Asymptotically Optimal Repeated Auctions for Sponsored Search

Nicolas Lambert
Computer Science Department
Stanford University
Stanford, CA 94305
nlambert@stanford.edu

Yoav Shoham
Computer Science Department
Stanford University
Stanford, CA 94305
shoham@ai.stanford.edu

ABSTRACT

We investigate asymptotically optimal keyword auctions, that is, auctions which maximize revenue as the number of bidders grows. We do so under two alternative behavioral assumptions. The first explicitly models the repeated nature of keyword auctions. It introduces a novel assumption on individual bidding, namely that bidders never overbid their value, and bid their actual value if shut out for long enough. Under these conditions we present a broad class of repeated auctions that are asymptotically optimal among all sequential auctions (a superset of repeated auctions). Those auctions have varying payment schemes but share the ranking method. The Google auction belongs to this class, but not the Yahoo auction, and indeed we show that the latter is not asymptotically optimal. (Nonetheless, with some additional distributional assumptions, the Yahoo auction can be shown to belong to a broad category of auctions that are asymptotically optimal among all auction mechanisms that do not rely on ad relevance.) We then look at the one-shot keyword auction, which can be taken to model repeated auctions in which relatively myopic bidders converge on the equilibrium of the full-information stage game. In this case we show that the Google auction remains asymptotically optimal and the Yahoo auction suboptimal. The distributional assumptions under which our theorems hold are quite general. We do however show that the Google auction is not asymptotically revenue-maximizing for general distributions.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and behavioral sciences—*economics*

General Terms

Economics, theory

Keywords

Auctions, sponsored search, keyword auctions

ICEC'07, August 19–22, 2007, Minneapolis, Minnesota, USA.

1. INTRODUCTION

The Internet advertising segment is growing steadily. In the first half of 2006, revenues from Internet advertising reached nearly $8B, a 37% increase over the same period in 2005. In the next few years, the Internet is expected to account for 20% of global advertising spending[1]. The share of sponsored search is by far the most lucrative, and accounts for 40% of total Internet advertising, more than twice as much as any other advertising technique[2].

Auctions have become the main mechanism used by search engines to allocate sponsored links. In sponsored search auctions, advertising firms specify target keywords and bids. The search engine runs an auction every time an user enters a new query, and assigns bidders to auctions depending on the keyword preferences, and possibly some other constraints, such as budget limits or location. The result is a listing of relevant ads displayed in a "sponsored links" section, alongside the main search results. Advertising space is allocated to merchants according to a criterion that depends on their bids and/or ad relevance. Merchants are only charged when their ad is clicked, not whenever their ad is selected for display.

Despite its apparent similarity to the second-price auction, the Google auction is quite complex; in particular, it is *neither truthful nor revenue-maximizing*. Yet Google is doing just fine, generating 40% more revenue that its direct competitor Yahoo[3]. Recently, Google announced that its third quarter profit almost doubled, while Yahoo said that selection of winning bidders would now be based on both bids and an ad quality score, similarly to Google Adwords auction mechanism. Is this an accident? We claim not.

Given the high stakes, the relative paucity of results on optimal keyword auctions is striking. One reason is undoubtedly the complex dependence of the auctions on both the bid amounts and click-through rates – that is, clicks per impression – which makes for complex auctions and even more complex analysis. For example, in [1] Aggarwal et al. design a truthful one-shot auction in dominant strategies, and prove that this is the only truthful (and optimal) auction in dominant strategies for a given weighted-bid ranking mechanism, but the auction is quite complex and not intuitive. For example, for the ranking method used by Google, the price would be a linear combination of bids of other merchants that depend on the expected click frequency in

[1]Source: The Economist, November 25, 2006
[2]Source: Business Wire, September 25, 2006
[3]Source: The New York Times, July 19, 2006

various slots. Optimal truthful one-shot auctions for more general ranking mechanisms have been characterized [5], but both payment schemes and winner selection rules remain quite complex.

Another source of difficulty is the fact that sponsored search auctions are repeated, usually with very high frequency, and most advertising firms use bidding proxies that make use of available information on past performance to decide on the next bid. For example, Edelman and Ostrovski [2] produce evidence of strategic behaviors with empirical data of the Yahoo auction during the period 2002-2003, after its change to second-bid pricing. While the current auction model is more stable than the previous, first-price mechanism, the amount of strategizing by the bidders remains significant. It is well known that analysis of repeated auctions is substantially more involved than that of one-shot auctions.

So how have researchers handled these challenges? To begin with, in spite of what was said in the previous paragraph, research to date has concentrated almost exclusively on one-shot auctions. One justification for this is demonstration that certain more sophisticated but plausible bidding strategies eventually converge to the equilibrium of the full-information stage game (this, for example, is claimed in [3]). This means that rather than adopt a Myerson-style analysis ([7]) of the full repeated game, one makes certain assumptions about bidders' behaviors. Perhaps less "clean" theoretically, this move is arguably more informative about the real world. The question is of course whether the behavioral assumptions are plausible.

We include analysis of the one-shot model and its implied behavioral restriction in the paper, but put more emphasis on an alternative assumption regarding bidders' behaviors, one that refers to the individual bidder rather than the set as a whole (as does the assumption about convergence to equilibrium). Specifically, we explore the assumption that merchants never overbid their value, and start bidding their actual value if they remain excluded for long enough. We find this a plausible restriction on the behavior of somewhat risk-averse bidders, more defensible than the convergence to equilibrium assumption.

This modeling assumption is one modest contribution of this paper. A more substantial modeling contribution is the focus on asymptotic optimality. One aspect of internet auctions that has not been exploited in research so far is the relatively high number of bidders. And so we ask how revenue can be optimized, as the number of bidders grows unboundedly. We call this asymptotic optimality.

The technical results in this paper all refer to asymptotic optimality, and are partitioned according to the two alternative modeling assumptions.

In the context of the new individual conditions on bidding strategies, we present a broad class of repeated auctions that maximize revenue among all sequential auctions as the number of bidders grows, and, at the limit, perform as well as an omniscient seller. Those auctions have varying payment schemes but share the ranking method. The Google auction belongs to this class, but so do many others. The Yahoo auction does not belong to the class, and indeed we show that it is not asymptotically optimal, as any other auction that allocates slots to the merchants with highest bids. We do however show that the Yahoo auction belongs to a broad category of auctions that are asymptotically revenue-maximizing among all sequential auctions with the same ranking method and whose payments do not rely on ad relevance, and, when ad quality and value are independent parameters, we extend this result to all ranking methods that do not use ad relevance.

In the context of the equilibrium assumption, we show that the Google auction remains asymptotically optimal in a strong sense, while the Yahoo auction is suboptimal.

Our theorems hold under certain traditional assumptions on the distributions of bidder parameters that are quite general, and, arguably, hold in most real-world situations of interest. Nonetheless we show that the Google auction is not asymptotically revenue-maximizing for general distributions.

Our paper is organized as follows. In section 2 we review the existing literature on sponsored search. In section 3 we introduce our model and various auction mechanisms. In section 4 we state our optimality and suboptimality results under individual conditions on bidders'dynamics, while in section 5 we treat the case of the equilibrium assumption. Finally in section 6 we show that the Google auction is not necessarily revenue-maximizing for general distributions on bidder parameters.

2. RELATED WORK

Most of the literature on sponsored search focuses on the one-shot auction. Lahaie [6] discusses the equilibria, efficiency and revenue properties of the first-price and second-price versions of keyword auctions, considered as one-shot auction games with either complete or incomplete information. Varian [8] focuses on the second-price static game auction model of complete information and analyses "symmetric" equilibria, a subset of Nash equilibria.

