Design Online Auctions with Past Performance Information*

De Liu
University of Kentucky

Jianqing Chen
Univeristy of Texas at Austin

Andrew B. Whinston
University of Texas at Austin

September 10, 2004

1 Introduction

In many resource allocation settings, the potential contenders for resources are recurring players. For example, in government procurement auctions, the pool of candidates may not change much over an extended period of time (think of the Boeing and Lockheed Martin in competing US military contracts). In keyword advertising business, the same group of online retailers may bid for a keyword repeatedly. Oftentimes, the principal can accumulate information about the recurring contenders by observing their past performance. This is especially true in online settings where data collections can be automated. There is considerable interest to understand whether such information is useful and, more importantly, how such information can be used, in improving the efficiency of resource allocation and/or increasing the revenue of the principal.

In this paper, we focus on studying recurring allocation of rights that involves the use of unit-price contracts (UPCs). In an UPC auction, each bidder agrees to pay a unit price for the yield from the resource (or in the procurement case, ask for a unit price for consumed inputs). Examples include highway contracting, government and private procurement auctions, timber auctions, publishing rights for books, and more recently, keyword advertising on search engines.

A central question studied in this paper is whether the principal should use the available historical information in allocating of rights. We investigate this question in a setting where agents not only differ in their marginal valuation of the output but also in the yield from the resource, which is determined by the inherent characteristics of the agents.

One example is the auctioning of keyword advertising slots. An advertiser's total payoff from occupying a keyword advertising slot is affected by two factors: the number of customers who click on the link and the expected profit from each customer. Under the UPC mechanism, each advertiser announces an unit price she is willing to pay for each click, which is sometimes called "cost per click" or CPC. Upon winning an advertisement slot, each advertiser pays according to her announced CPC and the actual number of clicks. While advertisers differ in their willingness to pay for each click, they may also differ in the click through rate (CTR), the number of clicks for a given number of page impressions.

*We thank Thomas Wiseman from University of Texas at Austin for insightful inputs to this research. This material is for reviewing only. Please keep it confidential. Corresponding email:de.liu@uky.edu

Different CTRs could arise for a few reasons. First, the CTR reflects the relevance of the advertisement to the search traffic. Second, the CTR is affected by the presentation of the advertisement. Third, the CTR may reflect the brand name effect. From the search engine's point of view, selecting advertisers with high CTR is also important because it represents the yield rate of using the advertising spots.

Although historical CTRs are readily observable by search engines, the major search engines (Google, Overture and MSN) have adopted different approaches to CTRs. Overture and MSN uses standard UPC auctions for allocating advertising spots whereas Google allocates advertising spots according to the CPC multiplied by corresponding historical CTR.

Besides the observable historical performance, our model setting has two other features. First, neither agents or the principal possess full information about the yield rate – both have to learn from observing the past. Other than the online advertising example, the symmetric uncertainty may occur in some special procurement settings, such as procurement contracts involving high-tech development. Second, the yield rate of an agent follows a markov process. This assumption both provides a rationale for learning from the past and asserts that there is no constant winners – so the principal has to keep picking good performers across time.

We compare three different UPC mechanisms, the standard UPC auction, the UPC auctions adjusted by last period yield (*adjusted UPC*), and the UPC auction adjusted by expected yield (*adjusted UPC II*). We are able to identify the optimal adjustment score, which turns out to be none of the above. The optimal adjustment UPC is more in favor of low-yield agents than adjusted UPC II, while we show that the latter mechanism achieves the most efficient resource allocation.

UPC auctions in a procurement setting has been studied by many authors. Ewerhard and Fieseler (2004) found that although UPC auctions lead to inefficient auctions ex post, it is however desirable for a payment-minimizing principal, because UPC auctions exploit "ex-post observable information which is not used in standard auction formats such as the first-price auction". Manelli and Vincent (1995) studied optimal procurement mechanisms in a general setting. They argued that auctions are not optimal when low cost sellers also provide low quality service. But they consider one period procurement and therefore there is no ex ante information available about sellers.

There are also a few papers which studies search engine resource allocation strategies but with very different emphasis. Hu (2004) studied the pricing of online advertising in principal-agent context. His main concern is the risk allocation and incentives for both parties to invest non-tractable efforts. He didn't model the competition among agents which characterizes at least the setting of keyword advertising. Feng and Bhargava (2004) studies several issues arising in search engine advertising, including the different strategies of Google and Overture, using numerical simulations. They reached a seemingly similar conclusion but based on vastly different arguments: they consider the visitor traffic as positive affected by the overall relevance of advertisements therefore giving Google's approach a nature advantage over Overture since Google tends to select more relevant advertisers. Feng (2003) also studied allocation mechanisms under a setting where advertisers have a consistent ranking of advertising positions but different rate of decrease in absolute valuation.

2 Model Setup

There is a pool of n risk-neutral advertisers. Each advertiser's valuation V of an advertising spot is a function of two factors, her click through rate (CTR) r and her valuation per click V. We assume $V = vr$.

