

Position Auctions with Consumer Search*

Susan Athey and Glenn Ellison

March 2008

Abstract

This paper examines a model in which advertisers bid for “sponsored-link” positions on a search engine. The value advertisers derive from each position is endogenized as coming from sales to a population of consumers who make rational inferences about firm qualities and search optimally. Consumer search strategies, equilibrium bidding, and the welfare benefits of position auctions are analyzed. Implications for reserve prices and a number of other auction design questions are discussed.

JEL Classification No.: D44, L86, M37

1 Introduction

Google, Yahoo! and Microsoft allocate the small “sponsored links” one sees at the top and on the right side of their search engine results via similar auction mechanisms. In just a few years, this has become one of the most practically important applications of auctions, with annual revenues surpassing \$10 billion. Several analyses of these “position auctions” have now appeared in the economics and computer science literature. The most fundamental result is that the benchmark auction in which the k^{th} highest bidder wins the k^{th} slot and pays the $k + 1^{st}$ highest bid is not equivalent to the VCG mechanism and thus does not induce truthful bidding, but does result in the same revenue as the VCG mechanism.¹ Other authors have explored variants and discussed click-through-weighting, budget constraints, click-fraud, and other issues.

The literature has focused on position auctions as a topic in auction theory. Most papers

*Harvard University and Massachusetts Institute of Technology. E-mail: athey@fas.harvard.edu, g Ellison@mit.edu. We thank Ben Edelman, John Morgan, Ilya Segal, Hal Varian, and various seminar audiences for their comments and Eduardo Azevedo for exceptional research assistance. This work was supported by NSF grants SES-0550897 and SES-0351500 and the Toulouse Network for Information Technology.

¹Edelman, Ostrovsky, and Schwarz (EOS) (2007), Aggarwal, Goel, and Motwani (2006), and Varian (2007) all contain versions of this result.

abstract away from the fact that the “objects” being auctioned are advertisements.² We feel that this is an important omission because when the value of a link is due to consumers’ clicking on the links and making purchases, it is natural to assume that consumer behavior and link values will be affected by the process by which links are selected for display. In this paper, we develop this line of analysis, present a number of results on bidding behavior, consumer search, and welfare, and address a number of auction-design topics. These include reserve prices, click-through weighting, fostering product diversity, obfuscation, search diversion, the effect of different payment schemes (e.g. pay per click vs. pay per action), and the possibility of using multi-stage auction mechanisms.

Section 2 of the paper presents our base model. The most important assumptions are that advertisers differ in quality (with high quality firms being more likely to meet each consumer’s need), that consumers incur costs of clicking on ads, and that consumers act rationally in deciding how many ads to click on and in what order. Section 2 also presents some basic results on search and welfare: we characterize optimal consumer search strategies and illustrate the welfare benefits that result when well sorted sponsored-link lists make consumer search more efficient.

Section 3 contains our analysis of the sponsored-link auction. Because the value of being in any given position on the search screen depends on the qualities of all of the other advertisers, the auction is not a private-values model and does not fit within the framework of Edelman, Ostrovsky and Schwarz (2007). But the analysis and equilibrium are nonetheless similar. We derive a symmetric perfect Bayesian equilibrium with monotone bidding functions and discuss some properties. Section 4 discusses reserve-price policies. A fundamental difference between our model and models with exogenous click-through rates is that reserve prices may enhance both auctioneer revenue and social welfare. Although reserve prices prevent some firms from being listed and thereby prevent some welfare enhancing purchases from occurring, they can enhance welfare in two ways: they help consumers avoid some of the inefficient search costs they incur when clicking on low quality links; and can increase the number of links that are examined in equilibrium. We present several results, including a demonstration that in one special case there is an alignment of private and socially optimal policies.