Aggarwal et al. [1] introduce a truthful one-shot auction when advertisers are allocated in decreasing order of their weighted bids, and show its uniqueness for a given ranking method. They demonstrate the existence of an equilibrium for the second-price version of the game of complete information that achieves the same revenue. Iyengar et al. [5] also consider truthful auctions in the static game, with more general ranking methods, and identify the optimal truthful mechanisms.

In [4], Feng et al. propose simulations to estimate the seller's revenue under different auction mechanisms and different model parameters. However they use a static setting and do not attempt to model bidding strategies, but rather focus their analysis on the effect of distributional parameters.

Edelman et al. [3] are among the first to examine bidding dynamics. They consider the second-price version of keyword auctions, both in a static and dynamic setting. They show that for certain bidding strategies, bids may converge to a "locally envy-free" equilibrium of the one-short game, for which the revenue is the same as for the VCG auction.

3. MODEL

We assume that n bidders compete for a fixed number of slots $s > 1$, with $n > s$ (without loss of generality as we are interested in asymptotic properties). The slots are numbered from top to bottom, slot 1 appears at the top of the page, followed by slot 2, etc. Our model is similar to that introduced in [1, 6]. Each bidder i has a *value* v_i, and a *relevance parameter* x_i, indicator of the quality, or relevance, of the advertisement. Both value and relevance may vary depending on the search keyword that is being associated with the advertisement. In this paper we consider one particular auction corresponding to one keyword. The relevance determines the number of clicks per ad impression, called *click-through rate*, which is the probability for an advertisement to receive a click when displayed. Click-through rates depend on the position of the slot, an advertisement that appears near the top is visited more frequently than when it is dragged to the bottom. Similarly to [1, 6], we assume that click-through rates depend on the slot, and are modulated by relevance: given the position parameters $\theta_1 \geq \cdots \geq \theta_s > 0$, the click through rate of merchant i in position j is $c_{i,j} = \theta_j x_i$. [4] shows that actual data fits our model particularly well. By convention, we assume $\theta_i = 0$ for $i > s$.

We take a probabilistic approach and make certain assumptions on the distribution of values and relevance. For each bidder i, the relevance x_i and value v_i are random variables, and the pairs (x_i, v_i) are identically and independently distributed with joint density $f(x, v)$. We note X and V the relevance and click-through rate of an arbitrary bidder. We use the short notation $f(v)$ to denote the marginal density associated with value v. We note $f_{|v}(x) = \frac{f(x,v)}{f(v)}$ the conditional density of relevance x given the value v, and we denote by $F_{|v}(x)$ the corresponding cumulative distribution. We assume f has compact support $S \subset \mathbb{R}^2_+$, and that the product xv when $(x, v) \in S$ is maximized at only one point in S, (X_M, V_M). We also note $\bar{X}$ and $\bar{V}$ the maximum value taken respectively by x and v in S. Finally we impose that S contains no singularities, that is, there is no isolated point[4]. These topological assumptions allow us to be rigorous, and are verified in most situations of interest. They include as a special case the conditions required in [6, 4]. For example, f might be defined on any rectangular domain and be positive on that domain. The case of general distributions will be discussed in section 6. As opposed to the traditional literature on one-shot auctions, we make minimal assumptions about the knowledge of the main actors. In particular, the density function and the click-through rate model with parameters θ_i's may, or not, be known.

3.1 General sponsored search auctions

We model sponsored search auctions as sequential auctions, and consider a sequencial auction associated with a specific keyword. At each period t, n bidders participate in an auction that sells advertisement slots for the given keyword. Winning bidders receive a slot for the length of the period. The auction mechanism may depend on the period, as we do not exclude the case for which the seller attempts to learn the bidder's value from past bidding information, and charges the winning merchants accordingly. We call them *sequential auctions*. In most cases of interest however, the auction is the same for every period, and depends only on the bids for the current period. We then talk about *repeated auctions*, a special case of sequential auctions.

We assume that each merchant's relevance is known by the seller.

At each period $t = 1, 2, \ldots,$

1. Advertisers submit bids.
2. The seller decides on the number of slots for sale, the allocation of slots to bidders, and the payments-per-click for the period.
3. The seller may decide to release information on bidders, bids, slot attribution, ad relevance or clicks per impression.

[4]Any ball centered on a point in S has positive mass.

We note, for a period t, $s_t \leq s$ the number of slots for sale, $\sigma_t(\cdot)$ the allocation function, and $p_t(\cdot)$ the payment function. The allocation function indicates, for an available slot $i \leq s_t$, the bidder $\sigma_t(i)$ to whom the slot is attributed for the period. Payments-per-clicked are determined by the payment function. A bidder assigned a slot i pays $p_t(i)$ for each click received. We assume that the auction never charges above the bid offer. Bidders who are not assigned a slot do not receive any click and are not charged.

We note $b_i^{(t)}$ the bid of bidder i for period t. We do not attempt to model specific bidding strategies, rather, we make general assumptions on bidding behaviors. Therefore the release of certain information to bidders is transparent in our model, and we impose no constraints on communications between the seller and the merchants. The seller may for example decide to release information regarding bids of competitors, to provide individual data on click frequencies to the merchants, or to give statistics on the average popularity of certain slots. In practice search engines often provide real-time information. For instance, Yahoo provides current bids for a given keyword and the popularity of a given query, while Google gives estimates on costs and positions for a bid.

In our setting, bids, allocations and payments for a given period may depend on past bids, ranks and allocations as well as value and ad relevance. Ultimately, each variable of each period is entirely determined by the initial values and relevance. In that sense, bids, ranks, and allocations, are random variables entirely determined by the initial value and relevance for each bidder. In the remaining of this paper, we consider the probability space formed by bidders' values and relevance, with a probability defined by the product of joint density f. To simplify our notations we consider only deterministic auctions and bidding strategies, however the extension of our analysis to randomized mechanisms is straightforward; it suffices to include a sequence of coin flips in the probability space.

3.2 Common auction formats

Several auction models have been proposed. We review the auctions with first- and second-price rules, and the VCG auction.

We distinguish between two allocation methods:

B-ranking Bidders are ranked in order of decreasing *bid* $b_i^{(t)}$.

RB-ranking Bidders are ranked in order of decreasing *relevance* $\times$ *bid* product $x_i b_i^{(t)}$.

In both cases, the s bidders with highest ranks are allocated an advertisement slot: a merchant with the i-th rank gets the i-th slot, when $i \leq s$.

We now consider payment schemes. Each may be used in conjunction with either allocation method. The Google auction can be modeled as a second-price scheme with RB-ranking, while the Yahoo auction uses the same pricing method but with B-ranking.

First-price Each assigned bidder i is charged her actual bid $b_i^{(t)}$.

Second-price Each assigned bidder $r_t(i)$ in position i is charged the minimum amount she would have to bid to retain her position, that is, $b^{(t)}_{r_t(i+1)}$ for the B-ranking, and $\frac{x_{r_t(i+1)} b^{(t)}_{r_t(i+1)}}{x_{r_t(i)}}$ for the RB-ranking.

VCG Each bidder is charged the externality she imposes on other bidders. For a bidder $r_t(i)$ in position i, payment-per-click is

$$p_t(i) = \frac{1}{\theta_i} \sum_{j=i}^{s} (\theta_j - \theta_{j+1}) b^{(t)}_{r_t(j+1)}$$

for the B-ranking, and

$$p_t(i) = \frac{1}{\theta_i x_{r_t(i)}} \sum_{j=i}^{s} (\theta_j - \theta_{j+1}) x_{r_t(j+1)} b^{(t)}_{r_t(j+1)}$$

for the RB-ranking.

The VCG auction is truthful. However, as several authors have remarked, despite the apparent similarity between the second-price rule and the single item second-price auction, the second-price rule is not a truthful mechanism.

