Revenue Analysis of a Family of Ranking Rules for Keyword Auctions

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

David M. Pennock
Yahoo! Research
New York, NY 10011
pennockd@yahoo-inc.com

ABSTRACT

Keyword auctions lie at the core of the business models of today's leading search engines. Advertisers bid for placement alongside search results, and are charged for clicks on their ads. Advertisers are typically ranked according to a score that takes into account their bids and potential click-through rates. We consider a family of ranking rules that contains those typically used to model Yahoo! and Google's auction designs as special cases. We find that in general neither of these is necessarily revenue-optimal in equilibrium, and that the choice of ranking rule can be guided by considering the correlation between bidders' values and click-through rates. We propose a simple approach to determine a revenue-optimal ranking rule within our family, taking into account effects on advertiser satisfaction and user experience. We illustrate the approach using Monte-Carlo simulations based on distributions fitted to Yahoo! bid and click-through rate data for a high-volume keyword.

Categories and Subject Descriptors

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

General Terms

Economics, Theory

Keywords

sponsored search, search engines, keyword auctions

1. INTRODUCTION

Major search engines like Google, Yahoo!, and MSN sell advertisements by auctioning off space on keyword search results pages. For example, when a user searches the web for "iPod", the highest paying advertisers (for example, Apple or Best Buy) for that keyword may appear in a separate "sponsored" section of the page above or to the right of the algorithmic results. The sponsored 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 web page. Generally, advertisements that appear in a higher position on the page garner more attention and more clicks from users. Thus, all else being equal, advertisers prefer higher positions to lower positions.

Advertisers bid for placement on the page in an auction-style format where the larger their bid the more likely their listing will appear above other ads on the page. By convention, sponsored search advertisers generally bid and 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. Overture Services, formerly GoTo.com and now owned by Yahoo! Inc., is credited with pioneering sponsored search advertising. Overture's success prompted a number of companies to adopt similar business models, most prominently Google, the leading web search engine today. Microsoft's MSN, previously an affiliate of Overture, now operates its own keyword auction marketplace. Sponsored search is one of the fastest growing, most effective, and most profitable forms of advertising, generating roughly $7 billion in revenue in 2005 after nearly doubling every year for the previous five years.

The search engine evaluates the advertisers' bids and allocates the positions on the page accordingly. Notice that, although bids are expressed as payments per click, the search engine cannot directly allocate clicks, but rather allocates *impressions*, or placements on the screen. Clicks relate only stochastically to impressions. Until recently, Yahoo! ranked bidders in decreasing order of advertisers' stated values per click, while Google ranks in decreasing order of advertisers' stated values per impression. In Google's case, value per impression is computed by multiplying the advertiser's (per-click) bid by the advertisement's expected click-through rate, where this expectation may consider a number of unspecified factors including historical click-through rate, position on the page, advertiser identity, user identity, and the context of other items on the page. We refer to these rules as "rank-by-bid" and "rank-by-revenue", respectively.[1]

We analyze a family of ranking rules that contains the Yahoo! and Google models as special cases. We consider rank-

*This work was done while the author was at Yahoo! Research.

[1]These are industry terms. We will see, however, that rank-by-revenue is not necessarily revenue-optimal.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
EC'07, June 11–15, 2007, San Diego, California, USA.
Copyright 2007 ACM 978-1-59593-653-0/07/0006 ...$5.00.

ing rules where bidders are ranked in decreasing order of score $e^q b$, where e denotes an advertiser's click-through rate (normalized for position) and b his bid. Notice that $q = 0$ corresponds to Yahoo!'s rank-by-bid rule and $q = 1$ corresponds to Google's rank-by-revenue rule. Our premise is that bidders are playing a symmetric equilibrium, as defined by Edelman, Ostrovsky, and Schwarz [3] and Varian [11]. We show through simulation that although $q = 1$ yields the efficient allocation, settings of q considerably less than 1 can yield superior revenue in equilibrium under certain conditions. The key parameter is the correlation between advertiser value and click-through rate. If this correlation is strongly positive, then smaller q are revenue-optimal. Our simulations are based on distributions fitted to data from Yahoo! keyword auctions. We propose that search engines set thresholds of acceptable loss in advertiser satisfaction and user experience, then choose the revenue-optimal q consistent with these constraints. We also compare the potential gains from tuning q with the gains from setting reserve prices, and find that the former may be much more significant.