Section 5 examines click-weighted auctions similar to those used by Google, Yahoo! and Microsoft. We note that some optimality arguments do not carry over to this setting, and

²Chen and He (2006) is a noteworthy exception – they develop a model with optimal consumer search and note that the fact that auctions lead to a sorting of advertisers by quality can rationalize top-down search and be a channel through which sponsored link auctions contribute to social welfare.

consumer search can be less efficient. We discuss more complicated bidding mechanisms that might be used to address this concern. We also discuss a variant of the model with pay-per-action bids, and show that this can lead to equivalent outcomes.

Section 6 discusses a few more auction design topics: the impact of search-diverting sites; consumer uncertainty about search engine quality; and obfuscation by advertisers that may impact click-through rates.

Our paper contributes to a growing literature. Edelman, Ostrovsky, and Schwarz (2007), Aggarwal, Goel, and Motwani (2006), and Varian (2007), all contain versions of the result that the standard unweighted position auction (which EOS call the generalized second price or GSP auction) is not equivalent to a VCG mechanism but can yield the same outcome in equilibrium. Such results can be derived in the context of a perfect information model under certain equilibrium selection conditions. EOS show they the equivalence can also be derived in an incomplete information ascending bid auction, and that in this case the VCG-equivalent equilibrium is the unique perfect Bayesian equilibrium. The papers also note conditions under which results would carry over to click-weighted auctions.

The above papers consider environments in which the value to advertiser i of being in position j on the screen is the product of a per-click value v_i and a position-specific click-through rate c_j . Borgers, Cox, Pesendorfer and Petricek (BCPP) (2006) extend this model to allow click-through rates and value per click to vary across positions in different ways for different advertisers and emphasize that there can be a great multiplicity of equilibrium outcomes in a perfect information setting.³ Our model does not fit in the BCPP framework either, however, because they maintain the assumption that advertiser i 's click-through rate in position j is independent of the characteristics of the other advertisers.

Chen and He (2006) is more closely related to our paper. They had previously introduced a model that is like ours in several respects. Advertisers are assumed to have different valuations because they have different probabilities of meeting consumers' needs. Consumers search optimally until their need is satisfied. They also include some desirable elements which we do not include: they endogenize the prices advertisers charge consumers; and allow firms to have different production costs.⁴ We are enthusiastic about these assumptions as introducing important issues and think that the observations they make about the basic economics of their model (and ours) are interesting and important: they note that

³BCPP also contains an empirical analysis which includes methodological innovations and estimates of how value-per-click changes with position in Yahoo! data.

⁴As in Diamond (1971) the equilibrium turns out to be that all firms charge the monopoly price.

sponsored links provide a welfare benefit by directing consumer search and making it more efficient; and note that without some refinement there will also be equilibria in which consumers pay no attention to the order of sponsored links and the links are therefore worthless.

Our work goes beyond theirs in several ways. They only consider what happens for one particular realization of firm qualities, whereas we assume the qualities are independently drawn from a distribution and examine an incomplete information game. They assume that consumers know the (unordered) set of realized qualities, so there is no updating about the quality of the set of advertisers as consumers move down the list. They have no heterogeneity in consumer search costs and obviate the search duration problem by assuming that search costs are such that all consumers will search all listed firms. As a consequence, most of our paper addresses issues that don't come up in their framework. They have no analog to our derivation of optimal consumer search strategies because their consumers simply click on all sponsored links (if necessary). They have no analog to our derivation of the equilibrium strategies in an asymmetric information bidding game – they compute Nash equilibria of a game where firms know each others values. And they do not discuss reserve prices or the many other auction design issues we consider. (Many of the auction design questions hinge on how the design affects the information consumers get about firm qualities and thereby influences consumer search, which is not something that comes up when search costs are such that all consumers view all ads.)

2 A Base Model

A continuum of consumers have a “need”. They receive a benefit of 1 if the need is met. To identify firms able to meet the need they visit a search site. The search site displays M sponsored links. Consumer j can click on any of these at cost s_j . Consumers click optimally until their need is met or until the expected benefit from an additional click falls below s_j . We assume the s_j have an atomless distribution G with support on $[0, 1]$.