3.3 RB(ρ)-ranking, bounded pricing, and regular auctions

B-ranking and RB-ranking methods are two special cases of a broad class we call *RB(ρ)-ranking*. RB(ρ)-ranking ranks the bidders in order of decreasing product $x_i^{\rho}(b_i^{(t)})^{1-\rho}$, in which ρ is a weight parameter, $0 \leq \rho < 1$. An auction that uses a ranking method allocates slots in ranking order, with first slot attributed to the first bidder. By varying the weight ρ, one gives more or less importance to bid or relevance. Setting weight $\rho = 0$ yields the B-ranking, whereas $\rho = \frac{1}{2}$, yields the RB-ranking. When an auction uses a ranking mechanism, we note $r_t(\cdot)$ the ranking function, which is a permutation of $\{1, \ldots, n\}$, and indicates that bidder $r_t(i)$ gets rank i for the period t. Naturally, for a slot $i \leq s$, $\sigma_t(i) = r_t(i)$.

Most payment schemes used in conjunction with a ranking function only make use of the relevance parameter and bids of the top k bidders for the current period, for some fixed k. We call such mechanisms *bounded pricing* schemes. Formally, the payments-per-click are defined by a function, q_i, for each slot i for sale. At period t, the bidder in slot i is charged $p_t(i) = q_i(b^{(t)}_{r_t(1)}, \ldots, b^{(t)}_{r_t(k)}, x_{r_t(1)}, \ldots, x_{r_t(k)})$ for each click received. We require that each function q_i be continuous and that, for all b, x, $q_i(b, \ldots, b, x, \ldots, x) = b$. The first-price, second-price, and truthful pricing rules are instances of bounded pricing schemes.

We call *Regular auctions* the auctions that use a bounded pricing method in conjunction with RB(ρ)-ranking. Most of the results of this paper apply to regular auctions, which include essentially all of the sponsored search auction models encountered in practice or in the literature, and in particular the Yahoo and Google auction models.

3.4 Seller's revenue and bidder's utility

The seller's revenue, or profit, for a certain period t is the sum of the payments for each slot during the period:

$$\Pi_t = \sum_{i=1}^{s_t} \theta_i x_{\sigma_t(i)} p_t(i)$$

We define the seller's revenue $\mathcal{R}$ as the expected per-period revenue, considering the probability space associated with merchants' values and relevance introduced at the beginning of this section. To ensure well-defined quantities, we consider the limit inf, which is the worst-case revenue (in certain situations the per-period revenue may not converge when the number of periods grows).

$$\mathcal{R} = \liminf_{T \to +\infty} E\left[\frac{1}{T}\sum_{t=1}^{T} \Pi_t\right]$$

If there is convergence, $\mathcal{R}$ is simply the limit of the expected per-period profit.

The bidder utility model is less important under our new behavioral assumptions, and we only use it in the equilibrium condition (section 5). For simplicity, we assume for the rest of the paper that bidders' utility functions are quasi-linear with limit average reward. Bidders who do not receive any slot at any given period have a null utility for the period, while the utility of a bidder i who gets slot j and pays p_i be click is $u_i = \theta_j x_i (v_i - p_i)$. However, most results hold without specific assumptions on the utility functions.

An omniscient seller knows the values for each bidder. In such circumstances, she can extract the full surplus by charging the winning bidders their actual value, and allocating the slots in decreasing order of expected payments, proportional to $x_i v_i$.

4. INDIVIDUAL CONDITIONS ON BIDDING STRATEGIES

In this section we assume that bidders' behaviors meet the following conditions:

ASSUMPTION 1. *We assume that bidders do not bid above their true value. We also assume that there exists a fixed number of periods T_0, such that any bidder left unassigned for the last T_0 consecutive periods bids her true value at the next auction.*

We find these assumptions quite reasonable in a competitive environment. In practice, merchants are often asked to bid the "maximum price they are willing to pay". Bidding at most one's actual value avoids putting oneself at risk. Moreover, in the Google and Yahoo auctions, overbidding is a weakly dominated strategy (for the one-shot auction game).

While our assumptions prevent overbidding, they allow merchants to bid under their true value. In general, the lower the bid, the lower the chances to win a slot. If a bidder underbids repeatedly and consequently keeps losing the auction, raising her bid will increase her chances to win a slot at the next round. If she keeps being excluded for many consecutive periods, that process will eventually lead her to bid her true value, maximizing her chances to get assigned with a nonnegative utility. Naturally, if she succeeds, she may decide in the following periods to re-adjust her bid at a lower price, in order to increase her utility, at the risk of losing the auction. Besides, in many auctions (e.g. Google or Yahoo) the price charged to a winning bidder increases with the bids of certain others (who may be excluded), and bidding one's true value may as a side effect raise the cost of a competitor, even if one is still left unassigned. Finally, regarding the Google and Yahoo auctions, our assumptions include as a special case those of [8], where the first excluded bidder is assumed to bid her true value.

For certain rank-based allocation methods, such as RB(ρ)-ranking, the conditions we impose on bidders' behaviors allows to bound the gap between the true value of a bidder and her bid. More precisely, for the RB(ρ)-ranking method, the "declared rank" $x_i^\rho (b_i^{(t)})^{1-\rho}$ of a high-ranked bidder is close to her "true rank" $x_i^\rho v_i^{1-\rho}$, which is also relatively high compared to that of other bidders.

LEMMA 1. *In an auction using RB(ρ)-ranking, for any $k > 0$, there exists $m > 0$ and a critical period T_c, both independent of the number of bidders, such that for $t > T_c$, every bidder with a rank $j \leq k$ has a ranking value $x_i^\rho (b_i^{(t)})^{1-\rho}$ and product value $x_i^\rho v_i^{1-\rho}$ as high as the least of the top m product values $x_i^\rho v_i^{1-\rho}$'s among all bidders.*

PROOF. Let $T_c = m = sT_0 + k$, and consider a period $t > T_c$. Let A be the set of the m bidders with the highest product values $x_i^\rho v_i^{1-\rho}$. In the previous T_0 periods, at most s bidders are allocated at each period, so that at most sT_0 members of A have been assigned, and by assumption at least k bidders of A bid their true value at t. Among these bidders, let i_0 be the bidder who has the lowest product value $x_i^\rho v_i^{1-\rho}$. Any bidder ranked k or better by the auction must have a ranking value $x_i^\rho (b_i^{(t)})^{1-\rho}$ at least as high as the ranking value of i_0, which is $x_{i_0}^\rho v_{i_0}^{1-\rho}$. As bidders do not overbid their true value, any bidder must have a product value $x_i^\rho v_i^{1-\rho}$ at least as high as her ranking value. So any bidder ranked at most k by the auction must have a ranking value and a product value as high as that of i_0, and so as high as the least product value of bidders in A . $\square$

If those merchants are ordered according to RB(ρ)-ranking on relevance and value – instead of relevance and bid –, the ranking value, that is, the product $x_i^\rho v_i^{1-\rho}$ of the top merchants will, on average, tend to a maximum. The result is intuitive and merely due to the fact that the distribution of value/relevance is bounded, hence so is the distribution of the $x_i^\rho v_i^{1-\rho}$'s. This is captured by the following lemma, whose elementary proof is left to the reader.

LEMMA 2. *Let $m > 0$ be a fixed number, and K be a bidder with the m-th highest value of the products $x_i^\rho v_i^{1-\rho}$. Then,*

$$\lim_{n \to +\infty} E[x_K^\rho v_K^{1-\rho}] = \max_{(x,v) \in S} x^\rho v^{1-\rho}$$

and, for any $\epsilon > 0$, and any $\delta > 0$,

$$P\left(\left| x_K^\rho v_K^{1-\rho} - \max_{(x,v)\in S} x^\rho v^{1-\rho} \right| \leq \delta \right) \geq 1 - \epsilon$$

for a number of bidders n large enough.