In Section 2 we give a formal model of keyword auctions, and establish its equilibrium properties in Section 3. In Section 4 we note that giving agents bidding credits can have the same effect as tuning the ranking rule explicitly. In Section 5 we give a general formulation of the optimal keyword auction design problem as an optimization problem, in a manner analogous to the single-item auction setting. We then provide some theoretical insight into how tuning q can improve revenue, and why the correlation between bidders' values and click-through rates is relevant. In Section 6 we consider the effect of q on advertiser satisfaction and user experience. In Section 7 we describe our simulations and interpret their results.

Related work. As mentioned the papers of Edelman et al. [3] and Varian [11] lay the groundwork for our study. Both papers independently define an appealing refinement of Nash equilibrium for keyword auctions and analyze its equilibrium properties. They called this refinement "locally envy-free equilibrium" and "symmetric equilibrium", respectively. Varian also provides some empirical analysis.

The general model of keyword auctions used here, where bidders are ranked according to a weight times their bid, was introduced by Aggarwal, Goel, and Motwani [1]. That paper also makes a connection between the revenue of keyword auctions in incomplete information settings with the revenue in symmetric equilibrium. Iyengar and Kumar [5] study the optimal keyword auction design problem in a setting of incomplete information, and also make the connection to symmetric equilibrium. We make use of this connection when formulating the optimal auction design problem in our setting.

The work most closely related to ours is that of Feng, Bhargava, and Pennock [4]. They were the first to realize that the correlation between bidder values and click-through rates should be a key parameter affecting the revenue performance of various ranking mechanisms. For simplicity, they assume bidders bid their true values, so their model is very different from ours and consequently so are their findings. According to their simulations, rank-by-revenue always (weakly) dominates rank-by-bid in terms of revenue, whereas our results suggest that rank-by-bid may do much better for negative correlations.

Lahaie [8] gives an example that suggests rank-by-bid should yield more revenue when values and click-through rates are positively correlated, whereas rank-by-revenue should do better when the correlation is negative. In this work we make a deeper study of this conjecture.

2. MODEL

There are K positions to be allocated among N bidders, where $N > K$. We assume that the (expected) click-through rate of bidder s in position t is of the form $e_s x_t$, i.e. separable into an advertiser effect $e_s \in [0, 1]$ and position effect $x_t \in [0, 1]$. We assume that $x_1 > x_2 > \ldots > x_K > 0$ and let $x_t = 0$ for $t > K$. We also refer to e_s as the *relevance* of bidder s. It is useful to interpret x_t as the probability that an ad in position t will be noticed, and e_s as the probability that it will be clicked on if noticed.

Bidder s has value v_s for each click. Bidders have quasi-linear utility, so that the utility to bidder s of obtaining position t at a price of p per click is

$$e_s x_t (v_s - p).$$

A weight w_s is associated with agent s, and agents bid for position. If agent s bids b_s, his corresponding *score* is $w_s b_s$. Agents are ranked by score, so that the agent with highest score is ranked first, and so on. We assume throughout that agents are numbered such that agent s obtains position s.

An agent pays per click the lowest bid necessary to retain his position, so that the agent in slot s pays $\frac{w_{s+1}}{w_s} b_{s+1}$. The auctioneer may introduce a reserve score of r, so that an agent's ad appears only if his score is at least r. For agent s, this translates into a reserve price (minimum bid) of r/w_s.

3. EQUILIBRIUM

We consider the pure-strategy Nash equilibria of the auction game. This is a full-information concept. The motivation for this choice is that in a keyword auction, bidders are allowed to continuously adjust their bids over time, and hence obtain estimates of their profits in various positions. As a result it is reasonable to assume that if bids stabilize, bidders should be playing best-responses to each other's bids [2, 3, 11]. Formally, in a Nash equilibrium of this game the following inequalities hold.

$$e_s x_s \left(v_s - \frac{w_{s+1}}{w_s} b_{s+1} \right) \geq e_s x_t \left(v_s - \frac{w_{t+1}}{w_s} b_{t+1} \right) \forall t > s \quad (1)$$

$$e_s x_s \left(v_s - \frac{w_{s+1}}{w_s} b_{s+1} \right) \geq e_s x_t \left(v_s - \frac{w_t}{w_s} b_t \right) \forall t < s \quad (2)$$

