An Analysis of Alternative Slot Auction Designs for Sponsored Search

Sébastien Lahaie*
Division of Engineering and Applied Sciences
Harvard University, Cambridge, MA 02138
slahaie@eecs.harvard.edu

ABSTRACT

Billions of dollars are spent each year on sponsored search, a form of advertising where merchants pay for placement alongside web search results. Slots for ad listings are allocated via an auction-style mechanism where the higher a merchant bids, the more likely his ad is to appear above other ads on the page. In this paper we analyze the incentive, efficiency, and revenue properties of two slot auction designs: "rank by bid" (RBB) and "rank by revenue" (RBR), which correspond to stylized versions of the mechanisms currently used by Yahoo! and Google, respectively. We also consider first- and second-price payment rules together with each of these allocation rules, as both have been used historically. We consider both the "short-run" incomplete information setting and the "long-run" complete information setting. With incomplete information, neither RBB nor RBR are truthful with either first or second pricing. We find that the informational requirements of RBB are much weaker than those of RBR, but that RBR is efficient whereas RBB is not. We also show that no revenue ranking of RBB and RBR is possible given an arbitrary distribution over bidder values and relevance. With complete information, we find that no equilibrium exists with first pricing using either RBB or RBR. We show that there typically exists a multitude of equilibria with second pricing, and we bound the divergence of (economic) value in such equilibria from the value obtained assuming all merchants bid truthfully.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

*Part of this work was done while the author was at Yahoo! Research.

Keywords

sponsored search, search engines, slot allocation, auction theory, rank by bid, rank by revenue

1. INTRODUCTION

Today, Internet giants Google and Yahoo! boast a combined market capitalization of over $150 billion, largely on the strength of *sponsored search*, the fastest growing component of a resurgent online advertising industry. PricewaterhouseCoopers estimates that 2004 industry-wide sponsored search revenues were $3.9 billion, or 40% of total Internet advertising revenues.[1] Industry watchers expect 2005 revenues to reach or exceed $7 billion.[2] Roughly 80% of Google's estimated $4 billion in 2005 revenue and roughly 45% of Yahoo!'s estimated $3.7 billion in 2005 revenue will likely be attributable to sponsored search.[3] A number of other companies—including LookSmart, FindWhat, InterActiveCorp (Ask Jeeves), and eBay (Shopping.com)—earn hundreds of millions of dollars of sponsored search revenue annually.

Sponsored search is a form of advertising where merchants pay to appear alongside web search results. For example, when a user searches for "used honda accord san diego" in a web search engine, a variety of commercial entities (San Diego car dealers, Honda Corp, automobile information portals, classified ad aggregators, eBay, etc...) may bid to to have their listings featured alongside the standard "algorithmic" search listings. Advertisers bid for placement on the page in an auction-style format where the higher they bid the more likely their listing will appear above other ads on the page. By convention, sponsored search advertisers generally pay *per click*, meaning that they pay only when a user clicks on their ad, and do not pay if their ad is displayed but not clicked. Though many people claim to systematically ignore sponsored search ads, Majestic Research reports that

[1] `www.iab.net/resources/adrevenue/pdf/IAB_PwC_2004full.pdf`
[2] `battellemedia.com/archives/002032.php`
[3] These are rough "back of the envelope" estimates. Google and Yahoo! 2005 revenue estimates were obtained from Yahoo! Finance. We assumed $7 billion in 2005 industry-wide sponsored search revenues. We used Nielsen/NetRatings estimates of search engine market share in the US, the most monetized market:
`wired-vig.wired.com/news/technology/0,1282,69291,00.html`
Using comScore's international search engine market share estimates would yield different estimates:
`www.comscore.com/press/release.asp?press=622`

as many as 17% of Google searches result in a paid click, and that Google earns roughly nine cents on average for every search query they process.[4]

Usually, sponsored search results appear in a separate section of the page designated as "sponsored" above or to the right of the algorithmic results. Sponsored search results are displayed in a format similar to algorithmic results: as a list of items each containing a title, a text description, and a hyperlink to a corresponding web page. We call each position in the list a *slot*. Generally, advertisements that appear in a higher ranked slot (higher on the page) garner more attention and more clicks from users. Thus, all else being equal, merchants generally prefer higher ranked slots to lower ranked slots.

Merchants bid for placement next to particular search queries; for example, Orbitz and Travelocity may bid for "las vegas hotel" while Dell and HP bid for "laptop computer". As mentioned, bids are expressed as a maximum willingness to pay per click. For example, a forty-cent bid by HostRocket for "web hosting" means HostRocket is willing to pay up to forty cents every time a user clicks on their ad.[5] The auctioneer (the search engine[6]) evaluates the bids and allocates slots to advertisers. In principle, the allocation decision can be altered with each new incoming search query, so in effect new auctions clear continuously over time as search queries arrive.

Many allocation rules are plausible. In this paper, we investigate two allocation rules, roughly corresponding to the two allocation rules used by Yahoo! and Google. The "rank by bid" (RBB) allocation assigns slots in order of bids, with higher ranked slots going to higher bidders. The "rank by revenue" (RBR) allocation assigns slots in order of the product of bid times expected relevance, where relevance is the proportion of users that click on the merchant's ad after viewing it. In our model, we assume that an ad's expected relevance is known to the auctioneer and the advertiser (but not necessarily to other advertisers), and that clickthrough rate decays monotonically with lower ranked slots. In practice, the expected clickthrough rate depends on a number of factors, including the position on the page, the ad text (which in turn depends on the identity of the bidder), the nature and intent of the user, and the context of other ads and algorithmic results on the page, and must be learned over time by both the auctioneer and the bidder [13]. As of this writing, to a rough first-order approximation, Yahoo! employs a RBB allocation and Google employs a RBR allocation, though numerous caveats apply in both cases when it comes to the vagaries of real-world implementations.[7]

Even when examining a one-shot version of a slot auction, the mechanism differs from a standard multi-item auction in subtle ways. First, a single bid per merchant is used to allocate multiple non-identical slots. Second, the bid is communicated not as a direct preference over slots, but as a preference for clicks that depend stochastically on slot allocation.

We investigate a number of economic properties of RBB and RBR slot auctions. We consider the "short-run" incomplete information case in Section 3, adapting and extending standard analyses of single-item auctions. In Section 4 we turn to the "long-run" complete information case; our characterization results here draw on techniques from linear programming. Throughout, important observations are highlighted as claims supported by examples. Our contributions are as follows:

- We show that with multiple slots, bidders do not reveal their true values with either RBB or RBR, and with either first- or second-pricing.
- With incomplete information, we find that the informational requirements of playing the equilibrium bid are much weaker for RBB than for RBR, because bidders need not know any information about each others' relevance (or even their own) with RBB.
- With incomplete information, we prove that RBR is efficient but that RBB is not.
- We show via a simple example that no general revenue ranking of RBB and RBR is possible.
- We prove that in a complete-information setting, first-price slot auctions have no pure strategy Nash equilibrium, but that there always exists a pure-strategy equilibrium with second pricing.
- We provide a constant-factor bound on the deviation from efficiency that can occur in the equilibrium of a second-price slot auction.