We assume each advertiser's CTR follows a markov process with two status, a high CTR r_H or a low CTR r_L. An advertiser's period $t+1$ CTR only depends on her period t CTR by a transition probability $P_{ij}, i, j \in \{H, L\}$, the probability of moving from i to j. We focus on a stable markov process with (expected) proportion H being α. The stability condition requires $P_{LH} = \frac{1-P_{HH}}{1-\alpha}\alpha$. We denote the $E\left[r_{t+1}|r_t = r_H\right] = E_H$ and $E\left[r_{t+1}|r_t = r_L\right] = E_L$. Assume $E_L < E_H$.

We assume an advertiser's valuation per click in each period is a random draw from a common distribution $F(v), v \in [0, 1]$. This means that the valuation per click is a transient characteristic. We assume $F(v)$ is differentiable and its density function, denoted as $f(v)$ is positive everywhere on its support.

Each advertiser learns her v but not her r for the coming period before bidding. She observes her last period CTR, and so does the intermediary. An advertiser does not observe other advertisers' last period status. Each advertiser holds the same belief that a competitor is H in the last period with probability α.

The intermediary runs a first-price UPC auction in each period. But it may also take advantage of the observed historical CTR by treating the bids of H-type and L-type differently. We assume the intermediary adjustment only affects the allocation of the advertising slot and the winner still has to pay her original bid multiplied by her realized CTR.

3 Model Analysis

Let $\beta_L(v)$ and $\beta_H(v)$ be the bidding functions for H-type and L-type respectively. Let $\phi_L(\cdot)$ and $\phi_H(\cdot)$ be their reverse bidding function. Let $\bar{b}_L$ and $\bar{b}_H$ denote the upper bound of L-type and H-type's bids. Let $\bar{B}_L = \gamma\bar{b}_L$ and $\bar{B}_H = \bar{b}_H$. We can safely assume $\bar{B}_L \leq \bar{B}_H$ and $\gamma \leq 1$[1].

Define $H_L(b)$ and $H_H(b)$ as bids' distribution of H-type and L-type respectively. Assuming monotonic increasing bidding functions (which can also be checked later), so $H_L(b) = F(\phi_L(b))$ and $H_H(b) = F(\phi_H(b))$, and the advertisers' utility functions can be written as:

$$U_L(v,b) = E_L(v-b)\left[\alpha H_H(\gamma b)) + (1-\alpha)H_L(b))\right]^{n-1} \quad (1)$$

$$U_H(v,b) = E_H(v-b)\left[\alpha H_H(b)) + (1-\alpha)H_L(\min(\frac{1}{\gamma}b, \overline{B}_L))\right]^{n-1} \quad (2)$$

Proposition 1 [2] *The equilibrium bidding functions are characterized by*

1. $\forall v \in [0,1]$, $\beta_H(\gamma v) = \gamma\beta_L(v)$

2. $\forall v \in [0,1]$,

$$\phi_L'(b) = \frac{[\alpha H_H(\gamma b))+(1-\alpha)H_L(b))]^{n-1}}{\sum_{i=0}^{n-1}\binom{n-1}{i}\left\{i\frac{\gamma h_H(\gamma b))}{H_H(\gamma b))}+(n-1-i)\frac{h_L(b))}{H_L(b))}\right\}[\alpha H_H(\gamma b))]^i[(1-\alpha)H_L(b))]^{n-1-i}}\frac{1}{(\phi_L(b)-b)}$$

3. $\forall v \in [\gamma, 1]$, $\beta_H(v) = v - \frac{E_H\int_\gamma^v[\alpha F(x)+(1-\alpha)]^{n-1}dx + U_H(\gamma)}{E_H[\alpha F(v)+(1-\alpha)]^{n-1}}$

where $U_H(\gamma)$ is the expected equilibrium utility of a H-type advertiser whose valuation per click is γ.

The above proposition completely characterizes the optimal bidding functions for each type of advertisers. The intuition is a L-type with value v and a H-type with γv will

[1] we can prove that $\bar{B}_L > \bar{B}_H$ is not in the intermediarie's interests.

[2] Due to limited space, all proofs are omitted

face the exactly same competition so that they will place the same "effective" bid, which means $b_H = \gamma b_L$. Because $\gamma \leq 1$, we can expect the H-type who has valuation $v \geq \gamma$ will outbid any L-type and faces only competition from other H-type advertisers.

To facilitate understanding, we assume v is uniformly distributed on the interval $[0, 1]$ and the reservation price is 0. In such a case, the bidding function becomes:

$$\beta_L(v) = \frac{n-1}{n} v \tag{3}$$

$$\beta_H(v) = \begin{cases} \frac{n-1}{n} v & v \in [0, \gamma] \\ v - \frac{E_H \int_\gamma^v [\alpha x + (1-\alpha)]^{n-1} dx + U_H(\gamma)}{E_H [\alpha v + (1-\alpha)]^{n-1}} & v \in [\gamma, 1] \end{cases} \tag{4}$$