Inequalities (1) and (2) state that bidder s does not prefer a lower or higher position to his own, respectively. It can be hard to derive any theoretical insight into the properties of these Nash equilibria—multiple allocations of positions to bidders can potentially arise in equilibrium [2]. Edelman, Ostrovsky, and Schwarz [3] introduced a refinement of Nash equilibrium called "locally envy-free equilibrium" that is more tractable to analyze; Varian [11] independently proposed this solution concept and called it "symmetric equilibrium". In a symmetric equilibrium, inequality (1) holds for all s, t rather than just for $t > s$. So for all s and all $t \neq s$, we have

$$e_s x_s \left(v_s - \frac{w_{s+1}}{w_s} b_{s+1} \right) \geq e_s x_t \left(v_s - \frac{w_{t+1}}{w_s} b_{t+1} \right),$$

or equivalently

$$x_s(w_s v_s - w_{s+1} b_{s+1}) \geq x_t(w_s v_s - w_{t+1} b_{t+1}). \quad (3)$$

Edelman et al. [3] note that this equilibrium arises if agents are raising their bids to increase the payments of those above them, a practice which is believed to be common in actual keyword auctions. Varian [11] provides some empirical evidence that Google bid data agrees well with the hypothesis that bidders are playing a symmetric equilibrium.

Varian does a thorough analysis of the properties of symmetric equilibrium, assuming $w_s = e_s = 1$ for all bidders. It is straightforward to adapt his analysis to the case where bidders are assigned arbitrary weights and have separable click-through rates.[2] As a result we find that in symmetric equilibrium, bidders are ranked in order of decreasing $w_s v_s$. To be clear, although the auctioneer only has access to the bids b_s and not the values v_s, in symmetric equilibrium the bids are such that ranking according to $w_s b_s$ is equivalent to ranking according to $w_s v_s$.

The smallest possible bid profile that can arise in symmetric equilibrium is given by the recursion

$$x_s w_{s+1} b_{s+1} = (x_s - x_{s+1}) w_{s+1} v_{s+1} + x_{s+1} w_{s+2} b_{s+2}.$$

In this work we assume that bidders are playing the smallest symmetric equilibrium. This is an appropriate selection for our purposes: by optimizing revenue in this equilibrium, we are optimizing a lower bound on the revenue in any symmetric equilibrium. Unraveling the recursion yields

$$x_s w_{s+1} b_{s+1} = \sum_{t=s}^{K} (x_t - x_{t+1}) w_{t+1} v_{t+1}. \quad (4)$$

Agent s's total expected payment is e_s/w_s times the quantity on the left-hand side of (4). The base case of the recursion occurs for $s = K$, where we find that the first excluded bidder bids his true value, as in the original analysis.

Multiplying each of the inequalities (4) by the corresponding e_s/w_s to obtain total payments, and summing over all positions, we obtain a total equilibrium revenue of

$$\sum_{s=1}^{K} \sum_{t=s}^{K} \frac{w_{t+1}}{w_s} e_s (x_t - x_{t+1}) v_{t+1}. \quad (5)$$

To summarize, the minimum possible revenue in symmetric equilibrium can be computed as follows, given the agents' relevance-value pairs (e_s, v_s): first rank the agents in decreasing order of $w_s v_s$, and then evaluate (5).

With a reserve score of r, it follows from inequality (3) that no bidder with $w_s v_s < r$ would want to participate in the auction. Let $K(r)$ be the number of bidders with $w_s v_s \geq r$, and assume it is at most K. We can impose a reserve score of r by introducing a bidder with value r and weight 1, and making him the first excluded bidder (who in symmetric equilibrium bids truthfully). In this case the recursion yields

$$x_s w_{s+1} b_{s+1} = \sum_{t=s}^{K(r)-1} (x_t - x_{t+1}) w_{t+1} v_{t+1} + x_{K(r)} r$$

and the revenue formula is adapted similarly.

[2]If we redefine $w_s v_s$ to be v_s and $w_s b_s$ to be b_s, we recover Varian's setup and his original analysis goes through unchanged.

4. BIDDING CREDITS

An indirect way to influence the allocation is to introduce bidding credits.[3] Suppose bidder s is only required to pay a fraction $c_s \in [0, 1]$ of the price he faces, or equivalently a $(1 - c_s)$ fraction of his clicks are received for free. Then in a symmetric equilibrium, we have

$$e_s x_s \left(v_s - \frac{w_{s+1}}{w_s} c_s b_{s+1} \right) \geq e_s x_t \left(v_s - \frac{w_{t+1}}{w_s} c_s b_{t+1} \right)$$

or equivalently

$$x_s \left(\frac{w_s}{c_s} v_s - w_{s+1} b_{s+1} \right) \geq x_t \left(\frac{w_s}{c_s} v_s - w_{t+1} b_{t+1} \right).$$

If we define $w'_s = \frac{w_s}{c_s}$ and $b'_s = c_s b_s$, we recover inequality (3). Hence the equilibrium revenue will be as if we had used weights w' rather than w. The bids will be scaled versions of the bids that arise with weights w' (and no credits), where each bid is scaled by the corresponding factor $1/c_s$.