N advertisers wish to advertise on a website. Firm i has probability q_i of meeting each consumer's need, which is private information. We assume that all firms draw their q_i independently from a common distribution, F , which is atomless and has support $[0, 1]$. Advertisers get a payoff of 1 every time they meet a need.

Informally, we follow EOS in assuming that the search site that conducts an ascending bid auction for the M positions: if the advertisers drop out at per-click bids b_1, \dots, b_N , the search engine selects the advertisers with the M highest bids and lists them in order from

top to bottom. The k^{th} highest bidder pays the $k + 1^{\text{st}}$ highest bid for each click it gets.⁵ To avoid some of the complications that arise in an infinite horizon continuous time models, however, we formalize the auction as a simpler M -stage game in which the firms are simply repeatedly asked to name the price at which they will next drop out if no other firm has yet dropped out.⁶ In the first stage, which we call stage $M + 1$, the firms simultaneously submit bids $b_{M1}, \dots, b_{MN} \in [0, \infty)$ specifying a per-click price they are willing to pay to be listed on the screen. The $N - M$ lowest bidders are eliminated.⁷ Write b^{M+1} for the highest bid among the firms that have been eliminated. In remaining stages k , which we'll index by the number of firms remaining, $k \in \{M, M - 1, \dots, 2\}$, the firms which have not yet been eliminated simultaneously submit bids $b_{kn} \in [b^{k+1}, \infty)$. The firm with the single lowest bid is assigned position k and eliminated from future bidding. We define b^k to be the bid of this player. At the end of the auction, the firms in positions $1, 2, \dots, M$ will make per-click payments of b^2, b^3, \dots, b^{M+1} for the clicks they receive.

Before proceeding, we pause to mention the main simplifications incorporated in the model. First, advertisers are symmetric except for their probability of meeting a need: profit-per-action is the same for all firms. Generalizing this would allow us to distinguish between the externality a firm creates on others by being higher on the list, which is related to the probability of meeting the need, and the value the firm gets from being in a position. Related to this, we could also allow firms to have value to being in a position due to impressions (not clicks). Incorporating such impression values (as in BCPP) would also put a wedge between the externality created by a firm and its value to being in a given position. Finally, consumers get no information about whether listed firms are more or less likely to meet their needs from reading the text of their ads. It would be more realistic to assume that firms are heterogeneous in a way that is recognized by consumers before clicking. We consider one extension along these lines in Section 5.

⁵Note that this model differs from the real-world auctions by Google, Yahoo!, and MSN in that it does not weight bids by clickthrough weights. We discuss such weighted auctions in Section 5. We present results first for the unweighted auction because the environment is easier to analyze. It should also be an approximation to real-world auctions in which differences in click-through rates across firms are minor, e.g. where the bidders are retailers with similar business models.

⁶Two examples of issues we avoid dealing with in this way are formalizing a clock-process in which firms can react instantaneously to dropouts and specifying what happens if two or more firms never drop out. See Demange, Gale and Sotomayor (1986) for more on extensive form specifications of multi-unit auctions.

⁷If two or more firms are tied for the M^{th} highest bid, we assume that the tie is broken randomly with each tied firm being equally likely to be eliminated.

2.1 Consumer welfare with sorted lists

One of the main ideas that we wish to bring out in our model is that one channel through which sponsored-link auctions affect consumer utility is through their effect on the efficiency of consumer search. We introduce this idea in this section by characterizing consumer welfare with sorted and unsorted lists. This section also contains important building blocks for all of our analyses: an analysis of the Bayesian updating that occurs whenever consumers find that a particular link does not meet their needs; and a derivation of optimal search strategies.

A benchmark for comparison is what happens if the advertisements are presented to consumers in a random order. Define $\bar{q} = E[q_i]$. In that case, the consumer expects each website to meet the need with probability \bar{q} .