For the RB-ranking, it means that those merchants with a very high product $x_i^\rho v_i^{1-\rho}$ will have an ad relevance and per-click value close to X_M and V_M respectively. It also implies that the revenue of an omniscient seller, who knows the true value of each bidder, converges to $X_M V_M \sum \theta_i$ as the number of bidders grows.

4.1 RB-ranking optimality

We now show that regular auctions that use RB-ranking maximize revenue asymptotically. The intuition is as follows: when many advertising firms bid for the same keyword, any given firm competing for a slot must bid near its own value, since many other firms with similar characteristics will be excluded at each auction, and are ready to bid as high as needed to get some advertising space with positive utility. Bidders with similar ranks will have similar value, and if the payment-per-click of a bidder is not too different from the bids of her neighbors in the ranking order, it will be close to her own value for a click. By using RB-ranking, the seller will select the winners among the bidders with the highest "*relevance* × *bid*". It that case, it is the same as the highest "*relevance* × *value*", which are proportional to the expected payments. As ranking values tend to be maximized when the competition is fierce, using RB-ranking allows for maximization of expected payments.

We will use the following elementary lemma that we state without proof:

LEMMA 3. *Let $u_{p,t}$ be a random variable, and $T_c > 0$. Suppose $E[u_{p,t}]$ converges to (resp. is asymptotically bounded above/below by) U uniformly*[5] *for all $t > T_c$, then $E[\frac{1}{T}\sum_{t=1}^{T} u_{p,t}]$ converges to (resp. is asymptotically bounded above/below by) U.*

THEOREM 1. *All regular auctions with RB-ranking are asymptotically revenue-maximizing. Furthermore, the revenue $\mathcal{R}_n^{RB}$ of such an auction gets arbitrarily close to the revenue $\mathcal{R}_n^{Omn}$ of an omniscient seller, as the number of bidders n grows:*

$$\lim_{n \to +\infty} \mathcal{R}_n^{RB} = \lim_{n \to +\infty} \mathcal{R}_n^{Omn}$$

PROOF. Consider a period t. Using the notation of section 3, for a regular auction with bounded pricing, the revenue at period t is:

$$\Pi_t = \sum_{i=1}^{s} \theta_i x_{r_t(i)} p_t(i)$$

with $p_t(i) = q_i(b_{r_t(1)}^{(t)}, \ldots, b_{r_t(k)}^{(t)}, x_{r_t(1)}, \ldots, x_{r_t(k)})$ for some k.

We first show that the expected payment of each winning bidder is arbitrarily close to the highest payment: for any $\epsilon > 0$, $|E[x_{r_t(i)} p_t(r_t(i)) - X_M V_M]| \leq \epsilon$, for n large enough.

As q_i is continuous at $(V_M, \ldots, V_M, X_M, \ldots, X_M)$, where it takes the value V_M, there exists $\delta > 0$, that we can choose smaller than ϵ, so that if $|b_1 - V_M| \leq \delta, \ldots, |b_k - V_M| \leq \delta$, $|x_1 - X_M| \leq \delta, \ldots, |x_k - X_M| \leq \delta$, then $|q_i(b_1, \ldots, b_k, x_1, \ldots, x_k) - V_M| \leq \epsilon$.

Let A be the event $\{|b_{r_t(j)}^{(t)} - V_M| \leq \delta, 1 \leq j \leq k\} \cup \{|x_{r_t(j)} - X_M| \leq \delta, 1 \leq j \leq k\}$.

As the maximum of xv, $(x,v) \in S$ is reached at a unique point (X_M, V_M), there exists $\delta > 0$ such that if $|xv - X_M V_M|$ is small enough, say, smaller than δ', then $|x - X_M| \leq \delta$

[5]We say that a sequence $a_{p,q}$ converges to L uniformly in q if for all $\epsilon > 0$, there exists p_0 such that if $p > p_0$, $|a_{q,p} - L| < \epsilon$ for all q.

and $|v - V_M| \leq \delta$. From lemma 1, there exists m such that if K is a bidder with the m-th highest value of the $x_i v_i$'s, then $X_M V_M \geq x_{r_t(i)} b^{(t)}_{r_t(i)} \geq x_K v_K$, for the top k ranks i. Then, by lemma 2, for a certain critical number of bidders N, if $n > N$, for t after some critical period, $P(|x_{r_t(i)} b^{(t)}_{r_t(i)} - X_M V_M| \leq \delta') \geq P(|x_K v_K - X_M V_M| \leq \delta') \geq 1 - \epsilon$. However, if $|x_{r_t(i)} b^{(t)}_{r_t(i)} - X_M V_M| \leq \delta'$, then $|x_{r_t(i)} - X_M| \leq \delta$ and $|b^{(t)}_{r_t(i)} - V_M| \leq \delta$, hence $P(|x_{r_t(i)} - X_M| \leq \delta) \geq 1 - \epsilon$ and $P(|b^{(t)}_{r_t(i)} - V_M| \leq \delta) \geq 1 - \epsilon$. Therefore, with simple arithmetic manipulation, $P(A) \geq 1 - 2k\epsilon$ for n large enough.

We now split the expectation term in two parts[6],

$$E[x_{r_t(i)} p_t(i) - X_M V_M] = E[(x_{r_t(i)} p_t(i) - X_M V_M)\mathbb{1}_A] + E[(x_{r_t(i)} p_t(i) - X_M V_M)\mathbb{1}_{\bar{A}}]$$

For n large enough and all t large enough,

$$|E[(x_{r_t(i)} p_t(i) - X_M V_M)\mathbb{1}_{\bar{A}}]| \leq (1 - P(A)) X_M (V_M + \bar{V}) \leq 2k(V_M + \bar{V})\epsilon$$

while for the second term,

$$\begin{aligned} &|E[(x_{r_t(i)} p_t(i) - X_M V_M)\mathbb{1}_A]| \\ &\leq |E[(p_t(i) - V_M)\mathbb{1}_A x_{r_t(i)}]| + V_M |E[(x_{r_t(i)} - X_M)\mathbb{1}_A]| \\ &\leq \epsilon E[\mathbb{1}_A x_{r_t(i)}] + V_M X_M P(|x_{r_t(i)} - X_M| > \delta) + V_M \delta \\ &\leq (X_M + V_M X_M + V_M)\epsilon \end{aligned}$$

We conclude that for all period after some critical period, for any $\epsilon > 0$, if the number of bidders is large enough, $|E[x_{r_t(i)} p_t(i) - X_M V_M]| \leq \epsilon$, and so $|E[\Pi_t] - X_M V_M \sum_i \theta_i| \leq (\sum \theta_i) s\epsilon$. By lemma 3, $\mathcal{R}^{RB}_n \rightarrow X_M V_M \sum_i \theta_i = \lim_n \mathcal{R}^{Omn}_n$ as n grows to infinity. □

In particular :

Corollary 1. *The Google auction is asymptotically revenue-maximizing.*

4.2 B-ranking suboptimality

In a very competitive environment, the gap between value-per-click and declared bids is considerably reduced. The seller may charge each bidder nearly the highest possible amount she would be willing to pay, and extract most of the utility surplus of the winners. However, the seller's revenue strongly depends on the popularity of the displayed ads. Bidders with high values, or high bids, may not be the most relevant.

Naturally there may exist a relationship between ad popularity and bids, as the value and relevance for a given merchant is in general correlated. If relevance increases with value, on average, B-ranking will indirectly favor advertisers with high relevance. Nevertheless, it will not select the most relevant bidders, and indeed we can bound the average relevance of the top-ranked bidders. The bound has a statistical meaning, it is the average maximum relevance among a fixed-size random sample of merchants that would all have maximum value.