In Section 2 we specify our model of bidders and the various slot auction formats.

In Section 3.1 we study the incentive properties of each format, asking in which cases agents would bid truthfully. There is possible confusion here because the "second-price" design for slot auctions is reminiscent of the Vickrey auction for a single item; we note that for slot auctions the Vickrey mechanism is in fact very different from the second-price mechanism, and so they have different incentive properties.[8]

In Section 3.2 we derive the Bayes-Nash equilibrium bids for the various auction formats. This is useful for the efficiency and revenue results in later sections. It should become clear in this section that slot auctions in our model are a straightforward generalization of single-item auctions. Sections 3.3 and 3.4 address questions of efficiency and revenue under incomplete information, respectively.

In Section 4.1 we determine whether pure-strategy equilibria exist for the various auction formats, under complete information. In Section 4.2 we derive bounds on the deviation from efficiency in the pure-strategy equilibria of second-price slot auctions.

Our approach is positive rather than normative. We aim to clarify the incentive, efficiency, and revenue properties of two slot auction designs currently in use, under settings of

[4]battellemedia.com/archives/001102.php

[5]Usually advertisers also set daily or monthly budget caps; in this paper we do not model budget constraints.

[6]In the sponsored search industry, the auctioneer and search engine are not always the same entity. For example Google runs the sponsored search ads for AOL web search, with revenue being shared. Similarly, Yahoo! currently runs the sponsored search ads for MSN web search, though Microsoft will begin independent operations soon.

[7]Here are two among many exceptions to the Yahoo! = RBB and Google = RBR assertion: (1) Yahoo! excludes ads deemed insufficiently relevant either by a human editor or due to poor historical click rate; (2) Google sets differing reserve prices depending on Google's estimate of ad quality.

[8]Other authors have also made this observation [5, 6].

incomplete and complete information. We do not attempt to derive the "optimal" mechanism for a slot auction.

Related work. Feng et al. [7] compare the revenue performance of various ranking mechanisms for slot auctions in a model with incomplete information, much as we do in Section 3.4, but they obtain their results via simulations whereas we perform an equilibrium analysis.

Liu and Chen [12] study properties of slot auctions under incomplete information. Their setting is essentially the same as ours, except they restrict their attention to a model with a single slot and a binary type for bidder relevance (high or low). They find that RBR is efficient, but that no general revenue ranking of RBB and RBR is possible, which agrees with our results. They also take a design approach and show how the auctioneer should assign relevance scores to optimize its revenue.

Edelman et al. [6] model the slot auction problem both as a static game of complete information and a dynamic game of incomplete information. They study the "locally envy-free equilibria" of the static game of complete information; this is a solution concept motivated by certain bidding behaviors that arise due to the presence of budget constraints. They do not view slot auctions as static games of incomplete information as we do, but do study them as dynamic games of incomplete information and derive results on the uniqueness and revenue properties of the resulting equilibria. They also provide a nice description of the evolution of the market for sponsored search.

Varian [18] also studies slot auctions under a setting of complete information. He focuses on "symmetric" equilibria, which are a refinement of Nash equilibria appropriate for slot auctions. He provides bounds on the revenue obtained in equilibrium. He also gives bounds that can be used to infer bidder values given their bids, and performs some empirical analysis using these results. In contrast, we focus instead on efficiency and provide bounds on the deviation from efficiency in complete-information equilibria.

2. PRELIMINARIES

We focus on a slot auction for a single keyword. In a setting of incomplete information, a bidder knows only distributions over others' private information (value per click and relevance). With complete information, a bidder knows others' private information, and so does not need to rely on distributions to strategize. We first describe the model for the case with incomplete information, and drop the distributional information from the model when we come to the complete-information case in Section 4.

2.1 The Model

There is a fixed number K of slots to be allocated among N bidders. We assume without loss of generality that $K \leq N$, since superfluous slots can remain blank. Bidder i assigns a *value* of X_i to each click received on its advertisement, regardless of this advertisement's rank.[9] The probability that i's advertisement will be clicked if viewed is $A_i \in [0,1]$. We refer to A_i as bidder i's *relevance*. We refer to $R_i = A_i X_i$ as bidder i's *revenue*. The X_i, A_i, and R_i are random variables and we denote their realizations by x_i, α_i, and r_i respectively. The probability that an advertisement will be viewed if placed in slot j is $\gamma_j \in [0,1]$. We assume $\gamma_1 > \gamma_2 > \ldots > \gamma_K$. Hence bidder i's advertisement will have a *clickthrough rate* of $\gamma_j \alpha_i$ if placed in slot j. Of course, an advertisement does not receive any clicks if it is not allocated a slot.

Each bidder's value and relevance pair (X_i, A_i) is independently and identically distributed on $[0,\bar{x}] \times [0,1]$ according to a continuous density function f that has full support on its domain. The density f and slot probabilities $\gamma_1, \ldots, \gamma_K$ are common knowledge. Only bidder i knows the realization x_i of its value per click X_i. Both bidder i and the seller know the realization α_i of A_i, but this realization remains unobservable to the other bidders.

We assume that bidders have quasi-linear utility functions. That is, the expected utility to bidder i of obtaining the slot of rank j at a price of b per click is

$$u_i(j,b) = \gamma_j \alpha_i (x_i - b)$$

If the advertising firms bidding in the slot auction are risk-neutral and have ample liquidity, quasi-linearity is a reasonable assumption.

The assumptions of independence, symmetry, and risk-neutrality made above are all quite standard in single-item auction theory [11, 19]. The assumption that clickthrough rate decays monotonically with lower slots—by the same factors for each agent—is unique to the slot auction problem. We view it as a main contribution of our work to show that this assumption allows for tractable analysis of the slot auction problem using standard tools from single-item auction theory. It also allows for interesting results in the complete information case. A common model of decaying clickthrough rate is the exponential decay model, where $\gamma_k = \frac{1}{\delta^{k-1}}$ with decay $\delta > 1$. Feng et al. [7] state that their actual clickthrough data is fitted extremely well by an exponential decay model with $\delta = 1.428$.

