) - U_H(v) = E_H \left[v \rho_H(v) - \int_0^v \rho_H(t) dt \right]. \quad (\text{A15})$$

The expected payment from one bidder is,

$$\begin{aligned}
& (1 - \alpha)E[M_L(v)] + \alpha E[M_H(v)] \\
&= (1 - \alpha)E_L \int_0^1 \left[v\rho_L(v) - \int_0^v \rho_L(t)dt \right] f(v)dv \\
&\quad + \alpha E_H \int_0^1 \left[v\rho_H(v) - \int_0^v \rho_H(t)dt \right] f(v)dv \\
&= (1 - \alpha)E_L \int_0^1 [v\rho_L(v)f(v) - \rho_L(v)(1 - F(v))]dv \\
&\quad + \alpha E_H \int_0^1 [v\rho_H(v)f(v) - \rho_H(v)(1 - F(v))]dv \\
&= (1 - \alpha)E_L \int_0^1 \rho_L(v) \left(v - \frac{1 - F(v)}{f(v)} \right) f(v)dv \\
&\quad + \alpha E_H \int_0^1 \rho_H(v) \left(v - \frac{1 - F(v)}{f(v)} \right) f(v)dv.
\end{aligned}$$

The total expected revenue from all bidders is n times the above. \square

A.4. Proof of Proposition 3

Since bidding functions for L-type and H-type are both increasing, the competition among bidders with the same type always ends with the highest-valuation bidder winning the slot. In inter-type competition, an L-type bidder with valuation-per-click v ties with an H-type bidder with valuation-per-click γv (if $\gamma \leq 1$), according to Lemma 1. Their total valuations for the advertising slot are vE_L and γvE_H , respectively. The ranking mechanism will be efficient as long as $vE_L = \gamma vE_H$, or $\gamma = E_L/E_H$.

Now we show an efficient UPC auction and a standard first-price auction where bidders have a valuation of $E_H v$ with probability α and $E_L v$ with probability $(1 - \alpha)$, and v is distributed according to F . Let \tilde{b}_L and \tilde{b}_H denote bids (random variables) from L-types and H-types, respectively, in the efficient UPC auction. The payoff function of an L-type bidder is:

$$\begin{aligned}
U_L(v, b) &= E_L(v - b) \left[\alpha Pr\left(\tilde{b}_H < \frac{E_L}{E_H} b\right) \right. \\
&\quad \left. + (1 - \alpha) Pr(\tilde{b}_L < b) \right]^{n-1} \\
&= (E_L v - E_L b) [\alpha Pr(E_H \tilde{b}_H < E_L b) \\
&\quad + (1 - \alpha) Pr(E_L \tilde{b}_L < E_L b)]^{n-1}. \quad (A16)
\end{aligned}$$

The above payoff function can also be regarded as the payoff function for a bidder who has a valuation of $E_L v$ and bids $E_L b$ in a standard first-price auction. Similarly, the payoff function of an H-type bidder

$$\begin{aligned}
U_H(v, b) &= (E_H v - E_H b) [\alpha Pr(E_H \tilde{b}_H < E_H b) \\
&\quad + (1 - \alpha) Pr(E_L \tilde{b}_L < E_H b)]^{n-1} \quad (A17)
\end{aligned}$$

can also be regarded as the payoff function a bidder who has a valuation of $E_H v$ and bids $E_H b$ in the standard first-price auction. Thus we can infer bidders' total payment in the standard auction is exactly the same as in the efficient UPC auction. In other words, they generate the same expected revenue to the intermediary. \square

A.5. Proof of Proposition 4

(a) Taking the first order derivative of the expected revenue (12) with respect to γ yields

$$\begin{aligned}
& (n - 1)(1 - \alpha)\alpha E_L \int_0^1 [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} \\
&\quad \times [vf(v) - (1 - F(v))]f(\gamma v)v dv - (n - 1)(1 - \alpha) \\
&\quad \times \alpha E_H \frac{1}{\gamma^2} \int_0^1 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} \\
&\quad \times [vf(v) - (1 - F(v))]f(v/\gamma)v dv. \quad (A18)
\end{aligned}$$

We only need to check the sign of the above for $0 \leq \gamma \leq E_L/E_H$. To do so, eliminate the common positive terms and change the dummy variable of the second term (notice that for $\gamma < 1$, the integrand in the second term is zero for $v \in [\gamma, 1]$):

$$\begin{aligned}
& E_L \int_0^1 [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} [vf(v) - (1 - F(v))] \\
&\quad \times f(\gamma v)v dv - E_H \int_0^1 [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} \\
&\quad \times [\gamma vf(\gamma v) - (1 - F(\gamma v))]f(v)v dv. \quad (A19)
\end{aligned}$$