This technique allows one to use credits instead of explicit changes in the weights to affect revenue. For instance, rank-by-revenue will yield the same revenue as rank-by-bid if we set credits to $c_s = e_s$.

5. REVENUE

We are interested in setting the weights w to achieve optimal expected revenue. The setup is as follows. The auctioneer chooses a function g so that the weighting scheme is $w_s \equiv g(e_s)$. We do not consider weights that also depend on the agents' bids because this would invalidate the equilibrium analysis of the previous section.[4] A pool of N bidders is then obtained by i.i.d. draws of value-relevance pairs from a common probability density $f(e_s, v_s)$. We assume the density is continuous and has full support on $[0, 1] \times [0, \infty)$. The revenue to the auctioneer is then the revenue generated in symmetric equilibrium under weighting scheme w. This assumes the auctioneer is patient enough not to care about revenue until bids have stabilized.

The problem of finding an optimal weighting scheme can be formulated as an optimization problem very similar to the one derived by Myerson [9] for the single-item auction case (with incomplete information). Let $Q_{sk}(e, v; w) = 1$ if agent s obtains slot k in equilibrium under weighting scheme w, where $e = (e_1, \ldots, e_N)$ and $v = (v_1, \ldots, v_N)$, and let it be 0 otherwise.

Note that the total payment of agent s in equilibrium is

$$e_s x_s \frac{w_{s+1}}{w_s} b_{s+1} = \sum_{t=s}^{K} e_s (x_t - x_{t+1}) \frac{w_{t+1}}{w_s} v_{t+1}$$

$$= e_s x_s v_s - \int_0^{v_s} \sum_{k=1}^{K} e_s x_k Q_{sk}(e_s, e_{-s}, y, v_{-s}; w)\, dy.$$

The derivation then continues just as in the case of a single-item auction [7, 9]. We take the expectation of this payment,

[3]Hal Varian suggested to us that bidding credits could be used to affect revenue in keyword auctions, which prompted us to look into this connection.

[4]The analysis does not generalize to weights that depend on bids. It is unclear whether an equilibrium would exist at all with such weights.

and sum over all agents to obtain the objective

$$\int_0^\infty \int_0^\infty \left[\sum_{s=1}^{N} \sum_{k=1}^{K} e_s x_k \psi(e_s, v_s) Q_{sk}(e, v; w) \right] f(e, v)\, dv\, de,$$

where ψ is the "virtual valuation"

$$\psi(e_s, v_s) = v_s - \frac{1 - F(v_s|e_s)}{f(v_s|e_s)}.$$

According to this analysis, we should rank bidders by "virtual score" $e_s\psi(e_s, v_s)$ to optimize revenue (and exclude any bidders with negative virtual score). However, unlike in the incomplete information setting, here we are constrained to ranking rules that correspond to a certain weighting scheme $w_s \equiv g(e_s)$. We remark that the virtual score cannot be reproduced exactly via a weighting scheme.

LEMMA 1. *There is no weighting scheme g such that the virtual score equals the score, for any density f.*

PROOF. Assume there is a g such that $e\psi(e, v) = g(e)v$. (The subscript s is suppressed for clarity.) This is equivalent to

$$\frac{d}{dv} \log(1 - F(v|e)) = h(e)/v, \tag{6}$$

where $h(e) = (g(e)/e - 1)^{-1}$. Let $\bar{v}$ be such that $F(\bar{v}|e) < 1$; under the assumption of full support, there is always such a $\bar{v}$. Integrating (6) with respect to v from 0 to $\bar{v}$, we find that the left-hand side converges whereas the right-hand side diverges, a contradiction. □

Of course, to rank bidders by virtual score, we only need $g(e_s)v_s = h(e_s\psi(e_s, v_s))$ for some monotonically increasing transformation h. (A necessary condition for this is that $\psi(e_s, v_s)$ be increasing in v_s for all e_s.) Absent this regularity condition, the optimization problem seems quite difficult because it is so general: we need to maximize expected revenue over the space of all functions g.

To simplify matters, we now restrict our attention to the family of weights $w_s = e_s^q$ for $q \in (-\infty, +\infty)$. It should be much simpler to find the optimum within this family, since it is just one-dimensional. Note that it covers rank-by-bid ($q = 0$) and rank-by-revenue ($q = 1$) as special cases.