Our model lacks budget constraints, which are an important feature of real slot auctions. With budget constraints keyword auctions cannot be considered independently of one another, because the budget must be allocated across multiple keywords—a single advertiser typically bids on multiple keywords relevant to his business. Introducing this element into the model is an important next step for future work.[10]

2.2 Auction Formats

In a slot auction a bidder provides to the seller a declared value per click $\tilde{x}_i(x_i, \alpha_i)$ which depends on his true value and relevance. We often denote this declared value (bid) by $\tilde{x}_i$ for short. Since a bidder's relevance α_i is observable to the seller, the bidder cannot misrepresent it. We denote the k^{th} highest of the N declared values by $\tilde{x}^{(k)}$, and the k^{th} highest of the N declared revenues by $\tilde{r}^{(k)}$, where the declared revenue of bidder i is $\tilde{r}_i = \alpha_i \tilde{x}_i$. We consider two types of allocation rules, "rank by bid" (RBB) and "rank by revenue" (RBR):

[9] Indeed Kitts et al. [10] find that in their sample of actual click data, the correlation between rank and conversion rate is not statistically significant. However, for the purposes of our model it is also important that bidders *believe* that conversion rate does not vary with rank.

[10] Models with budget constraints have begun to appear in this research area. Abrams [1] and Borgs et al. [3] design multi-unit auctions for budget-constrained bidders, which can be interpreted as slot auctions, with a focus on revenue optimization and truthfulness. Mehta et al. [14] address the problem of matching user queries to budget-constrained advertisers so as to maximize revenue.

RBB. Slot k goes to bidder i if and only if $\tilde{x}_i = \tilde{x}^{(k)}$.

RBR. Slot k goes to bidder i if and only if $\tilde{r}_i = \tilde{r}^{(k)}$.

We will commonly represent an allocation by a one-to-one function $\sigma : [K] \to [N]$, where $[n]$ is the set of integers $\{1, 2, \ldots, n\}$. Hence slot k goes to bidder $\sigma(k)$.

We also consider two different types of payment rules. Note that no matter what the payment rule, a bidder that is not allocated a slot will pay 0 since his listing cannot receive any clicks.

First-price. The bidder allocated slot k, namely $\sigma(k)$, pays $\tilde{x}_{\sigma(k)}$ per click under both the RBB and RBR allocation rules.

Second-price. If $k < N$, bidder $\sigma(k)$ pays $\tilde{x}_{\sigma(k+1)}$ per click under the RBB rule, and pays $\tilde{r}_{\sigma(k+1)}/\alpha_{\sigma(k)}$ per click under the RBR rule. If $k = N$, bidder $\sigma(k)$ pays 0 per click.[11]

Intuitively, a second-price payment rule sets a bidder's payment to the lowest bid it could have declared while maintaining the same ranking, given the allocation rule used.

Overture introduced the first slot auction design in 1997, using a first-price RBB scheme. Google then followed in 2000 with a second-price RBR scheme. In 2002, Overture (at this point acquired by Yahoo!) then switched to second pricing but still allocates using RBB. One possible reason for the switch is given in Section 4.

We assume that ties are broken as follows in the event that two agents make the exact same bid or declare the same revenue. There is a permutation of the agents $\kappa : [N] \to [N]$ that is fixed beforehand. If the bids of agents i and j are tied, then agent i obtains a higher slot if and only if $\kappa(i) < \kappa(j)$. This is consistent with the practice in real slot auctions where ties are broken by the bidders' order of arrival.

3. INCOMPLETE INFORMATION

3.1 Incentives

It should be clear that with a first-price payment rule, truthful bidding is neither a dominant strategy nor an *ex post* Nash equilibrium using either RBB or RBR, because this guarantees a payoff of 0. There is always an incentive to shade true values with first pricing.

The second-price payment rule is reminiscent of the second-price (Vickrey) auction used for selling a single item, and in a Vickrey auction it is a dominant strategy for a bidder to reveal his true value for the item [19]. However, using a second-price rule in a slot auction together with either allocation rule above does not yield an incentive-compatible mechanism, either in dominant strategies or *ex post* Nash equilibrium.[12] With a second-price rule there is no incentive for a bidder to bid higher than his true value per click using either RBB or RBR: this either leads to no change in the outcome, or a situation in which he will have to pay more than his value per click for each click received, resulting in a negative payoff.[13] However, with either allocation rule there may be an incentive to shade true values with second pricing.

Claim 1. *With second pricing and $K \geq 2$, truthful bidding is not a dominant strategy nor an* ex post *Nash equilibrium for either RBB or RBR.*

Example. There are two agents and two slots. The agents have relevance $\alpha_1 = \alpha_2 = 1$, whereas $\gamma_1 = 1$ and $\gamma_2 = 1/2$. Agent 1 has a value of $x_1 = 6$ per click, and agent 2 has a value of $x_2 = 4$ per click. Let us first consider the RBB rule. Suppose agent 2 bids truthfully. If agent 1 also bids truthfully, he wins the first slot and obtains a payoff of 2. However, if he shades his bid down below 4, he obtains the second slot at a cost of 0 per click yielding a payoff of 3. Since the agents have equal relevance, the exact same situation holds with the RBR rule. Hence truthful bidding is not a dominant strategy in either format, and neither is it an *ex post* Nash equilibrium.

To find payments that make RBB and RBR dominant-strategy incentive-compatible, we can apply Holmstrom's lemma [9] (see also chapter 3 in Milgrom [15]). Under the restriction that a bidder with value 0 per click does not pay anything (even if he obtains a slot, which can occur if there are as many slots as bidders), this lemma implies that there is a unique payment rule that achieves dominant-strategy incentive compatibility for either allocation rule. For RBB, the bidder allocated slot k is charged per click

$$\sum_{i=k+1}^{K} (\gamma_{i-1} - \gamma_i)\tilde{x}^{(i)} + \gamma_K \tilde{x}^{(K+1)} \tag{1}$$

Note that if $K = N$, $\tilde{x}^{(K+1)} = 0$ since there is no $K + 1^{\text{th}}$ bidder. For RBR, the bidder allocated slot k is charged per click

$$\frac{1}{\alpha_{\sigma(k)}} \left(\sum_{i=k+1}^{K} (\gamma_{i-1} - \gamma_i)\tilde{r}^{(i)} + \gamma_K \tilde{r}^{(K+1)} \right) \tag{2}$$

Using payment rule (2) and RBR, the auctioneer is aware of the true revenues of the bidders (since they reveal their values truthfully), and hence ranks them according to their true revenues. We show in Section 3.3 that this allocation is in fact efficient. Since the VCG mechanism is the unique mechanism that is efficient, truthful, and ensures bidders with value 0 pay nothing (by the Green-Laffont theorem [8]), the RBR rule and payment scheme (2) constitute exactly the VCG mechanism.

In the VCG mechanism an agent pays the externality he imposes on others. To understand payment (2) in this sense, note that the first term is the added utility (due to an increased clickthrough rate) agents in slots $k + 1$ to K would receive if they were all to move up a slot; the last term is the utility that the agent with the $K + 1^{\text{st}}$ revenue would receive by obtaining the last slot as opposed to nothing. The leading coefficient simply reduces the agent's expected payment to a payment per click.