[6]We use $\mathbb{1}_A$ as the indicator function, it equals 1 when the event A is realized.

Lemma 4. *For an auction that allocates slots based on a B-ranking scheme, for any slot $i \leq s$,*

$$\limsup_{n \rightarrow +\infty} E[x_{\sigma_t(i)}] \leq m \int_x x F_{|\bar{V}}(x)^{m-1} f_{|\bar{V}}(x) dx$$

for some number $m > 0$, uniformly for all time beyond some critical period. The bound is tight.

In order to maximize revenue, the seller must be able to select the bidders with high values, but also with popular ads. B-ranking does not seem a good candidate, as it does not succeed in separating the bidders with highest relevance. In fact, we show that no auction with B-ranking can achieve an asymptotically maximal revenue.

Theorem 2. *Auctions that use B-ranking to allocate slots to advertisers never maximize revenue when the number of bidders is large. If $\mathcal{R}^B_n$ is the revenue of an auction using B-ranking auction, and $\mathcal{R}^{Opt}_n$ the revenue of an asymptotically optimal auction, then, when the number of bidders n is above some critical number N,*

$$\frac{\mathcal{R}^B_n}{\mathcal{R}^{Opt}_n} < C$$

with $C < 1$, and for any set of bidding strategies that verify assumption 1. C depends on the distribution of values and can be arbitrarily low for certain distributions.

Proof. For a sequential repeated auction as defined in section 3, the revenue at period t is:

$$\Pi_t = \sum_{i=1}^{s_t} \theta_i x_{\sigma_t(i)} p_t(i) \leq \bar{V} \sum_{i=1}^{s} \theta_i x_{\sigma_t(i)}$$

By lemma 4, if $t > T_c$ for some fixed period T_c, and some fixed $m > 0$,

$$\limsup_{n \rightarrow +\infty} E[\Pi_t] \;\leq\; \bar{V} \left[m \int_x x F_{|\bar{V}}(x)^{m-1} f_{|\bar{V}}(x) dx \right] \sum_i \theta_i$$

Note that the in-bracket part of the bound corresponds to the expected maximum of m variables independently distributed according to the marginal density $f_{|\bar{V}}$. If x lies in the support of the marginal distribution, $(x, \bar{V})$ is in the support of the joint distribution, and so $x\bar{V}$ is bounded above by $X_M V_M$, and may reach that upper value at one point at most according to the distributional assumptions made in section 3. This bound can be set arbitrarily low for well-chosen conditional distributions associated with strong positive skew, such as negatively correlated value and relevance.

Hence $\limsup_{n \rightarrow +\infty} E[\Pi_t] < X_M V_M \sum \theta_i$ for $t > T_c$, uniformly in t. Applying lemma 3, we obtain $\limsup_{n \rightarrow +\infty} \mathcal{R}^R_n < X_M V_M \sum \theta_i = \lim_{n \rightarrow +\infty} \mathcal{R}^{Opt}_n$, which concludes the proof. □

In particular, we have the following:

Corollary 2. *The Yahoo auction is not asymptotically revenue maximizing, and, if the number of bidders is large enough, its revenue is strictly lower than that of the Google auction.*

Note that, even though RB-ranking is sometimes referred to as "rank by revenue", the above result does not hold in general. This is obvious from our general assumptions on bidding strategies, but even in the more restrictive Nash equilibrium assumptions, Lahaie ([6]) shows that, with a fixed number of bidders, there are many cases for which the Yahoo auction revenue dominates that of Google.

The bound given in lemma 4 is tight, however it will only occur when the winners coordinate their bids in a certain way, function of their respective relevance. In many practical situations, however, merchants would not cooperate. They also may not have access to click frequencies or indicators of ad quality. If the seller decides of slot attributions and payments-per-click solely from bidding information, then relevance is not a useful parameter to consider for the bidders. Rather, it is plausible to assume that merchants will make bids by accounting for their own value, and possibly past allocations and bids if available. In other circumstances bidders may not get sufficiently accurate estimates on their own ad quality, and may prefer to restrain the impact of those uncertain quantities on their bidding strategy.

In those cases, the only parameter that influences the relevance of winning bidders is the possible correlation between value and relevance. Consequently it becomes possible to characterize the relevance of the top bidders in a precise way.

LEMMA 5. *For an auction using B-ranking, and whose payments depend solely on bids, assuming bids are not determined by relevance,*

$$\lim_{n\to\infty} E[x_{\sigma_t(i)}] = E[X|V = \bar{V}]$$

for any fixed rank i.

We now can quantify the average ad relevance of the winning bidders in an auction with B-ranking. B-ranking promotes bidders with high values, and participating in an auction that attracts many advertisers forces them to make competitive offers close to their own value. Consequently, for a regular auction, payments are maximized, and their revenue is the best possible among that of all auctions that use B-ranking. We prove our claim in the theorem below, after calculating the revenue of a regular auction with B-ranking.

LEMMA 6. *Assuming bids are not determined by relevance, the asymptotic revenue of a regular auction with B-ranking and whose payment scheme does not rely on ad relevance is*

$$\bar{V}E[X|V = \bar{V}]\sum_{i=1}^{s}\theta_i$$

THEOREM 3. *All regular auctions with B-ranking are asymptotically revenue-maximizing among auctions that allocate according to B-ranking and whose payments are determined by bids only.*

PROOF. The revenue at period t of such an auction is bounded above by $\bar{V}(\sum\theta_i x_{r_t(i)})$. Applying lemma 5, $E[x_{r_t(i)}]$ converges to $E[X|V = \bar{V}]$ as the number of bidders grows, uniformly when $t > T_c$, for some fixed T_c. Therefore, for $t > T_c$, the per-period revenue is asymptotically bounded above by $\bar{V}E[X|V = \bar{V}]\sum\theta_i$, uniformly in t, which, according to lemma 6, is the asymptotic revenue of a regular auction with B-ranking. By lemma 3, the revenue of the auction under consideration is bounded above by the same quantity, which concludes the proof. □

In general, regular auctions with B-ranking are not optimal among auction mechanisms with different allocation methods, even when they do not take relevance into account. When relevance and value are statistically related, allocation methods based on bids are, indirectly, also based on relevance. For example if bidders with very high value are associated with poor relevance it may be more profitable to allocate slots to bidders with lower value. Nonetheless, when relevance and value are unrelated, our result of asymptotic optimality of regular auctions with B-ranking generalizes over all auctions that do not rely on the bidders' relevance.

THEOREM 4. *If value and relevance are independent for each bidder, then all regular auctions with B-ranking are asymptotically revenue-maximizing among auctions whose allocation method and payments are solely determined by bids.*

PROOF. Consider any such auction. Assuming that bids are only function of initial values, past bids, ranks, and payments-per-click, as allocations and payments are determined by bids, and since value is independent of quality, bids are random variables independent of quality. Each $\sigma_t(i)$ being determined by bids, is also independent of quality. Hence,

$$\begin{aligned} E[x_{\sigma_t(i)}] &= \sum_j E[x_j \mathbb{1}_{\sigma_t(i)=j}] \\ &= \sum_j E[X]E[\mathbb{1}_{\sigma_t(i)=j}] = E[X] \end{aligned}$$

Since the expected profit of period t is bounded above by $\bar{V}(\sum\theta_i E[x_{r_t(i)}]) = \bar{V}E[X](\sum\theta_i)$, so by lemma 3, its revenue is no higher than $\bar{V}E[X](\sum\theta_i)$, the asymptotic revenue of a regular auction with B-ranking. □

In particular, for the Yahoo auction:

COROLLARY 3. *The Yahoo auction is asymptotically optimal among all auctions that use the same ranking method and with payments that are not function of ad quality. When there is no relationship between value and ad quality, the Yahoo auction is asymptotically optimal among all auctions for which payments and winner selection are determined by bids only.*

Our analysis shows that, when bidding strategies are not based on ad relevance, the ratio between the revenue of a B-ranking regular auction, such as the Yahoo auction, and an RB-ranking regular auction, such as the Google auction, converges to $E[X|V = \bar{V}]/\bar{X}$, or $E[X]/\bar{X}$ when relevance and value are independent. This suggests that the loss of revenue may be significant: in a very competitive environment, when relevance has an approximately symmetric distribution, a search engine that uses ranking by bids may lose half of the revenue it would have obtained by considering relevance-weighted bids.