To see how tuning q can improve matters, consider again the equilibrium revenue:

$$R(q) = \sum_{s=1}^{K} \sum_{t=s}^{K} \left(\frac{e_{t+1}}{e_s} \right)^q e_s(x_t - x_{t+1})v_{t+1}. \tag{7}$$

If the bidders are ranked in decreasing order of relevance, then $\frac{e_t}{e_s} \leq 1$ for $t > s$ and decreasing q slightly without affecting the allocation will increase revenue. Similarly, if bidders are ranked in increasing order of relevance, increasing q slightly will yield an improvement. Now suppose there is perfect positive correlation between value and relevance. In this case, rank-by-bid will always lead to the same allocation as rank-by-revenue, and bidders will always be ranked in decreasing order of relevance. It then follows from (7) that $q = 0$ will yield more revenue in equilibrium than $q = 1$.[5]

[5]It may appear that this contradicts the revenue-equivalence theorem [7, 9], because mechanisms that always lead to the same allocation in equilibrium should yield the same revenue. Note though that with perfect correlation, there are α, β such that $v_s = \alpha e_s + \beta$. So the assumption of full support is violated, which is necessary for revenue equivalence. Recall that a density has full support over a given domain if every point in the domain has positive density.

If a good estimate of f is available, Monte-Carlo simulations can be used to estimate the revenue curve as a function of q, and the optimum can be located. Simulations can also be used to quantify the effect of correlation on the location of the optimum. We do this in Section 7.

6. EFFICIENCY AND RELEVANCE

In principle the revenue-optimal parameter q may lie anywhere in $(-\infty, \infty)$. However, tuning the ranking rule also has consequences for advertiser satisfaction and user experience, and taking these into account reduces the range of allowable q.

The *total relevance* of the equilibrium allocation is

$$L(q) \quad = \quad \sum_{s=1}^{K} e_s x_s,$$

i.e. the aggregate click-through rate. Presumably users find the ad display more interesting and less of a nuisance if they are more inclined to click on the ads, so we adopt total relevance as a measure of user experience.

Let $p_s = \frac{w_{s+1}}{w_s} b_{s+1}$ be the price per click faced by bidder s. The *total value (efficiency)* generated by the auction in equilibrium is

$$\begin{aligned} V(q) \quad &= \quad \sum_{s=1}^{K} e_s x_s v_s \\ &= \quad \sum_{s=1}^{K} e_s x_s (v_s - p_s) + \sum_{s=1}^{K} e_s x_s p_s. \end{aligned}$$

As we see total value can be reinterpreted as total profits to the bidders and auctioneer combined. Since we only consider deviations from maximum efficiency that increase the auctioneer's profits, any decrease in efficiency in our setting corresponds to a decrease in bidder profits. We therefore adopt efficiency as a measure of advertiser satisfaction.

We would expect total relevance to increase with q, since more weight is placed on each bidder's individual relevance. We would expect efficiency to be maximized at $q = 1$, since in this case a bidder's weight is exactly his relevance.

PROPOSITION 1. *Total relevance is non-decreasing in q.*

PROOF. Recall that in symmetric equilibrium, bidders are ranked in order of decreasing $w_s v_s$. Let $\epsilon > 0$. Perform an exchange sort to obtain the ranking that arises with $q + \epsilon$ starting from the ranking that arises with q (for a description of exchange sort and its properties, see Knuth [6] pp. 106–110). Assume that ϵ is large enough to make the rankings distinct. Agents s and t, where s is initially ranked lower than t, are swapped in the process if and only if the following conditions hold:

$$\begin{aligned} e_s^q v_s \quad &\leq \quad e_t^q v_t \\ e_s^{q+\epsilon} v_s \quad &> \quad e_t^{q+\epsilon} v_t \end{aligned}$$

which together imply that $e_s^\epsilon > e_t^\epsilon$ and hence $e_s > e_t$ as $\epsilon > 0$. At some point in the sort, agent s occupies some slot

k while agent t occupies slot $k-1$. After the swap, total relevance will have changed by the amount

$$\begin{aligned} & e_s x_{k-1} + e_t x_k - e_t x_{k-1} - e_s x_k \\ = \; & (e_s - e_t)(x_{k-1} - x_k) > 0 \end{aligned}$$

As relevance strictly increases with each swap in the sort, total relevance is strictly greater when using $q+\epsilon$ rather than q. □

PROPOSITION 2. *Total value is non-decreasing in q for $q \leq 1$ and non-increasing in q for $q \geq 1$.*