[11]We are effectively assuming a reserve price of zero, but in practice search engines charge a non-zero reserve price per click.

[12]Unless of course there is only a single slot available, since this is the single-item case. With a single slot both RBB and RBR with a second-price payment rule are dominant-strategy incentive-compatible.

[13]In a dynamic setting with second pricing, there may be an incentive to bid higher than one's true value in order to exhaust competitors' budgets. This phenomenon is commonly called "bid jamming" or "antisocial bidding" [4].

3.2 Equilibrium Analysis

To understand the efficiency and revenue properties of the various auction formats, we must first understand which rankings of the bidders occur in equilibrium with different allocation and payment rule combinations. The following lemma essentially follows from the Monotonic Selection Theorem by Milgrom and Shannon [16].

LEMMA 1. *In a RBB (RBR) auction with either a first- or second-price payment rule, the symmetric Bayes-Nash equilibrium bid is strictly increasing with* value *(*revenue*).*

As a consequence of this lemma, we find that RBB and RBR auctions allocate the slots greedily by the *true* values and revenues of the agents, respectively (whether using first- or second-price payment rules). This will be relevant in Section 3.3 below. For a first-price payment rule, we can explicitly derive the symmetric Bayes-Nash equilibrium bid functions for RBB and RBR auctions. The purpose of this exercise is to lend qualitative insights into the parameters that influence an agent's bidding, and to derive formulae for the expected revenue in RBB and RBR auctions in order to make a revenue ranking of these two allocation rules (in Section 3.4).

Let $G(y)$ be the expected resulting clickthrough rate, in a symmetric equilibrium of the RBB auction (with either payment rule), to a bidder with value y and relevance $\alpha = 1$. Let $H(y)$ be the analogous quantity for a bidder with revenue y and relevance 1 in a RBR auction. By Lemma 1, a bidder with value y will obtain slot k in a RBB auction if y is the k^{th} highest of the true realized values. The same applies in a RBR auction when y is the k^{th} highest of the true realized revenues. Let $F_X(y)$ be the distribution function for value, and let $F_R(y)$ be the distribution function for revenue. The probability that y is the k^{th} highest out of N values is

$$\binom{N-1}{k-1}(1 - F_X(y))^{k-1} F_X(y)^{N-k}$$

whereas the probability that y is the k^{th} highest out of N revenues is the same formula with F_R replacing F_X. Hence we have

$$G(y) = \sum_{k=1}^{K} \gamma_k \binom{N-1}{k-1}(1 - F_X(y))^{k-1} F_X(y)^{N-k}$$

The H function is analogous to G with F_R replacing F_X. In the two propositions that follow, g and h are the derivatives of G and H respectively. We omit the proof of the next proposition, because it is almost identical to the derivation of the equilibrium bid in the single-item case (see Krishna [11], Proposition 2.2).

PROPOSITION 1. *The symmetric Bayes-Nash equilibrium strategies in a first-price RBB auction are given by*

$$\tilde{x}^B(x, \alpha) = \frac{1}{G(x)} \int_0^x y\, g(y)\, dy$$

The first-price equilibrium above closely parallels the first-price equilibrium in the single-item model. With a single item g is the density of the second highest value among all N agent values, whereas in a slot auction it is a weighted combination of the densities for the second, third, etc. highest values.

Note that the symmetric Bayes-Nash equilibrium bid in a first-price RBB auction does not depend on a bidder's relevance α. To see clearly why, note that a bidder chooses a bid b so as to maximize the objective

$$\alpha G(\tilde{x}^{-1}(b))(x - b)$$

and here α is just a leading constant factor. So dropping it does not change the set of optimal solutions. Hence the equilibrium bid depends only on the value x and function G, and G in turn depends only on the marginal cumulative distribution of value F_X. So really only the latter needs to be common knowledge to the bidders. On the other hand, we will now see that information about relevance is needed for bidders to play the equilibrium in the first-price RBR auction. So the informational requirements for a first-price RBB auction are much weaker than for a first-price RBR auction: in the RBB auction a bidder need not know his own relevance, and need not know any distributional information over others' relevance in order to play the equilibrium.

Again we omit the next proposition's proof since it is so similar to the one above.

PROPOSITION 2. *The symmetric Bayes-Nash equilibrium strategies in a first-price RBR auction are given by*

$$\tilde{x}^R(x, \alpha) = \frac{1}{\alpha H(\alpha x)} \int_0^{\alpha x} y\, h(y)\, dy$$

Here it can be seen that the equilibrium bid is increasing with x, but not necessarily with α. This should not be much of a concern to the auctioneer, however, because in any case the declared revenue in equilibrium is always increasing in the true revenue.

It would be interesting to obtain the equilibrium bids when using a second-price payment rule, but it appears that the resulting differential equations for this case do not have a neat analytical solution. Nonetheless, the same conclusions about the informational requirements of the RBB and RBR rules still hold, as can be seen simply by inspecting the objective function associated with an agent's bidding problem for the second-price case.

3.3 Efficiency

A slot auction is efficient if in equilibrium the sum of the bidders' revenues from their allocated slots is maximized. Using symmetry as our equilibrium selection criterion, we find that the RBB auction is not efficient with either payment rule.

CLAIM 2. *The RBB auction is not efficient with either first or second pricing.*

EXAMPLE. There are two agents and one slot, with $\gamma_1 = 1$. Agent 1 has a value of $x_1 = 6$ per click and relevance $\alpha_1 = 1/2$. Agent 2 has a value of $x_2 = 4$ per click and relevance $\alpha_2 = 1$. By Lemma 1, agents are ranked greedily by value. Hence agent 1 obtains the lone slot, for a total revenue of 3 to the agents. However, it is most efficient to allocate the slot to agent 2, for a total revenue of 4.

Examples with more agents or more slots are simple to construct along the same lines. On the other hand, under our assumptions on how clickthrough rate decreases with lower rank, the RBR auction is efficient with either payment rule.

THEOREM 1. *The RBR auction is efficient with either first- or second-price payments rules.*

PROOF. Since by Lemma 1 the agents' equilibrium bids are increasing functions of their revenues in the RBR auction, slots are allocated greedily according to true revenues. Let σ be a non-greedy allocation. Then there are slots s, t with $s < t$ and $r_{\sigma(s)} < r_{\sigma(t)}$. We can switch the agents in slots s and t to obtain a new allocation, and the difference between the total revenue in this new allocation and the original allocation's total revenue is

$$\begin{aligned} & \left(\gamma_t r_{\sigma(s)} + \gamma_s r_{\sigma(t)}\right) - \left(\gamma_s r_{\sigma(s)} + \gamma_t r_{\sigma(t)}\right) \\ = \quad & \left(\gamma_s - \gamma_t\right)\left(r_{\sigma(t)} - r_{\sigma(s)}\right) \end{aligned}$$

Both parenthesized terms above are positive. Hence the switch has increased the total revenue to the bidders. If we continue to perform such switches, we will eventually reach a greedy allocation of greater revenue than the initial allocation. Since the initial allocation was arbitrary, it follows that a greedy allocation is always efficient, and hence the RBR auction's allocation is efficient. □

Note that the assumption that clickthrough rate decays montonically by the same factors $\gamma_1, \ldots, \gamma_K$ for all agents is crucial to this result. A greedy allocation scheme does not necessarily find an efficient solution if the clickthrough rates are monotonically decreasing in an independent fashion for each agent.