Proposition 1 *If the ads are sorted randomly, then consumers with $s > \bar{q}$ don't click on any ads. Consumers with $s < \bar{q}$ click on ads until their need is met or they run out of ads. Expected consumer surplus is*

$$E(CS(s)) = \begin{cases} 0 & \text{if } s \geq \bar{q} \\ (\bar{q} - s) \frac{1 - (1 - \bar{q})^M}{\bar{q}} & \text{if } s < \bar{q} \end{cases}$$

Proof: The clicking strategies are obvious. Consumers who are willing to search get $(\bar{q} - s)$ from the first search. If this is unsuccessful (which happens with probability $(1 - \bar{q})$) they get $(\bar{q} - s)$ from their second search. The total payoff is $(\bar{q} - s)(1 + (1 - \bar{q}) + (1 - \bar{q})^2 + \dots + (1 - \bar{q})^{M-1})$.

QED

Suppose now that the bidding model has an equilibrium in which strategies are strictly monotone in q . Then, in equilibrium the firms will be sorted so that the firm with the highest q is on top. Consumers know this, so the expected utility from clicking on the top firm is the highest order statistic, $q^{1:N}$.⁸ Their expected payoff for any additional clicks must be determined by Bayesian updating: the fact that the first website didn't meet their needs makes them reduce their estimate of its quality and of all lower websites' qualities.

Let $q^{1:N}, \dots, q^{N:N}$ be the order statistics of the N firms' qualities and let z^1, \dots, z^N be Bernoulli random variables equal to one with these probabilities. Let \bar{q}_k be the expected quality of website k in a sorted list, given that the consumer has failed to fulfill his need

⁸ As is Ellison, Fudenberg, and Möbius (2004) we write $q^{1:N}$ for the highest value, in contrast to the usual convention in statistics, which is to call the highest value the N^{th} order statistic.

from the first $k - 1$ advertisers:

$$\bar{q}_k = E(q^{k:N} | z^1 = \dots = z^{k-1} = 0).$$

Proposition 2 *If the firms are sorted by quality in equilibrium, then consumers follow a top-down strategy: they start at the top and continue clicking until their need is met or until the expected quality of the next website is below the search cost: $\bar{q}_k < s$. The numbers \bar{q}_k are given by*

$$\begin{aligned} \bar{q}_k &= E(q^{k:N} | z^1 = \dots = z^{k-1} = 0) \\ &= \frac{\int_0^1 x f^{k:N}(x) \text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} dx}{\int_0^1 f^{k:N}(x) \text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} dx} \end{aligned}$$

A firm in position k will receive $(1 - q^{1:N}) \dots (1 - q^{k-1:N}) G(\bar{q}_k)$ clicks.

Proof: Consumers search in a top-down manner because the likelihood of that a site meets a consumer's need is consumer-independent, and hence maximized for each consumer at the site with the highest q . A consumer searches the k^{th} site if and only if the probability of success at this site is greater than s . The expected payoff to a consumer from searching the k^{th} site conditional on having gotten failures from the first $k - 1$ is $E(q^{k:N} | z^1 = \dots = z^{k-1} = 0)$. (The $f^{k:N}$ in the formula is the PDF of the k^{th} order statistic of F .)

QED

One feature of our model which we'll exploit at several points is that the special case with uniformly-distributed valuations is surprisingly tractable: fairly simple closed-form expressions can be given for many of the terms that come up in the propositions. We present these formulas at several points to help build intuition and to present more explicit results.

Corollary 1 *If in addition the distribution F of firm qualities is uniform, then consumer with search cost s stops clicking when she reaches position $k^{\text{max}}(s)$, where*

$$k^{\text{max}}(s) = \lceil \frac{1-s}{1+s} N + \frac{1}{1+s} \rceil.$$

Proof: When F is uniform, $f^{k:N}(x) = \frac{N!}{(N-k)!(k-1)!} (1-x)^{k-1} x^{N-k}$. Then the two probabilities from the previous expression are:

$$\text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} = \left(\frac{1-x}{2} \right)^{k-1}.$$

$$\text{Prob}\{z^1 = \dots = z^{k-1} = 0\} = \int_0^1 \left(\frac{1-x}{2}\right)^{k-1} f^{k:N}(x) dx$$