5. GLOBAL CONDITIONS ON BIDDING STRATEGIES

Certain authors have suggested that the notion of Nash-equilibria of the stage game may, in certain situations, be relevant, since there exists some bidding strategies that oftentimes converge to such equilibria[7]. In this section we focus on second-price payment schemes and assume that bids converge in time to an equilibrium of the stage game, in a sense made precise by the following definition (see also [3]):

Definition 1. The strategies of bidders form *an eventual best-response equilibrium* for a sponsored search auction with second-price rule using a RB(ρ)-ranking if, for all bidders' values $v_1, \ldots, v_n$, there exists a fixed period T_e such that for all period $t > T_e$, bids remain constant and no bidder has incentive to change positions:
For all i,j, $j < i$,

$$\theta_i(v_{r(i)} - p(i) \geq \theta_j(v_{r(i)} - p(j-1))$$

and for all i,j, $j > i$,

$$\theta_i(v_{r(i)} - p(i) \geq \theta_j(v_{r(i)} - p(j))$$

where $r(\cdot)$ and $p(\cdot)$ are respectively the ranking and payment functions at any period $t > T_e$. (By convention, $p(0)$ is the price that a bidder would pay if she were ranked higher than the first bidder.)

ASSUMPTION 2. *Bidders' strategies form an eventual best-response equilibrium.*

In second-price payment schemes, winning bidders pay the minimum price they would need to bid to retain their position. In a competitive environment, the profits of the winners of the first positions tend to be small. The gap between the bids of the last winning bidder and the first excluded bidder must be relatively short at periods of equilibrium, otherwise one of the first assigned bidders could increase her profit by lowering her bid in order to obtain the last slot. Subsequently the first unassigned bidder tends to have a relatively high ranking value:

LEMMA 7. *At an equilibrium period, given any $\epsilon > 0$ and $\delta > 0$,*

$$P(|x^{\rho}_{r(s+1)} b^{1-\rho}_{r(s+1)} - \max_{(x,v)\in S} x^{\rho} v^{1-\rho}| \leq \delta) \geq 1 - \epsilon$$

when there are sufficiently many bidders, where $r(\cdot)$ and $\mathbf{b}$ are respectively the ranking function and bid vector at equilibrium, for $t > T_e$.

It follows from the previous probabilistic bound that winning bidders will pay nearly their own value. As we showed in the previous section, the RB-ranking method used by Google allows the seller to maximize revenue, while the B-ranking method used by Yahoo is suboptimal:

THEOREM 5. *When bidding strategies form an eventual best-response equilibrium,*

- *The Google auction is asymptotically optimal, and performs as well as an omniscient seller.*
- *The Yahoo auction is not asymptotically optimal.*

PROOF. We omit the proof due to space restrictions. □

[7]In our experiments, we found that myopic bidders who play best-response to the last round of bids converge to an equilibrium most of the time.

6. GENERAL DISTRIBUTIONS

Our optimality results make use of the fact that in competitive environments many advertisers will share similar characteristics. When that is not the case, or ranking alone is not sufficient to select bidders with the desired properties, the auction may fail to be optimal asymptotically.

THEOREM 6. *There exists distributions such that the Google auction is not asymptotically revenue-maximizing, when we consider both global and individual assumptions on bidders' behaviors.*

PROOF SKETCH. We assume that value and relevance are independent, that is, $f(x, v) = g(x)h(v)$. For relevance, we choose a uniform distribution on $[\underline{X}, \bar{X}]$. We assume that values are chosen in $[1, +\infty[$, according to the density with cumulative distribution $H(v) = 1 - v^{-1/\alpha}$, for any $\alpha < 1$. We assume w.l.o.g. that bidder k is the bidder with the k-th highest value. After some simplifications,

$$\begin{aligned} E[v_k] &= n\binom{n-1}{k-1}\int_1^{+\infty}(1-H(x))^{k-1}(H(x))^{n-k}h(x)dx \\ &= n\binom{n-1}{k-1}\int_0^1(1-y)^{k-1-\alpha}y^{n-k}dy \\ &= \frac{n!}{(k-1)!(k-\alpha)\cdot(k-\alpha+1)\cdots(n-\alpha)} \end{aligned}$$

by changing variables and integrating by parts. In particular,

$$\frac{E_n[v_k]}{E_n[v_{k+1}]} \geq \frac{k}{k-\alpha}$$

for any number of bidders.

We consider two available slots. The revenue of the truthful auction with RB-ranking is $(\theta_1 - \theta_2)E[x_2v_2] + 2\theta_2E[x_3v_3]$, bounded below by $\underline{X}((\theta_1 - \theta_2)E[v_2] + 2\theta_2E[v_3])$. In the Google auction, suppose that bidders bid their actual value, except bidder 2, who bids ϵ above bidder 3. These bidding strategies trivially meet the requirements of assumptions 1 and 2, and the revenue is bounded above by $\bar{X}((\theta_1 + \theta_2)(E[v_3] + \epsilon))$. Therefore, if ϵ and $\bar{X} - \underline{X}$ are small enough, the revenue of the truthful auction is higher than that of the Google auction. □

7. SUMMARY AND FUTURE WORK

In this paper, we made two main contributions. First we introduced a new model of bidders' behaviors, along with a new criterion for asymptotic optimality in a dynamic auction setting. Our model's novelty lies in its generality. Second, we obtained a collection of results under our new behavioral assumptions. We identified a broad class of auctions that are asymptotically revenue-maximizing, and that includes as a special case the Google auction, but not the Yahoo auction. Indeed, no auction that ranks bidders in *bid* order maximizes revenue when the number of bidder exceeds a certain threshold. Nevertheless, the Yahoo auction belongs a broad class of auctions that are asymptotically optimal in a weak sense: their revenue is at least as large as 1) that of any auction whose payments do not rely on ad relevance and uses *bid* ranking, and, if *relevance* and *bid* are unrelated, 2) that of any auction whose ranking or payments do not rely on ad relevance. We also showed that in the more commonly explored case of equilibrium of the one-shot auction, the Google auction remains optimal, while the Yahoo auction is suboptimal, asymptotically.

While our results suggest that the Google auction would generate more revenue that the Yahoo auction when many bidders are competing for a keyword, no such ranking exists when only few bidders are present: under both sets of assumptions, there are cases in which Google's revenue exceeds that of Yahoo and vice-versa.

There are many aspects of sponsored search auctions that we did not explore. Space limitations do not allow a detailed discussion of these, but in follow-up papers we will address the following questions, to which we have partial answers:

1. Asymptotic optimality is desirable when many bidders participate in the auction. For practical purposes, it is important to know how many bidders are needed to reach a near-optimal revenue. What are the convergence rates to optimal revenue? How do they compare for the different auctions?

2. In addition to selecting keywords of interests, bidders also specify their maximum daily budget. How are the results affected by considering budget-constraints, and multiple keywords?