3.4 Revenue

To obtain possible revenue rankings for the different auction formats, we first note that when the allocation rule is fixed to RBB, then using either a first-price, second-price, or truthful payment rule leads to the same expected revenue in a symmetric, increasing Bayes-Nash equilibrium. Because a RBB auction ranks agents by their true values in equilibrium for any of these payment rules (by Lemma 1), it follows that expected revenue is the same for all these payment rules, following arguments that are virtually identical to those used to establish revenue equivalence in the single-item case (see e.g. Proposition 3.1 in Krishna [11]). The same holds for RBR auctions; however, the revenue ranking of the RBB and RBR allocation rules is still unclear. Because of this revenue equivalence principle, we can choose whichever payment rule is most convenient for the purpose of making revenue comparisons.

Using Propositions 1 and 2, it is a simple matter to derive formulae for the expected revenue under both allocation rules. The payment of an agent in a RBB auction is

$$m^B(x, \alpha) = \alpha G(x)\tilde{x}^V(x, \alpha)$$

The expected revenue is then $N \cdot E\left[m^V(X, A)\right]$, where the expectation is taken with respect to the joint density of value and relevance. The expected revenue formula for RBR auctions is entirely analogous using $\tilde{x}^R(x, \alpha)$ and the H function. With these in hand we can obtain revenue rankings for specific numbers of bidders and slots, and specific distributions over values and relevance.

CLAIM 3. *For fixed K, N, and fixed $\gamma_1, \ldots, \gamma_K$, no revenue ranking of RBB and RBR is possible for an arbitrary density f.*

EXAMPLE. Assume there are 2 bidders, 2 slots, and that $\gamma_1 = 1$, $\gamma_2 = 1/2$. Assume that value-relevance pairs are uniformly distributed over $[0, 1] \times [0, 1]$. For such a distribution with a closed-form formula, it is most convenient to use the revenue formulae just derived. RBB dominates RBR in terms of revenue for these parameters. The formula for the expected revenue in a RBB auction yields $1/12$, whereas for RBR auctions we have $7/108$.

Assume instead that with probability $1/2$ an agent's value-relevance pair is $(1, 1/2)$, and that with probability $1/2$ it is $(1/2, 1)$. In this scenario it is more convenient to appeal to formulae (1) and (2). In a truthful auction the second agent will always pay 0. According to (1), in a truthful RBB auction the first agent makes an expected payment of

$$E\left[(\gamma_1 - \gamma_2)A_{\sigma(1)}X_{\sigma(2)}\right] = \frac{1}{2}E\left[A_{\sigma(1)}\right]E\left[X_{\sigma(2)}\right]$$

where we have used the fact that value and relevance are independently distributed for different agents. The expected relevance of the agent with the highest value is $E\left[A_{\sigma(1)}\right] = 5/8$. The expected second highest value is also $E\left[X_{\sigma(2)}\right] = 5/8$. The expected revenue for a RBB auction here is then $25/128$. According to (2), in a truthful RBR auction the first agent makes an expected payment of

$$E\left[(\gamma_1 - \gamma_2)R_{\sigma(2)}\right] = \frac{1}{2}E\left[R_{\sigma(2)}\right]$$

In expectation the second highest revenue is $E\left[R_{\sigma(2)}\right] = 1/2$, so the expected revenue for a RBR auction is $1/4$. Hence in this case the RBR auction yields higher expected revenue.[14][15]

This example suggests the following conjecture: when value and relevance are either uncorrelated or positively correlated, RBB dominates RBR in terms of revenue. When value and relevance are negatively correlated, RBR dominates.

4. COMPLETE INFORMATION

In typical slot auctions such as those run by Yahoo! and Google, bidders can adjust their bids up or down at any time. As Börgers et al. [2] and Edelman et al. [6] have noted, this can be viewed as a continuous-time process in which bidders learn each other's bids. If the process stabilizes the result can then be modeled as a Nash equilibrium in pure strategies of the static one-shot game of complete information, since each bidder will be playing a best-response to the others' bids.[16] This argument seems especially appropriate for Yahoo!'s slot auction design where all bids are

[14] To be entirely rigorous and consistent with our initial assumptions, we should have constructed a continuous probability density with full support over an appropriate domain. Taking the domain to be e.g. $[0, 1] \times [0, 1]$ and a continuous density with full support that is sufficiently concentrated around $(1, 1/2)$ and $(1/2, 1)$, with roughly equal mass around both, would yield the same conclusion.

[15] Claim 3 should serve as a word of caution, because Feng et al. [7] find through their simulations that with a bivariate normal distribution over value-relevance pairs, and with 5 slots, 15 bidders, and $\delta = 2$, RBR dominates RBB in terms of revenue for any level of correlation between value and relevance. However, they assume that bidding behavior in a second-price slot auction can be well approximated by truthful bidding.

[16] We do not claim that bidders will actually learn each others' private information (value and relevance), just that for a stable set of bids there is a corresponding equilibrium of the complete information game.

made public. Google keeps bids private, but experimentation can allow one to discover other bids, especially since second pricing automatically reveals to an agent the bid of the agent ranked directly below him.

4.1 Equilibrium Analysis

In this section we ask whether a pure-strategy Nash equilibrium exists in a RBB or RBR slot auction, with either first or second pricing.

Before dealing with the first-price case there is a technical issue involving ties. In our model we allow bids to be nonnegative real numbers for mathematical convenience, but this can become problematic because there is then no bid that is "just higher" than another. We brush over such issues by assuming that an agent can bid "infinitesimally higher" than another. This is imprecise but allows us to focus on the intuition behind the result that follows. See Reny [17] for a full treatment of such issues.

For the remainder of the paper, we assume that there are as many slots as bidders. The following result shows that there can be no pure-strategy Nash equilibrium with first pricing.[17] Note that the argument holds for both RBB and RBR allocation rules. For RBB, bids should be interpreted as declared values, and for RBR as declared revenues.