Denoting the above as $G(\gamma)$ and reorganizing terms,

$$\begin{aligned}
G(\gamma) &= \int_0^1 [\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} \\
&\quad \times \{E_L[vf(v) - (1 - F(v))]f(\gamma v) - E_H[\gamma vf(\gamma v) \\
&\quad - (1 - F(\gamma v))]f(v)\}v dv. \quad (A20)
\end{aligned}$$

Notice the term in curly brackets,

$$\begin{aligned}
 & E_L[vf(v) - (1 - F(v))]f(\gamma v) - E_H[v\gamma f(\gamma v) \\
 & \quad - (1 - F(\gamma v))]f(v) \\
 & = E_Lvf(v)f(\gamma v) - E_L(1 - F(v))f(\gamma v) \\
 & \quad - E_H\gamma vf(\gamma v)f(v) + E_H(1 - F(\gamma v))f(v) \\
 & = (E_L - E_H\gamma)vf(v)f(\gamma v) + f(\gamma v)f(v) \\
 & \quad \times \left[E_H \frac{1 - F(\gamma v)}{f(\gamma v)} - E_L \frac{1 - F(v)}{f(v)} \right] \\
 & = f(v)f(\gamma v) \left\{ (E_L - E_H\gamma)v + \left[E_H \frac{1 - F(\gamma v)}{f(\gamma v)} \right. \right. \\
 & \quad \left. \left. - E_L \frac{1 - F(v)}{f(v)} \right] \right\}.
 \end{aligned}$$

When $\gamma \leq E_L/E_H$, the first term in the curly brackets is non-negative. By the IHR property,

$$E_H \frac{1 - F(\gamma v)}{f(\gamma v)} - E_L \frac{1 - F(v)}{f(v)} > 0. \quad (A21)$$

So, $G(\gamma) > 0$ for any $\gamma \leq E_L/E_H$, which implies $\gamma^{**} > E_L/E_H$.

(b) A sufficient condition for $\gamma^{**} < 1$ is $G(\gamma) < 0$ for any $\gamma \geq 1$.

When $\gamma > 1$, change the dummy variable for $G(\gamma)$ above and reorganize it as:

$$\begin{aligned}
 G(\gamma) & = \int_0^1 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} \\
 & \quad \times \left\{ E_L \left[v/\gamma - \frac{1 - F(v/\gamma)}{f(v/\gamma)} \right] - E_H \left[v - \frac{1 - F(v)}{f(v)} \right] \right\} \\
 & \quad \times f(v)f(v/\gamma) \frac{v}{\gamma} \frac{1}{\gamma} dv. \quad (A22)
 \end{aligned}$$

When the IHR property is satisfied, we have:

$$\begin{aligned}
 & E_L \left[v/\gamma - \frac{1 - F(v/\gamma)}{f(v/\gamma)} \right] - E_H \left[v - \frac{1 - F(v)}{f(v)} \right] < (E_L - E_H) \\
 & \quad \times \left[v - \frac{1 - F(v)}{f(v)} \right]. \quad (A23)
 \end{aligned}$$

So,

$$\begin{aligned}
 G(\gamma) & < \int_0^1 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} (E_L - E_H) \\
 & \quad \times \left[v - \frac{1 - F(v)}{f(v)} \right] f(v)f(\gamma v) \frac{v}{\gamma} \frac{1}{\gamma} dv \\
 & = -(E_H - E_L) \int_0^1 [vf(v) - (1 - F(v))] \\
 & \quad \times [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} f(v/\gamma) \frac{v}{\gamma} \frac{1}{\gamma} dv. \quad (A24)
 \end{aligned}$$

Notice that $\int_0^1 [vf(v) - (1 - F(v))]dv = 0$ and $vf(v) - (1 - F(v))$ crosses zero once: first negative, then pos-

itive. Assume the crossing point is at v_0 . If $-[\alpha F(v) + (1 - \alpha)F(v/r)]^{n-2} f(v/r) \frac{v}{\gamma^2}$ is an decreasing function of v , or

$$\partial^2 [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-1} / \partial \gamma \partial v \leq 0, \quad (A25)$$

we can have

$$\begin{aligned}
 & \int_0^1 [vf(v) - (1 - F(v))] [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} \\
 & \quad \times f(v/\gamma) \frac{v}{\gamma} \frac{1}{\gamma} dv \\
 & = \int_0^{v_0} [vf(v) - (1 - F(v))] [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} \\
 & \quad \times f(v/\gamma) \frac{v}{\gamma} \frac{1}{\gamma} dv + \int_{v_0}^1 [vf(v) - (1 - F(v))] \\
 & \quad \times [\alpha F(v) + (1 - \alpha)F(v/\gamma)]^{n-2} f(v/\gamma) \frac{v}{\gamma} \frac{1}{\gamma} dv \\
 & > [\alpha F(v_0) + (1 - \alpha)F(v_0/\gamma)]^{n-2} f(v_0/\gamma) \\
 & \quad \times \frac{v_0}{\gamma} \frac{1}{\gamma} \int_0^1 [vf(v) - (1 - F(v))] dv = 0. \quad (A26)
 \end{aligned}$$