Note that both the numerator and the denominator in the expression we're evaluating are equal to a constant times an integral of the form $\int_0^1 x^a(1-x)^b dx$. Integrating by parts one can show that this is equal to $a!b!/(a+b+1)!$. Evaluating the integrals gives

Lemma 1 *For uniform F , if consumers search an ordered list from the top down, then*

$$\begin{aligned} E(q^{k:N} | z^1 = \dots = z^{k-1} = 0) &= \frac{N+1-k}{N+k} \\ \text{Prob}\{z^1 = \dots = z^{k-1} = 0\} &= \prod_{j=1}^{k-1} \frac{2j-1}{N+j} \end{aligned}$$

The second part of the lemma can also be proved more quickly by noting that

$$\begin{aligned} \text{Prob}\{z^1 = \dots = z^{k-1} = 0\} &= \text{Prob}\{z^1 = 0\} \text{Prob}\{z^2 = 0 | z^1 = 0\} \dots \\ &= \prod_{j=1}^{k-1} (1 - E(q^{j:N} | z^1 = \dots = z^{j-1} = 0)) \end{aligned}$$

The consumer will want to search the k^{th} website if $(N+1-k)/(N+k) > s$. This holds for $k < k^{\text{max}}(s)$.

QED

Assuming that the quality distribution is uniform also makes it easy to compute expected consumer surplus. The expected payoff from clicking on the top link is $E(q^{1:N}) - s = N/(N+1) - s$. If the first link is unsuccessful, which happens with probability $1/(N+1)$, then (using Lemma 1) the consumer gets utility $E(q^{2:N} | z^1 = 0) - s = (N-1)/(N+2) - s$ from clicking on the second. Adding up these payoffs over the number of searches that will be done gives

Proposition 3 *If the distribution of firm quality F is uniform, the expected utility of a consumer with search cost s is:*

$$E(CS(s)) = \begin{cases} 0 & \text{if } s \in \left[\frac{N}{N+1}, 1\right] \\ \frac{N}{N+1} - s & \text{if } s \in \left[\frac{N-1}{N+2}, \frac{N}{N+1}\right] \\ \frac{N}{N+1} - s + \frac{1}{N+1} \left(\frac{N-1}{N+2} - s\right) & \text{if } s \in \left[\frac{N-2}{N+3}, \frac{N-1}{N+2}\right] \\ \frac{N}{N+1} - s + \frac{1}{N+1} \left(\frac{N-1}{N+2} - s\right) + \frac{1}{N+1} \frac{3}{N+2} \left(\frac{N-2}{N+3} - s\right) & \text{if } s \in \left[\frac{N-3}{N+4}, \frac{N-2}{N+3}\right] \\ \dots & \\ 1 - 1/2^M & \text{if } s \approx 0 \end{cases}$$

When N is large the graph of the function above approaches $1 - s$ whereas the unordered payoff is approximately $1 - 2s$. N doesn't need to be very large at all for the function to be close to its limiting value. For example, just looking at the first term we know that for $N = 5$ we have $E(CS(s)) > 5/6 - s$ for all s . The figure below graphs the relationship between $E(CS)$ and s for $N = 4$.




Figure 1: Consumer surplus with sorted and unsorted links: $N = 4$

3 Equilibrium Analysis

In this section we solve for the equilibrium of our base model taking both consumer and advertiser behavior into account. Advertisers' bids are influenced by click-through rates, so we start with an analysis of consumer behavior. We then analyze the bidding among advertisers.