3. Our results suggest that, with many bidders, keyword auction's revenues are not as affected by the payment method as by the allocation method. How do different allocation methods – other than the two studied in this paper – affect asymptotic optimality?

Acknowledgments

The authors would like to thank Hal Varian and the reviewers for their useful comments.

8. REFERENCES

[1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. *Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, 2006.

[2] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Working paper*, 2005.

[3] 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*, 2005.

[4] J. Feng, H. Bhargava, and D. Pennock. Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms. *Informs Journal on Computing*, 2006.

[5] G. Iyengar and A. Kumar. Characterizing optimal keyword auctions. *Second Workshop on Sponsored Search Auctions*, 2006.

[6] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. *Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, 2006.

[7] R. Myerson. Optimal Auction Design. *Math. Oper. Res.*, 6(1):58–73, 1981.

[8] H. Varian. Position auctions. *Working Paper*, 2006.

APPENDIX

PROOF OF LEMMA 4. Let $R_t = \sigma_t(i)$ be the bidder with rank $i \le s$ at period t, and let B_j be a bidder with the j-th highest value.

Let $\epsilon > 0$. First, note that $\prod_{j\le m} \frac{f(\tilde{x}_j, \tilde{v}_j)}{f(\tilde{v}_j)}$ is continuous on the compact S^m, and thus uniformly continuous by the Heine-Cantor theorem. So there exists some $\delta > 0$ so that if $|\tilde{v}_j - \bar{V}| \le \delta$, $j \le m$, then $|\prod_j f_{|\tilde{v}_j}(\tilde{x}_j) - \prod_j f_{|\bar{V}}(\tilde{x}_j)| \le \epsilon$ for all $\tilde{x}_1, \ldots, \tilde{x}_m$.

By lemma 1, R_t is one of the bidders with the top m values, where m is a fixed number independent of the number of bidders. So,

$$\begin{aligned} E[x_{R_t}] &\le E[\max_{j\le m} x_{B_j}] \\ &= E[\max_{j\le m} x_{B_j} \mathbb{1}_{v_{B_m} \ge \bar{V}-\delta}] \\ &\quad + E[\max_{j\le m} x_{B_j} \mathbb{1}_{v_{B_m} < \bar{V}-\delta}] \end{aligned}$$

By lemma 2, $|E[\max_{j\le m} x_{B_j} \mathbb{1}_{v_{B_m} < \bar{V}-\delta}]| \le \bar{X} P(v_{B_m} < \bar{V} - \delta) \le \bar{X}\epsilon$ for n large enough.

Besides, by the tower property of conditional expectations,

$$\begin{aligned} &E[\max_{j\le m} x_{B_j} \mathbb{1}_{v_{B_m} \ge \bar{V}-\delta}] = \\ &E[E[\max_{j\le m} x_{B_j} \mathbb{1}_{v_{B_m} \ge \bar{V}-\delta} | v_{B_1}, \ldots, v_{B_m}]] = \int_{\tilde{v}_1 \ge \cdots \ge \tilde{v}_m \ge \bar{V}-\delta} \\ &\left[\int_{\tilde{x}_1,\ldots,\tilde{x}_m} \max_j \tilde{x}_j \prod_j f_{|\tilde{v}_j}(\tilde{x}_j) d\tilde{x}_1 \ldots d\tilde{x}_m\right] dQ(\tilde{v}_1, \ldots, \tilde{v}_m) \\ &= E_1 + E_2 \end{aligned}$$

where Q is the probability measure associated with the top m values, in decreasing order, and with

$$\begin{aligned} |E_1| &= \left| \int_{\tilde{v}_1 \ge \cdots \ge \tilde{v}_m \ge \bar{V}-\delta} \left[\int_{\tilde{x}_1,\ldots,\tilde{x}_m} \max_j \tilde{x}_j \right.\right. \\ &\quad \left.\left. \prod_j (f_{|\tilde{v}_j}(\tilde{x}_j) - f_{|\bar{V}}(\tilde{x}_j)) d\tilde{x}_1 \ldots d\tilde{x}_m \right] dQ(\tilde{v}_1, \ldots, \tilde{v}_m) \right| \\ &\le \epsilon \int_{\tilde{v}_1,\ldots,\tilde{v}_m} \left[\int_{\tilde{x}_1,\ldots,\tilde{x}_m} \max_j \tilde{x}_j \right] dQ(\tilde{v}_1, \ldots, \tilde{v}_m) \\ &\le \epsilon \bar{X}^{m+1} \end{aligned}$$

and

$$\begin{aligned} E_2 &= \int_{\tilde{v}_1 \ge \cdots \ge \tilde{v}_m \ge \bar{V}-\delta} \\ &\quad \left[\int_{\tilde{x}_1,\ldots,\tilde{x}_m} \max_j \tilde{x}_j \prod_j f_{|\bar{V}}(\tilde{x}_j) d\tilde{x}_1 \ldots d\tilde{x}_m \right] dQ(\tilde{v}_1, \ldots, \tilde{v}_m) \\ &= \left[m \int_{\tilde{x}} \tilde{x} F_{|\bar{V}}^{m-1}(\tilde{x}) f_{|\bar{V}}(\tilde{x}) d\tilde{x} \right] . \\ &\qquad \int_{\tilde{v}_1 \ge \cdots \ge \tilde{v}_m \ge \bar{V}-\delta} dQ(\tilde{v}_1, \ldots, \tilde{v}_m) \\ &= \left[m \int_{\tilde{x}} \tilde{x} F_{|\bar{V}}^{m-1}(\tilde{x}) f_{|\bar{V}}(\tilde{x}) d\tilde{x} \right] P(v_{B_m} \ge \bar{V} - \delta) \\ &\le m \int_{\tilde{x}} \tilde{x} F_{|\bar{V}}^{m-1}(\tilde{x}) f_{|\bar{V}}(\tilde{x}) d\tilde{x} \end{aligned}$$

thus, any period $t > T_c$, for some period T_c, given any $\epsilon > 0$ if n is large enough,

$$E[R_t] \leq \bar{X}\epsilon + \bar{X}^{m+1}\epsilon + m\int_{\tilde{x}} \tilde{x}F^{m-1}_{|\bar{V}}(\tilde{x})f_{|\bar{V}}(\tilde{x})d\tilde{x}$$

hence,

$$\limsup_{n\to+\infty} E[x_{\sigma_t(i)}] \leq m\int_x xF_{|\bar{V}}(x)^{m-1}f_{|\bar{V}}(x)dx$$

uniformly in $t > T_c$. Note that the bound is tight (obtained when excluded bidders bids their true value and bids of the winners are ordered so that the i-th bidder has the highest ad relevance). □

Proof of lemma 5. Let $R_t = r_t(i)$ be the bidder of rank i for period t. Consider any $\epsilon > 0$. As $\frac{f(x,v)}{f(v)}$ is continuous on the compact S, it is uniformly continuous by the Heine-Cantor theorem, and there exists $\delta > 0$ such that $|v - \bar{V}| \leq \delta$ implies $|\frac{f(x,v)}{f(v)} - \frac{f(x,\bar{V})}{f(\bar{V})}| \leq \epsilon$ for all x.

We decompose the expectation:

$$\begin{aligned} E[x_{R_t}] &= E[x_{R_t}\mathbb{1}_{v_{R_t}<\bar{V}-\delta}] + E[x_{R_t}\mathbb{1}_{v_{R_t}\geq\bar{V}-\delta}] \\ &= E[x_{R_t}\mathbb{1}_{v_{R_t}<\bar{V}-\delta}] + \sum_{j=1}^{n} E[x_j\mathbb{1}_{R_t=j}\mathbb{1}_{v_j\geq\bar{V}-\delta}] \end{aligned}$$

Combining lemma 1 and lemma 2, we have $E[x_{R_t}\mathbb{1}_{v_{R_t}<\bar{V}-\delta}] \leq \bar{X}P(v_{r_t(i)} < \bar{V}-\delta) \leq \bar{X}\epsilon$ for n large enough, if $t > T_c$ for some T_c independent of ϵ.