Hence IHR and (A25) are sufficient conditions for $\gamma^{**} < 1$. \square

A.6. Proof of Corollary 1

It can be verified that $\gamma^{**} < 1$. Substituting the general distribution with an uniform one in (A20),

$$\begin{aligned}
 G(\gamma) & = \int_0^1 [\alpha \gamma + (1 - \alpha)]^{n-2} v^{n-2} \\
 & \quad \times \{E_L[2v - 1] - E_H[2\gamma v - 1]\} v dv = 0. \quad (A27)
 \end{aligned}$$

After integrating,

$$\begin{aligned}
 G(\gamma) & = [\alpha \gamma + (1 - \alpha)]^{n-2} \\
 & \quad \times \left(\frac{2E_L}{n+1} - \frac{E_L}{n} - \frac{2\gamma E_H}{n+1} + \frac{E_H}{n} \right) = 0, \quad (A28)
 \end{aligned}$$

implying $\gamma^{**} = \frac{(n-1)E_L + (n+1)E_H}{2nE_H}$. \square

A.7. The marginal impact of changing γ on winning probabilities

The marginal impact of changing γ on the winning probability of an L-type bidder with valuation v is

$$\begin{aligned}
 & \frac{\partial}{\partial \gamma} \left([\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1} \right) \\
 & = \alpha(n-1)[\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} f(\gamma v)v,
 \end{aligned}$$

and on the winning probability of an H-type bidder with valuation γv is

$$\begin{aligned} \frac{\partial}{\partial \gamma} \left([\alpha F(x) + (1 - \alpha)F(x/\gamma)]^{n-1} \right) \Big|_{x=\gamma v} \\ = -(1 - \alpha)(n - 1)[\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-2} f(v) \frac{v}{\gamma}. \end{aligned}$$

The number of L-type bidders on a small valuation segment $[v, v + dv]$ is $(1 - \alpha)f(v)dv$. H-type bidders who bid the same score as the former are on the valuation segment $[\gamma v, \gamma v + \gamma dv]$, with an expected number of $\alpha f(\gamma v)\gamma dv$. We can easily verify that the marginal impact of changing γ on the aggregate winning probability of the two segments:

$$\begin{aligned} \frac{\partial}{\partial \gamma} \left([\alpha F(\gamma v) + (1 - \alpha)F(v)]^{n-1} \right) (1 - \alpha)f(v)dv \\ + \frac{\partial}{\partial \gamma} \left([\alpha F(x) + (1 - \alpha)F(x/\gamma)]^{n-1} \right) \Big|_{x=\gamma v} \alpha f(\gamma v)\gamma dv = 0. \end{aligned}$$

In other words, a marginal increase in γ leads to a redistribution of winning probabilities from H-type bidders to their L-type counterpart who bid the same score.

References

- [1] J. Asker, E. Cantillon, Properties of scoring auctions, CEPR Discussion Paper No. 4734, 2004.
- [2] S. Baker, The online ad surge: brand advertising online has taken off—and it's shaking up Madison Ave, Business Week (November 22, 2004).
- [3] F. Branco, The design of multidimensional auctions, RAND Journal of Economics 28 (1) (1997) 63–81.
- [4] J. Bulow, J. Roberts, The simple economics of optimal auctions, Journal of Political Economy 97 (5) (1989) 1060–1090.
- [5] Y. Che, Design competition through multidimensional auctions, RAND Journal of Economics 24 (4) (1993) 668–680.
- [6] C. Ewerhart, K. Fieseler, Procurement auctions and unit-price contracts, RAND Journal of Economics 34 (3) (2003) 569–581.
- [7] J. Feng, Optimal allocation mechanisms when bidders ranking for the objects is common, Working paper, University of Florida, 2004.
- [8] J. Feng, H.K. Bhargava, D.M. Pennock, Implementing sponsored search in web search engines: computational evaluation of alternative mechanisms, forthcoming in INFORMS Journal on Computing (2006).
- [9] Y. Hu, Performance-based pricing models in online advertising, Working paper, MIT, 2004.
- [10] E. Maskin, J. Riley, Asymmetric auction, Review of Economic Studies 67 (3) (2000) 413–438.
- [11] P.R. Milgrom, R.J. Weber, A theory of auctions and competitive bidding, Econometrica 50 (5) (1982) 1089–1122.
- [12] J.G. Riley, Ex post information in auctions, Review of Economic Studies 55 (3) (1988) 409–429.
- [13] T.A. Weber, Z. Zheng, Paid referrals and search intermediary design, Working paper, Stanford University, 2004.