Theorem 2. *There exists no complete information Nash equilibrium in pure strategies in the first-price slot auction, for any possible values of the agents, whether using a RBB or RBR allocation rule.*

Proof. Let $\sigma : [K] \to [N]$ be the allocation of slots to the agents resulting from their bids. Let r_i and b_i be the revenue and bid of the agent ranked i^{th}, respectively. Note that we cannot have $b_i > b_{i+1}$, or else the agent in slot i can make a profitable deviation by instead bidding $b_i - \epsilon > b_{i+1}$ for small enough $\epsilon > 0$. This does not change its allocation, but increases its profit. Hence we must have $b_i = b_{i+1}$ (i.e. with one bidder bidding infinitesimally higher than the other). Since this holds for any two consecutive bidders, it follows that in a Nash equilibrium all bidders must be bidding 0 (since the bidder ranked last matches the bid directly below him, which is 0 by default because there is no such bid). But this is impossible: consider the bidder ranked last. The identity of this bidder is always clear given the deterministic tie-breaking rule. This bidder can obtain the top spot and increase his revenue by $(\gamma_1 - \gamma_K) r_K > 0$ by bidding some $\epsilon > 0$, and for small enough ϵ this is necessarily a profitable deviation. Hence there is no Nash equilibrium in pure strategies. □

On the other hand, we find that in a second-price slot auction there can be a multitude of pure strategy Nash equilibria. The next two lemmas give conditions that characterize the allocations that can occur as a result of an equilibrium profile of bids, given fixed agent values and revenues. Then if we can exhibit an allocation that satisfies these conditions, there must exist at least one equilibrium. We first consider the RBR case.

[17] Börgers et al. [2] have proven this result in a model with three bidders and three slots, and we generalize their argument. Edelman et al. [6] also point out this non-existence phenomenon. They only illustrate the fact with an example because the result is quite immediate.

Lemma 2. *Given an allocation σ, there exists a Nash equilibrium profile of bids b leading to σ in a second-price RBR slot auction if and only if*

$$\left(1 - \frac{\gamma_i}{\gamma_{j+1}}\right) r_{\sigma(i)} \le r_{\sigma(j)}$$

for $1 \le j \le N-2$ and $i \ge j+2$.

Proof. There exists a desired vector b which constitutes a Nash equilibrium if and only if the following set of inequalities can be satisfied (the variables are the π_i and b_j):

$$\pi_i \ge \gamma_j (r_{\sigma(i)} - b_j) \quad \forall i, \forall j < i \tag{3}$$

$$\pi_i \ge \gamma_j (r_{\sigma(i)} - b_{j+1}) \quad \forall i, \forall j > i \tag{4}$$

$$\pi_i = \gamma_i (r_{\sigma(i)} - b_{i+1}) \quad \forall i \tag{5}$$

$$b_i \ge b_{i+1} \quad 1 \le i \le N-1 \tag{6}$$

$$\pi_i \ge 0,\ b_i \ge 0 \quad \forall i$$

Here $r_{\sigma(i)}$ is the revenue of the agent allocated slot i, and π_i and b_i may be interpreted as this agent's surplus and declared revenue, respectively. We first argue that constraints (6) can be removed, because the inequalities above can be satisfied if and only if the inequalities without (6) can be satisfied. The necessary direction is immediate. Assume we have a vector (π, b) which satisfies all inequalities above except (6). Then there is some i for which $b_i < b_{i+1}$. Construct a new vector (π, b') identical to the original except with $b'_{i+1} = b_i$. We now have $b'_i = b'_{i+1}$. An agent in slot $k < i$ sees the price of slot i decrease from b_{i+1} to $b'_{i+1} = b_i$, but this does not make i more preferred than k to this agent because we have $\pi_k \ge \gamma_{i-1}(r_{\sigma(k)} - b_i) \ge \gamma_i(r_{\sigma(k)} - b_i) = \gamma_i(r_{\sigma(k)} - b'_{i+1})$ (i.e. because the agent in slot k did not originally prefer slot $i-1$ at price b_i, he will not prefer slot i at price b_i). A similar argument applies for agents in slots $k > i+1$. The agent in slot i sees the price of this slot go down, which only makes it more preferred. Finally, the agent in slot $i+1$ sees no change in the price of any slot, so his slot remains most preferred. Hence inequalities (3)–(5) remain valid at (π, b'). We first make this change to the b_{i+1} where $b_i < b_{i+1}$ and index i is smallest. We then recursively apply the change until we eventually obtain a vector that satisfies all inequalities.

We safely ignore inequalities (6) from now on. By the Farkas lemma, the remaining inequalities can be satisfied if and only if there is no vector z such that

$$\sum_{i,j} (\gamma_j r_{\sigma(i)})\, z_{\sigma(i)j} > 0$$

$$\sum_{i>j} \gamma_j z_{\sigma(i)j} + \sum_{i<j} \gamma_{j-1} z_{\sigma(i)j-1} \le 0 \quad \forall j \tag{7}$$

$$\sum_j z_{\sigma(i)j} \le 0 \quad \forall i \tag{8}$$

$$z_{\sigma(i)j} \ge 0 \quad \forall i, \forall j \ne i$$

$$z_{\sigma(i)i} \text{ free} \quad \forall i$$

Note that a variable of the form $z_{\sigma(i)i}$ appears at most once in a constraint of type (8), so such a variable can never be positive. Also, $z_{\sigma(i)1} = 0$ for all $i \ne 1$ by constraint (7), since such variables never appear with another of the form $z_{\sigma(i)i}$.

Now if we wish to raise $z_{\sigma(i)j}$ above 0 by one unit for $j \ne i$, we must lower $z_{\sigma(i)i}$ by one unit because of the constraint of type (8). Because $\gamma_j r_{\sigma(i)} \le \gamma_i r_{\sigma(i)}$ for $i < j$, raising

$z_{\sigma(i)j}$ with $i < j$ while adjusting other variables to maintain feasibility cannot make the objective $\sum_{i,j}(\gamma_j r_{\sigma(i)})z_{\sigma(i)j}$ positive. If this objective is positive, then this is due to some component $z_{\sigma(i)j}$ with $i > j$ being positive.

Now for the constraints of type (7), if $i > j$ then $z_{\sigma(i)j}$ appears with $z_{\sigma(j-1)j-1}$ (for $1 < j < N$). So to raise the former variable γ_j^{-1} units and maintain feasibility, we must (I) lower $z_{\sigma(i)i}$ by γ_j^{-1} units, and (II) lower $z_{\sigma(j-1)j-1}$ by γ_{j-1}^{-1} units. Hence if the following inequalities hold:

$$r_{\sigma(i)} \leq \left(\frac{\gamma_i}{\gamma_j}\right) r_{\sigma(i)} + r_{\sigma(j-1)} \qquad (9)$$

for $2 \leq j \leq N-1$ and $i > j$, raising some $z_{\sigma(i)j}$ with $i > j$ cannot make the objective positive, and there is no z that satisfies all inequalities above. Conversely, if some inequality (9) does not hold, the objective can be made positive by raising the corresponding $z_{\sigma(i)j}$ and adjusting other variables so that feasibility is just maintained. By a slight reindexing, inequalities (9) yield the statement of the theorem. □

The RBB case is entirely analogous.