R_t is a random variable that gives the bidder of rank i. Ranks are determined by bids, which are themselves determined by initial values, bidding history, previous allocations or payments, which in turn may depend on past bids, allocations and payments, but as none of these quantities depend directly on ad relevance, they ultimately are determined by initial values only. This means that R_t is measurable with respect to values, and that the indicator function $\mathbb{1}_{R_t=j}$ is defined on the set of values. Hence[8]:

$$\begin{aligned} &E[x_j\mathbb{1}_{R_t=j}\mathbb{1}_{v_j\geq\bar{V}-\delta}] = \\ &\quad\int_{\tilde{\mathbf{x}}}\int_{\tilde{\mathbf{v}},v_j\geq\bar{X}-\delta} \tilde{x}_j\mathbb{1}_{R_t=j}(\tilde{v}_1,\ldots,\tilde{v}_n)\prod_l f(\tilde{x}_l,\tilde{v}_l)d\tilde{\mathbf{x}}d\tilde{\mathbf{v}} \\ &= \int_{\tilde{\mathbf{v}},v_j\geq\bar{X}-\delta} \mathbb{1}_{R_t=j}(\tilde{v}_1,\ldots,\tilde{v}_n)\prod_l f(\tilde{v}_l)\left[\int_{\tilde{x}}\tilde{x}f_{|\tilde{v}_i}(\tilde{x})d\tilde{x}\right]d\tilde{\mathbf{v}} \\ &= E_1 + E_2 \end{aligned}$$

with

$$\begin{aligned} E_1 = \int_{\tilde{\mathbf{v}},v_j\geq\bar{X}-\delta} &\mathbb{1}_{R_t=j}(\tilde{v}_1,\ldots,\tilde{v}_n)\prod f(\tilde{v}_l) \\ &\left[\int_{\tilde{x}}\tilde{x}\left(f_{|\tilde{v}_i}(\tilde{x}) - f_{|\bar{V}}(\tilde{x})\right)d\tilde{x}\right]d\tilde{\mathbf{v}} \end{aligned}$$

noting that

$$\begin{aligned} |E_1| &\leq \epsilon\bar{X}^2\int_{\tilde{\mathbf{v}}}\mathbb{1}_{R_t=j}(\tilde{v}_1,\ldots,\tilde{v}_n)\prod f(\tilde{v}_l)d\tilde{\mathbf{v}} \\ &\leq \bar{X}^2P(R_t=j)\epsilon \end{aligned}$$

[8]We use the vector notation $\mathbf{a} = (a_1,\ldots,a_n)$.

while

$$\begin{aligned} E_2 &= \left[\int_{\tilde{\mathbf{v}},\tilde{v}_j\geq\bar{V}-\delta}\mathbb{1}_{R_t=j}(\tilde{v}_1,\ldots,\tilde{v}_n)\prod f(\tilde{v}_l)d\mathbf{v}\right]. \\ &\quad\left[\int_{\tilde{x}}\tilde{x}f_{|\bar{V}}(\tilde{x})d\tilde{x}\right] \\ &= P(\{R_t=j\}\cap\{v_j\geq\bar{V}-\delta\})E[X|V=\bar{V}] \end{aligned}$$

As $\sum_{j=1}^{n}P(R_t=j) = 1$ and $\sum_{j=1}^{n}P(\{R_t=j\}\cap\{v_j\geq\bar{V}-\delta\}) = P(v_{R_t}\geq\bar{V}-\delta)$, by summation, we get:

$$\begin{aligned} &|E[x_{R_t}\mathbb{1}_{v_{R_t}\geq\bar{V}-\delta}] - E[X|V=\bar{V}]| \\ &\qquad\leq \bar{X}^2\epsilon + E[X|V=\bar{V}]P(v_{R_t}<\bar{V}-\delta) \\ &\qquad\leq \bar{X}^2\epsilon + \bar{X}\epsilon \end{aligned}$$

by lemma 1 and 2, for n large enough.

Thus, for any $\epsilon > 0$, $|E[x_{R_t}] - E[X|V=\bar{V}]| \leq (\bar{X}+\bar{X}^2+\bar{X})\epsilon$ for all $t > T_c$ when n is large enough, and we conclude $E[x_{R_t}] \to E[X|V=\bar{V}]$ as $n\to+\infty$, when t is large enough, uniformly in t. □

Proof of lemma 6. The profit at t is:

$$\Pi_t = \sum_{i=1}^{s}\theta_i x_{r_t(i)}p_t(i)$$

with $p_t(i) = q_i(b^{(t)}_{r_t(1)},\ldots,b^{(t)}_{r_t(k)})$. We use the same event A as in theorem 1, and note that

$$\begin{aligned} E[x_{r_t(i)}p_t(i)] = \bar{V}&E[x_{r_t(i)}] - \bar{V}E[\mathbb{1}_A x_{r_t(i)}] \\ &+ E[(p_t(i)-\bar{V})x_{r_t(i)}\mathbb{1}_A] \\ &+ E[x_{r_t(i)}p_t(i)\mathbb{1}_{\bar{A}}] \end{aligned}$$

Each of the terms converges to 0 except for the first one, which converges to $\bar{V}E[X|V=\bar{V}]$. The rest of the proof is similar to that of theorem 1. □

Proof of lemma 7. Let $\epsilon,\delta > 0$, and $M = \max_{(x,v)\in S}x^\rho v^{1-\rho}$. Let K be the bidder with the $s+1$-th highest product value $x_i^\rho v_i^{1-\rho}$. At equilibrium, for any slot i, $x^\rho_{r(i)}b^{1-\rho}_{r(i)} \geq x_K v_K$, otherwise bidder K would be able to get a slot and a positive profit by raising her bid. Additionally, $x^\rho_{r(s+1)}b^{1-\rho}_{r(s+1)} \leq M$, otherwise the winning bidders would have a negative utility.

Consider a position $i < s$, and suppose that $|x^\rho_K v^{1-\rho}_K - M| \leq \delta'$, with $\delta' = \delta(1+\frac{\theta_i}{\theta_s})^{-1}$. By definition 1, $\theta_i(x_{r(i)}v_{r(i)} - x_{r(i+1)}b_{r(i+1)}) \geq \theta_s(x_{r(i)}v_{r(i)} - x_{r(s+1)}v_{r(s+1)}$. Therefore

$$\begin{aligned} M \geq x^\rho_{r(s+1)}b^{1-\rho}_{r(s+1)} &\geq x^\rho_{r(i)}v^{1-\rho}_{r(i)} \\ &\quad - \frac{\theta_i}{\theta_s}(x^\rho_{r(i)}v^{1-\rho}_{r(i)} - x^\rho_{r(i+1)}b^{1-\rho}_{r(i+1)}) \\ &\geq x^\rho_K v^{1-\rho}_K - \frac{\theta_i}{\theta_s}(M - x^\rho_K v^{1-\rho}_K) \\ &\geq M - \delta' - \frac{\theta_i}{\theta_s}(M-(M-\delta')) \\ &\geq M - \delta \end{aligned}$$

By lemma 2, if the number of bidders n is large enough, then $P(|x^\rho_K v^{1-\rho}_K - M| \leq \delta') \geq 1-\epsilon$, and so $P(|x^\rho_{r(s+1)}b^{1-\rho}_{r(s+1)} - M| \leq \delta) \geq 1-\epsilon)$. □