We can continue to evaluate the expected revenue of the intermediary as:

$$\begin{aligned} \pi(\gamma) &= (\frac{1-\alpha}{\gamma} + \alpha)^{n-1} (\frac{1-\alpha}{\gamma} \frac{E_L}{\gamma} + \alpha E_H) \int_0^\gamma (n-1) x^n dx \\ &\quad + n\alpha E_H \int_\gamma^1 \left[(2x-1)(\alpha x + 1 - \alpha)^{n-1} - \frac{1}{n} \gamma (\alpha\gamma + 1 - \alpha)^{n-1} \right] dx \end{aligned} \tag{5}$$

We can compare the performance of three stylized mechanisms. The case $\gamma = 1$ corresponds to standard UPC auctions. Overture is one of the adopters. We call the case $\gamma = r_L / r_H$ as *Adjusted UPC*, which has been adopted by Google. This mechanism directly uses historical CTR as multiplying factors. We call the case $\gamma = E_L / E_H$ as *Adjusted UPC II.* Notice that the third case is equivalent to a regular auction which elicits bids on total payment.

The following example illustrates bidding functions in three UPC mechanism. We can notice that L-type advertisers bid exactly the same under three mechanisms while H-type advertisers bid less aggressive under adjusted UPC, and even less agress under Adjusted UPC II.

Example 1 *We use* $n = 10$*,* $r_L = \frac{1}{8}, r_H = \frac{1}{2}, P_{LH} = \frac{1}{5}$*, and* $P_{HH} = \frac{4}{5}$*. The following figures illustrate the bidding functions under different UPC mechanisms.*

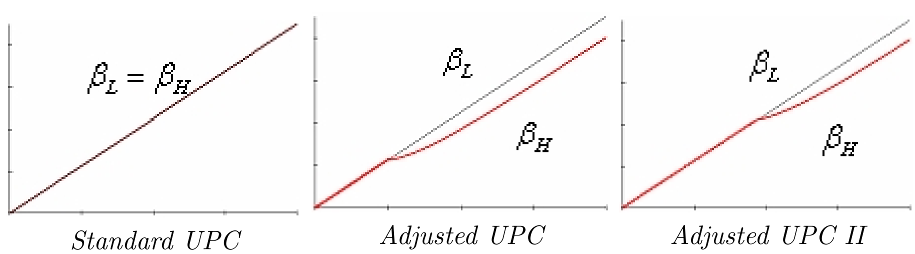

Figure 1: Bidding Functions under Three UPC Mechanisms

The expected profits of three different UPC mechanisms are as following.

$$\pi_{UPC} = 0.25568 < \pi_{AdjustedUPC} = 0.274\,03 < \pi_{AdjustedUPCII} = 0.281\,09$$

The reason that profits are higher under adjusted UPC mechanisms, even though H-type advertisers bid less aggressively, is that H-type advertisers are chosen more often. Now we turn to consider the optimal ranking score. We denote γ^* as the *efficient* ranking score which maximizes the aggregate surplus of advertisers and the intermediary. We denote γ^{**} as the *optimal* ranking score which maximizes the profits of the intermediary.

Proposition 2 *The efficient ranking factor is given by* $\gamma^* = E_L/E_H$. *The optimal ranking factor is given by* $\gamma^{**} = \frac{(n-1)E_L+(n+1)E_H}{2nE_H}$

It is clear from the above proposition that

$$\gamma^* < \gamma^{**} < 1$$

It can also be seen from above that γ^{**} moves further away from γ^* towards 1 as E_L/E_H decreases and as n increases. The intuition of the former is that when E_L/E_H is low, it is more profitable to squeeze high-type for rents. The intuition of the latter is that when n increases, the competition within high-type advertisers will increase, which reduces the need to use L-type advertisers to induce competition. There are two implications for key-word advertising businesses. The Overture's approach is in favor of low CTR advertisers and the Google's approach is in favor of high CTR advertisers. Overture's approach tends to perform better (worse) than Google's approach when there are few (many) advertisers and the difference in CTR is high (low). However, neither approach is optimal.

4 Future Research

This model setup can be expanded in several interesting directions. First, we can study the robustness of different mechanisms to search engine spam, where advertisers receives malicious non-productive clicks. Second, related to the first issue is the moral hazard issue, i.e. what if the advertisers could invest in non-contractible effort to influence her CTR as well as valuation? Another possibility is auction designs with new entrants. Lastly and ultimately, we could pursue a general optimal mechanism problem under such a model setting.

References

[1] Ewerhart, Christian and Fieseler, Karsten (2004). Procurement Auctions and Unit-price Contracts, forthcoming Rand Journal of Economics.

[2] Feng, Juan (2004). Optimal Allocation Mechanisms When Bidders Ranking for the Objects is Common. Working paper, October 2003.

[3] Feng, Juan and Bhargava, Hemant K. etc (2004). Implementing Paid Placement in Web Search Engines: Computational Evaluation of Alternative Mechanisms. Working paper, January 2004.

[4] Hu, Jeffrey (2004).Performance-based Pricing Models in Online Advertising, MIT working paper. March 2004.

[5] Manelli, Alejandro M. and Vincent, Daniel R (1995). Optimal Procurement Mechanisms, Econometirca 63(3): 591-620