We restrict our attention to equilibria in which advertisers' bids are monotone increasing in quality, so that consumers expect the list of firms to be sorted from highest to lowest quality and search in a top-down manner.⁹

3.1 Clickthrough rates with uniform distributions

Clickthrough rates (CTRs) are more complicated in our model than in the standard position auction model because the number of clicks that a firm receives depends not only on its

⁹In a model with endogenous search there will also be other equilibria. For example, if all remaining bidders drop out immediately once M firms remain and are ordered arbitrarily by an auctioneer that cannot distinguish among them, then consumers beliefs will be that the ordering of firms is meaningless, so it would be rational for consumers to ignore the order in which the firms appear and for firms to drop out of the bidding as soon as possible.

position on the list, but also on consumer beliefs about its quality and on the realized qualities of other firms. In this section we characterize both unconditional CTRs and conditional CTRs taking into account what a firm infers about the firm-quality distribution when a change to its bid is pivotal.

In a search model, the clicks received by the k^{th} firm is decreasing in k for two reasons: some consumers will have already met their need before getting the k^{th} position on the list; and a lower position signals to consumers that the firm's quality is lower, which reduces the number of consumers willing to click on the link. When consumer search costs are uniformly distributed, the probability that a consumer whose needs have not been met by the first $k - 1$ websites will click on the k^{th} website is just the expected quality of the k^{th} website conditional on the consumer having had $k - 1$ unsuccessful experiences, which we derived in the previous section.

Proposition 4 *Assume s and $q \sim U[0, 1]$. Write $D(k)$ for the ex ante expected clicks received by the k^{th} website and $D(k, q)$ for the number of clicks a website of quality q expects to receive if it plays the equilibrium strategy and ends up in the k^{th} position. We have*

$$\begin{aligned} D(k, q) &= \left(\frac{1+q}{2}\right)^{k-1} \frac{N+1-k}{N+k} \\ D(k) &= \frac{1 \cdot 3 \cdot \dots \cdot (2k-3)}{(N+1)(N+2)\dots(N+k-1)} \frac{N+1-k}{N+k} \end{aligned}$$

Proof: The first expression is derived by noting that the firm will receive a click only if all higher firms do not meet the consumer's need and the consumer will decide to click on site k if he or she gets that far. The probability that site j will be unsuccessful for the consumer conditioning on $q_j > q$ is $1 - E(q_j | q_j > q) = (1+q)/2$.¹⁰ The probability that all $k - 1$ clicks will be unsuccessful is $((1+q)/2)^{k-1}$. The probability that the consumer would click on the k^{th} site is

$$\begin{aligned} \text{Prob}\{s < E(q^{k:N} | z^1 = \dots = z^{k-1} = 0)\} &= E(q^{k:N} | z^1 = \dots = z^{k-1} = 0) \\ &= (N+1-k)/(N+k) \end{aligned}$$

by Lemma 1.

The second expression is simply the expression for $\text{Prob}\{z^1 = \dots = z^{k-1} = 0\}$ in Lemma 1 multiplied by the consumer's conditional expectation for $q^{k:N}$.

QED

¹⁰Note that the j in this expression is a generic index and does not denote the j^{th} highest value.

3.2 Equilibrium in the bidding game

Consider now our formalization of an “ascending auction” in which the N firms bid for the $M < N$ positions.

Note that conditional on being clicked on, a firm will be able to meet a consumer’s need with probability q . We’ve exogenously fixed the per-consumer profit at one, so q is like the value of a click in a standard position auction model.

Although one can think of our auction game as being like the EOS model with endogenous click-through rates, the auction part of our model cannot be made to fit within the EOS framework. The reason is that the click-through rates are a function of the bidders’ types as well as of the positions on the list.¹¹ The equilibrium derivation, however, is similar to that of EOS.

Our first observation is that, as in that model, firms will bid up to their true value to get on the list, but will then shade their bids in the subsequent bidding for higher positions on the list.

In the initial stage (stage $M + 1$) when $N > M$ firms remain, firms will get zero if they are eliminated. Hence, for a firm with quality q it is a weakly dominant strategy to bid q . We assume that all bidders behave in this way.