PROOF. Let $q \geq 1$ and let $\epsilon > 0$. Perform an exchange sort to obtain the second ranking from the first as in the previous proof. If agents s and t are swapped, where s was initially ranked lower than t, then $e_s > e_t$. This follows by the same reasoning as in the previous proof. Now $e_s^{1-q} \leq e_t^{1-q}$ as $1-q \leq 0$. This together with $e_s^q v_s \leq e_t^q v_t$ implies that $e_s v_s \leq e_t v_t$. Hence after swapping agents s and t, total value has not increased. The case for $q \leq 1$ is similar. □

Since the trends described in Propositions 1 and 2 hold pointwise (i.e. for any set of bidders), they also hold in expectation. Proposition 2 confirms that efficiency is indeed maximized at $q = 1$.

These results motivate the following approach. Although tuning q can optimize current revenue, this may come at the price of future revenue because advertisers and users may be lost, seeing as their satisfaction decreases. To guarantee future revenue will not be hurt too much, the auctioneer can impose bounds on the percent efficiency and relevance loss he is willing to tolerate, with $q = 1$ being a natural baseline. By Proposition 2, a lower bound on efficiency will yield upper and lower bounds on the search space for q. By Proposition 1, a lower bound on relevance will yield another lower bound on q. The revenue curve can then be plotted within the allowable range of q to find the revenue-optimal setting.

7. SIMULATIONS

To add a measure of reality to our simulations, we fit distributions for value and relevance to Yahoo! bid and click-through rate data for a certain keyword that draws over a million searches per month. (We do not reveal the identity of the keyword to respect the privacy of the advertisers.)

We obtained click and impression data for the advertisers bidding on the keyword. From this we estimated advertiser and position effects using a maximum-likelihood criterion. We found that, indeed, position effects are monotonically decreasing with lower rank. We then fit a beta distribution to the advertiser effects resulting in parameters $a = 2.71$ and $b = 25.43$.

We obtained bids of advertisers for the keyword. Using Varian's [11] technique, we derived bounds on the bidders' actual values given these bids. By this technique, upper and lower bounds are obtained on bidder values given the bids according to inequality (3). If the interval for a given value is empty, i.e. its upper bound lies below its lower bound, then we compute the smallest perturbation to the bids necessary to make the interval non-empty, which involves solving a quadratic program. We found that the mean absolute deviation required to fit bids to symmetric equilibrium was always at most 0.08, and usually significantly less, over different days in a period of two weeks.[6] We fit a lognormal distribution to the lower bounds on the bidders' values, resulting in parameters $\mu = 0.35$ and $\sigma = 0.71$.

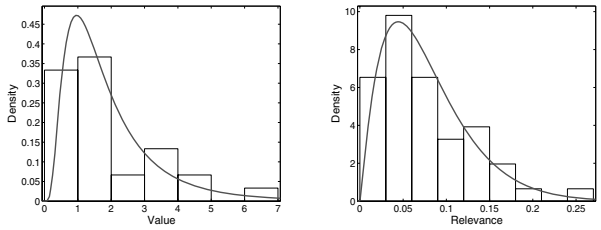

Figure 1: Empirical marginal distributions of value and relevance.

The empirical distributions of value and relevance together with the fitted lognormal and beta curves are given in Figure 1. It appears that mixtures of beta and lognormal distributions might be better fits, but since these distributions are used mainly for illustration purposes, we err on the side of simplicity.

We used a Gaussian copula to create dependence between value and relevance.[7] Given the marginal distributions for value and relevance together with this copula, we simulated the revenue effect of varying q for different levels of Spearman correlation, with 12 slots and 13 bidders. The results are shown in Figure 2.[8]

It is apparent from the figure that the optimal choice of q moves to the right as correlation decreases; this agrees with our intuition from Section 5. The choice is very sensitive to the level of correlation. If choosing only between rank-by-bid and rank-by-revenue, rank-by-bid is best for positive correlation whereas rank-by-revenue is best for negative correlation. At zero correlation, they give about the same expected revenue in this instance. Figure 2 also shows that in principle, the optimal q may be negative. It may also occur beyond 1 for different distributions, but we do not know if these would be realistic. The trends in efficiency and relevance are as described in the results from Section 6. (Any small deviation from these trends is due to the randomness inherent in the simulations.) The curves level off as $q \to +\infty$ because eventually agents are ranked purely according to relevance, and similarly as $q \to -\infty$.