LEMMA 3. *Given an allocation σ, there exists a Nash equilibrium profile of bids b leading to σ in a second-price RBB slot auction if and only if*

$$\left(1 - \frac{\gamma_i}{\gamma_{j+1}}\right) x_{\sigma(i)} \leq x_{\sigma(j)}$$

for $1 \leq j \leq N-2$ and $i \geq j+2$.

PROOF SKETCH. The proof technique is the same as in the previous lemma. The desired Nash equilibrium exists if and only if a related set of inequalities can be satisfied; by the Farkas lemma, this occurs if and only if an alternate set of inequalities cannot be satisfied. The conditions that determine whether the latter holds are given in the statement of the lemma. □

The two lemmas above immediately lead to the following result.

THEOREM 3. *There always exists a complete information Nash equilibrium in pure strategies in the second-price RBB slot auction. There always exists an* efficient *complete information Nash equilibrium in pure strategies in the second-price RBR slot auction.*

PROOF. First consider RBB. Suppose agents are ranked according to their true values. Since $x_{\sigma(i)} \leq x_{\sigma(j)}$ for $i > j$, the system of inequalities in Lemma 3 is satisfied, and the allocation is the result of some Nash equilibrium bid profile. By the same type of argument but appealing to Lemma 2 for RBR, there exists a Nash equilibrium bid profile such that bidders are ranked according to their true revenues. By Theorem 1, this latter allocation is efficient. □

This theorem establishes existence but not uniqueness. Indeed we expect that in many cases there will be multiple allocations (and hence equilibria) which satisfy the conditions of Lemmas 2 and 3. In particular, not all equilibria of a second-price RBR auction will be efficient. For instance, according to Lemma 2, with two agents and two slots any allocation can arise in a RBR equilibrium because no constraints apply.

Theorems 2 and 3 taken together provide a possible explanation for Yahoo!'s switch from first to second pricing. We saw in Section 3.1 that this does not induce truthfulness from bidders. With first pricing, there will always be some bidder that feels compelled to adjust his bid. Second pricing is more convenient because an equilibrium can be reached, and this reduces the cost of bid management.

4.2 Efficiency

For a given allocation rule, we call the allocation that would result if the bidders reported their values truthfully the *standard* allocation. Hence in the standard RBB allocation bidders are ranked by true values, and in the standard RBR allocation they are ranked by true revenues. According to Lemmas 2 and 3, a ranking that results from a Nash equilibrium profile can only deviate from the standard allocation by having agents with relatively similar values or revenues switch places. That is, if $r_i > r_j$ then with RBR agent j can be ranked higher than i only if the ratio r_j/r_i is sufficiently large; similarly for RBB. This suggests that the value of an equilibrium allocation cannot differ too much from the value obtained in the standard allocation, and the following theorems confirms this.

For an allocation σ of slots to agents, we denote its total value by $f(\sigma) = \sum_{i=1}^{N} \gamma_i r_{\sigma(i)}$. We denote by $g(\sigma) = \sum_{i=1}^{N} \gamma_i x_{\sigma(i)}$ allocation σ's value when assuming all agents have identical relevance, normalized to 1. Let

$$L = \min_{i=1,\ldots,N-1} \min \left\{ \frac{\gamma_{i+1}}{\gamma_i}, 1 - \frac{\gamma_{i+2}}{\gamma_{i+1}} \right\}$$

(where by default $\gamma_{N+1} = 0$). Let η_x and η_r be the standard allocations when using RBB and RBR, respectively.

THEOREM 4. *For an allocation σ that results from a pure-strategy Nash equilibrium of a second-price RBR slot auction, we have $f(\sigma) \geq L f(\eta_r)$.*

PROOF. We number the agents so that agent i has the i^{th} highest revenue, so $r_1 \geq r_2 \geq \ldots \geq r_N$. Hence the standard allocation has value $f(\eta_r) = \sum_{i=1}^{N} \gamma_i r_i$. To prove the theorem, we will make repeated use of the fact that $\frac{\sum_k a_k}{\sum_k b_k} \geq \min_k \frac{a_k}{b_k}$ when the a_k and b_k are positive. Note that according to Lemma 2, if agent i lies at least two slots below slot j, then $r_{\sigma(j)} \geq r_i \left(1 - \frac{\gamma_{j+2}}{\gamma_{j+1}}\right)$.

It may be the case that for some slot i, we have $\sigma(i) > i$ and for slots $k > i+1$ we have $\sigma(k) > i$. We then say that slot i is *inverted*. Let S be the set of agents with indices at least $i+1$; there are $N-i$ of these. If slot i is inverted, it is occupied by some agent from S. Also all slots strictly lower than $i+1$ must be occupied by the remaining agents from S, since $\sigma(k) > i$ for $k \geq i+2$. The agent in slot $i+1$ must then have an index $\sigma(i+1) \leq i$ (note this means slot $i+1$ cannot be inverted). Now there are two cases. In the first case we have $\sigma(i) = i+1$. Then

$$\begin{aligned} \frac{\gamma_i r_{\sigma(i)} + \gamma_{i+1} r_{\sigma(i+1)}}{\gamma_i r_i + \gamma_{i+1} r_{i+1}} &\geq \frac{\gamma_{i+1} r_i + \gamma_i r_{i+1}}{\gamma_i r_i + \gamma_{i+1} r_{i+1}} \\ &\geq \min \left\{ \frac{\gamma_{i+1}}{\gamma_i}, \frac{\gamma_i}{\gamma_{i+1}} \right\} \\ &= \frac{\gamma_{i+1}}{\gamma_i} \end{aligned}$$

In the second case we have $\sigma(i) > i+1$. Then since all agents in S except the one in slot i lie strictly below slot $i+1$, and the agent in slot i is not agent $i+1$, it must be that agent $i+1$ is in a slot strictly below slot $i+1$. This means that it is at least two slots below the agent that actually occupies slot i, and by Lemma 2 we then have $r_{\sigma(i)} \geq r_{i+1}\left(1 - \frac{\gamma_{i+2}}{\gamma_{i+1}}\right)$. Thus,

$$\begin{aligned} \frac{\gamma_i r_{\sigma(i)} + \gamma_{i+1} r_{\sigma(i+1)}}{\gamma_i r_i + \gamma_{i+1} r_{i+1}} &\geq \frac{\gamma_{i+1} r_i + \gamma_i r_{\sigma(i)}}{\gamma_i r_i + \gamma_{i+1} r_{i+1}} \\ &\geq \min\left\{\frac{\gamma_{i+1}}{\gamma_i}, 1 - \frac{\gamma_{i+2}}{\gamma_{i+1}}\right\} \end{aligned}$$

If slot i is not inverted, then on one hand we may have $\sigma(i) \leq i$, in which case $r_{\sigma(i)}/r_i \geq 1$. On the other hand we may have $\sigma(i) > i$ but there is some agent with index $j \leq i$ that lies at least two slots below slot i. Then by Lemma 2, $r_{\sigma(i)} \geq r_j\left(1 - \frac{\gamma_{i+2}}{\gamma_{i+1}}\right) \geq r_i\left(1 - \frac{\gamma_{i+2}}{\gamma_{i+1}}\right)$.