A typical Spearman correlation between value and relevance for the keyword was about 0.4—for different days in a week the correlation lay within $[0.36, 0.55]$. Simulation results with this correlation are in Figure 3. In this instance rank-by-bid is in fact optimal, yielding 25% more revenue than rank-by-revenue. However, at $q = 0$ efficiency and relevance are 9% and 17% lower than at $q = 1$, respectively. Imposing a bound of, say, 5% on efficiency and relevance loss from the baseline at $q = 1$, the optimal setting is $q = 0.6$, yielding 11% more revenue than the baseline.

[6] See Varian [11] for a definition of mean absolute deviation.
[7] A copula is a function that takes marginal distributions and gives a joint distribution with these marginals. It can be designed so that the variables are correlated. See for example Nelsen [10].
[8] The y-axes in Figures 2–4 have been normalized because the simulations are based on proprietary data. Only relative values are meaningful.

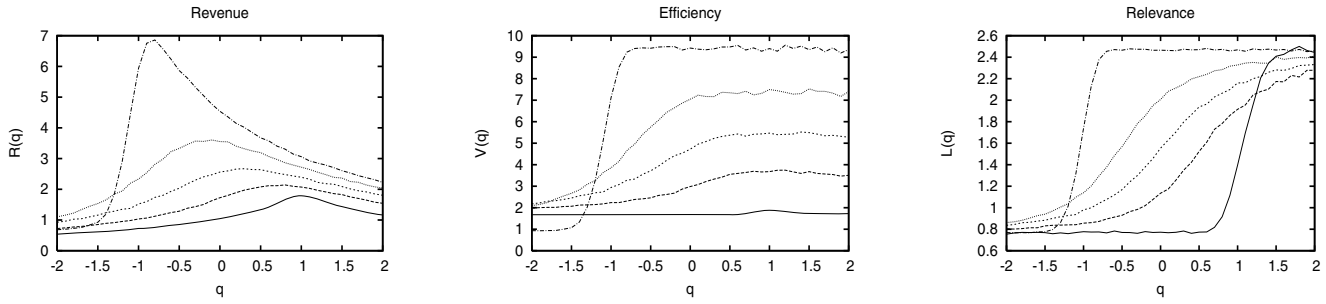

Figure 2: Revenue, efficiency, and relevance for different parameters q under varying Spearman correlation (key at right). Estimated standard errors are less than 1% of the values shown.

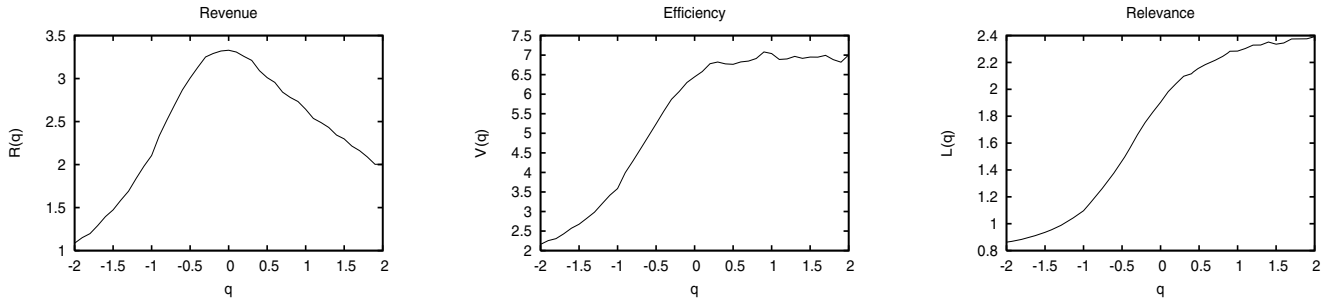

Figure 3: Revenue, efficiency, and relevance for different parameters q with Spearman correlation of 0.4. Estimated standard errors are less than 1% of the values shown.

We also looked into the effect of introducing a reserve score. Results are shown in Figure 4. Naturally, both efficiency and relevance suffer with an increasing reserve score. The optimal setting is $r = 0.2$, which gives only an 8% increase in revenue from $r = 0$. However, it results in a 13% efficiency loss and a 26% relevance loss. Tuning weights seems to be a much more desirable approach than introducing a reserve score in this instance.

The reason why efficiency and relevance suffer more with a reserve score is that this approach will often exclude bidders entirely, whereas this never occurs when tuning weights. The two approaches are not mutually exclusive, however, and some combination of the two might prove better than either alone, although we did not investigate this possibility.