We write $i \in I$ if slot i is inverted, and $i \notin I$ if neither i nor $i-1$ are inverted. By our arguments above two consecutive slots cannot be inverted, so we can write

$$\begin{aligned} \frac{f(\sigma)}{f(\gamma_r)} &= \frac{\sum_{i\in I}\left(\gamma_i r_{\sigma(i)} + \gamma_{i+1} r_{\sigma(i+1)}\right) + \sum_{i\notin I}\gamma_i r_{\sigma(i)}}{\sum_{i\in I}\left(\gamma_i r_i + \gamma_{i+1} r_{i+1}\right) + \sum_{i\notin I}\gamma_i r_i} \\ &\geq \min\left\{\min_{i\in I}\left\{\frac{\gamma_i r_{\sigma(i)} + \gamma_{i+1} r_{\sigma(i+1)}}{\gamma_i r_i + \gamma_{i+1} r_{i+1}}\right\}, \min_{i\notin I}\left\{\frac{\gamma_i r_{\sigma(i)}}{\gamma_i r_i}\right\}\right\} \\ &\geq L \end{aligned}$$

and this completes the proof. □

Note that for RBR, the standard value is also the efficient value by Theorem 1. Also note that for an exponential decay model, $L = \min\left\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\right\}$. With $\delta = 1.428$ (see Section 2.1), the factor is $L \approx 1/3.34$, so the total value in a pure-strategy Nash equilibrium of a second-price RBR slot auction is always within a factor of 3.34 of the efficient value with such a discount.

Again for RBB we have an analogous result.

THEOREM 5. *For an allocation σ that results from a pure-strategy Nash equilibrium of a second-price RBB slot auction, we have $g(\sigma) \geq Lg(\eta_x)$.*

PROOF SKETCH. Simply substitute bidder values for bidder revenues in the proof of Theorem 4, and appeal to Lemma 3. □

5. CONCLUSIONS

This paper analyzed stylized versions of the slot auction designs currently used by Yahoo! and Google, namely "rank by bid" (RBB) and "rank by revenue" (RBR), respectively. We also considered first and second pricing rules together with each of these allocation rules, since both have been used historically. We first studied the "short-run" setting with incomplete information, corresponding to the case where agents have just approached the mechanism. Our equilibrium analysis revealed that RBB has much weaker informational requirements than RBR, because bidders need not know any information about relevance (even their own) to play the Bayes-Nash equilibrium. However, RBR leads to an efficient allocation in equilibrium, whereas RBB does not. We showed that for an arbitrary distribution over value and relevance, no revenue ranking of RBB and RBR is possible. We hope that the tools we used to establish these results (revenue equivalence, the form of first-price equilibria, the truthful payments rules) will help others wanting to pursue further analyses of slot auctions.

We also studied the "long-run" case where agents have experimented with their bids and each settled on one they find optimal. We argued that a stable set of bids in this setting can be modeled as a pure-strategy Nash equilibrium of the static game of complete information. We showed that no pure-strategy equilibrium exists with either RBB or RBR using first pricing, but that with second pricing there always exists such an equilibrium (in the case of RBR, an *efficient* equilibrium). In general second pricing allows for multiple pure-strategy equilibria, but we showed that the value of such equilibria diverges by only a constant factor from the value obtained if all agents bid truthfully (which in the case of RBR is the efficient value).

6. FUTURE WORK

Introducing budget constraints into the model is a natural next step for future work. The complication here lies in the fact that budgets are often set for entire campaigns rather than single keywords. Assuming that the optimal choice of budget can be made independent of the choice of bid for a specific keyword, it can be shown that it is a dominant-strategy to report this optimal budget with one's bid. The problem is then to ascertain that bids and budgets can indeed be optimized separately, or to find a plausible model where deriving equilibrium bids and budgets together is tractable.

Identifying a condition on the distribution over value and relevance that actually does yield a revenue ranking of RBB and RBR (such as correlation between value and relevance, perhaps) would yield a more satisfactory characterization of their relative revenue properties. Placing bounds on the revenue obtained in a complete information equilibrium is also a relevant question.

Because the incomplete information case is such a close generalization of the most basic single-item auction model, it would be interesting to see which standard results from single-item auction theory (e.g. results with risk-averse bidders, an endogenous number of bidders, asymmetries, etc...) automatically generalize and which do not, to fully understand the structural differences between single-item and slot auctions.

Acknowledgements

David Pennock provided valuable guidance throughout this project. I would also like to thank David Parkes for helpful comments.

7. REFERENCES

[1] Z. Abrams. Revenue maximization when bidders have budgets. In *Proc. the ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[2] T. Börgers, I. Cox, and M. Pesendorfer. Personal Communication.

[3] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proc. the Sixth ACM Conference on Electronic Commerce*, Vancouver, BC, 2005.

[4] F. Brandt and G. Weiß. Antisocial agents and Vickrey auctions. In J.-J. C. Meyer and M. Tambe, editors,

Intelligent Agents VIII, volume 2333 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 2001.

[5] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. In *Workshop on Sponsored Search Auctions, ACM Electronic Commerce*, 2005.

[6] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. NBER working paper 11765, November 2005.

[7] J. Feng, H. K. Bhargava, and D. M. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 2005. Forthcoming.

[8] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[9] B. Holmstrom. Groves schemes on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[10] B. Kitts, P. Laxminarayan, B. LeBlanc, and R. Meech. A formal analysis of search auctions including predictions on click fraud and bidding tactics. In *Workshop on Sponsored Search Auctions, ACM Electronic Commerce*, 2005.

[11] V. Krishna. *Auction Theory*. Academic Press, 2002.

[12] D. Liu and J. Chen. Designing online auctions with past performance information. *Decision Support Systems*, 2005. Forthcoming.

[13] C. Meek, D. M. Chickering, and D. B. Wilson. Stochastic and contingent payment auctions. In *Workshop on Sponsored Search Auctions, ACM Electronic Commerce*, 2005.

[14] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized on-line matching. In *Proc. 46th IEEE Symposium on Foundations of Computer Science*, 2005.

[15] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[16] P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.

[17] P. J. Reny. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67(5):1029–1056, 1999.

[18] H. R. Varian. Position auctions. Working Paper, February 2006.

[19] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.