8. CONCLUSIONS

In this work we looked into the revenue properties of a family of ranking rules that contains the Yahoo! and Google models as special cases. In practice, it should be very simple to move between rules within the family: this simply involves changing the exponent q applied to advertiser effects. We also showed that, in principle, the same effect could be obtained by using bidding credits. Despite the simplicity of the rule change, simulations revealed that properly tuning q can significantly improve revenue. In the simulations, the revenue improvements were greater than what could be obtained using reserve prices.

On the other hand, we showed that advertiser satisfaction and user experience could suffer if q is made too small. We proposed that the auctioneer set bounds on the decrease in advertiser and user satisfaction he is willing to tolerate, which would imply bounds on the range of allowable q. With appropriate estimates for the distributions of value and relevance, and knowledge of their correlation, the revenue curve can then be plotted within this range to locate the optimum.

There are several ways to push this research further. It would be interesting to do this analysis for a variety of keywords, to see if the optimal setting of q is always so sensitive to the level of correlation. If it is, then simply using rank-by-bid where there is positive correlation, and rank-by-revenue where there is negative correlation, could be fine to a first approximation and already improve revenue. It would also be interesting to compare the effects of tuning q versus reserve pricing for keywords that have few bidders. In this instance reserve pricing should be more competitive, but this is still an open question.

In principle the minimum revenue in Nash equilibrium can be found by linear programming. However, many allocations can arise in Nash equilibrium, and a linear program needs to be solved for each of these. There is as yet no efficient way to enumerate all possible Nash allocations, so finding the minimum revenue is currently infeasible. If this problem could be solved, we could run simulations for Nash equilibrium instead of symmetric equilibrium, to see if our insights are robust to the choice of solution concept.

Larger classes of ranking rules could be relevant. For instance, it is possible to introduce discounts d_s and rank according to $w_s b_s - d_s$; the equilibrium analysis generalizes to this case as well. With this larger class the virtual score can equal the score, e.g. in the case of a uniform marginal distribution over values. It is unclear, though, whether such extensions help with more realistic distributions.

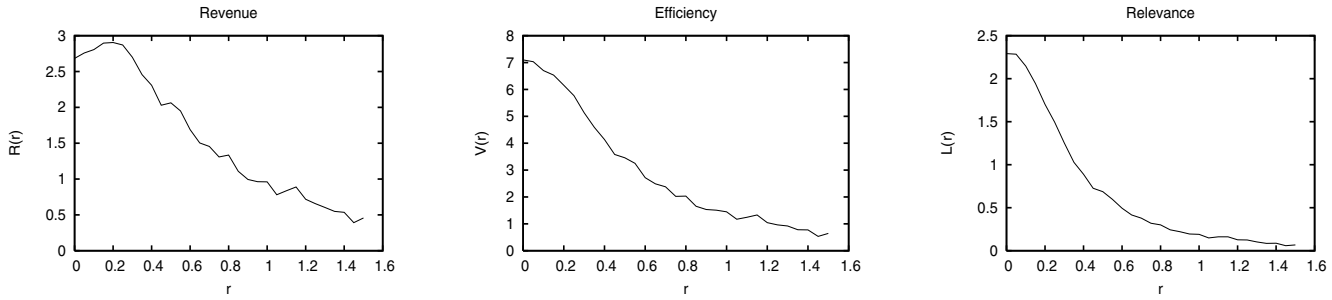

Figure 4: Revenue, efficiency, and relevance for different reserve scores r, with Spearman correlation of 0.4 and $q = 1$. Estimates are averaged over 1000 samples.

Acknowledgements

We thank Pavel Berkhin, Chad Carson, Yiling Chen, Ashvin Kannan, Darshan Kantak, Chris LuVogt, Jan Pedersen, Michael Schwarz, Tong Zhang, and other members of Yahoo! Research and Yahoo! Search Marketing.

9. REFERENCES

[1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.

[2] T. Börgers, I. Cox, M. Pesendorfer, and V. Petricek. Equilibrium bids in auctions of sponsored links: Theory and evidence. Working paper, November 2006.

[3] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the Generalized Second Price auction: Selling billions of dollars worth of keywords. *American Economic Review*, forthcoming.

[4] 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*, forthcoming.

[5] G. Iyengar and A. Kumar. Characterizing optimal keyword auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions*, Ann Arbor, MI, 2006.

[6] D. Knuth. *The Art of Computer Programming*, volume 3. Addison-Wesley, 1997.

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

[8] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.

[9] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1), February 1981.

[10] R. B. Nelsen. *An Introduction to Copulas*. Springer, 2006.

[11] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, forthcoming.