Once firms are sure to be on the list, however, they will not want to remain in the bidding until it reaches their value. To see this, suppose that k firms remain and the $k + 1^{st}$ firm dropped out at b^{k+1} . As the bid level b approaches q a firm knows that it will get $q - b^{k+1}$ per click if it drops out now. If it stays in and no one else drops out before b reaches q nothing will change. If another firm drops out at $q - \epsilon$, however, the firm would do much worse: it will get more clicks, but its payoff per click will just be $q - (q - \epsilon) = \epsilon$. Hence, the firm must drop out before the bid reaches its value.

Assume for now that the model has a symmetric strictly monotone equilibrium in which drop out points $b^*(k, b^{k+1}; q)$ are only a function of (1) the number of firms k that remain; (2) the current $k + 1^{st}$ highest bid, b^{k+1} ; and (3) the firm’s privately known quality q .¹²

Suppose that the equilibrium is such that a firm will be indifferent between dropping out at $b^*(k, b^{k+1}; q)$ and remaining in the auction for an extra db and then dropping out at $b^*(k, b^{k+1}; q) + db$. This change in the strategy does not affect the firm’s payoff if no other firm drops out in the db bid interval. Hence, to be locally indifferent the firm must

¹¹BCPP have a more general setup, but they still assume that click-through rates do not depend on the types of the other bidders.

¹²In principle, drop out points could condition on the history of drop out points in other ways. One can set $b^{k+1} = 0$ in the initial stage when no firm has yet to drop out.

be indifferent between remaining for the extra db conditional on having another firm drop out at $b^*(k, b^{k+1}; q)$. In this case the firm's expected payoff if it is the first to drop out is

$$E \left((1 - q^{1:N})(1 - q^{2:N}) \cdots (1 - q^{k-2:N})(1 - q) | q^{k-1:N} = q \right) \cdot G(\bar{q}_k) \cdot (q - b^{k+1}).$$

The first term in this expression is the probability that all higher websites will not meet a consumer's need. The second is the demand term coming from the expected quality. The third is the per-click profit. If the firm is the second to drop out in this db interval then its payoff is

$$E \left((1 - q^{1:N})(1 - q^{2:N}) \cdots (1 - q^{k-2:N}) | q^{k-1:N} = q \right) \cdot G(\bar{q}_{k-1}) \cdot (q - b^*).$$

The first two terms in this expression are greater reflecting the two mechanisms by which higher positions lead to more clicks. The final is smaller reflecting the lower markup. Indifference gives

$$G(\bar{q}_k)(1 - q)(q - b^{k+1}) = G(\bar{q}_{k-1})(q - b^*)$$

This can be solved for b^* .

Proposition 5 *The auction game has a symmetric strictly monotone pure strategy equilibrium. In particular, it is a Perfect Bayesian equilibrium for firms to choose their dropout points according to*

$$b^*(k, b^{k+1}, q) = \begin{cases} q & \text{if } k > M \\ b^{k+1} + (q - b^{k+1}) \left(1 - \frac{G(\bar{q}_k)}{G(\bar{q}_{k-1})}(1 - q) \right) & \text{if } k \leq M. \end{cases}$$

When both qualities and search costs are uniform, the latter expression is

$$b^*(k, b^{k+1}; q) = b^{k+1} + (q - b^{k+1}) \left(1 - (1 - q) \left(1 - \frac{2N + 1}{(N + 1)^2 - (k - 1)^2} \right) \right).$$

Sketch of proof: First, it is easy to show by induction on k that the strategies defined in the proposition are symmetric strictly monotone increasing and always have $q_i \geq b^{k+1}$ on the equilibrium path. The calculations above establish that the given bidding functions satisfy a first-order condition.

To show that the solution to the first-order condition is indeed a global best response we combine a natural single-crossing property of the payoff functions – the marginal benefit of a higher bid is greater for a higher quality firm – and the indifference on which the bidding

strategies are based. For example, we can show that the change in profits when a type q' bidder increases his bid from $b^*(q')$ to $b^*(\hat{q})$ is negative using

$$\begin{aligned}\pi(b^*(q'); q') - \pi(b^*(\hat{q}); q') &= \int_{q'}^{\hat{q}} \frac{\partial \pi}{\partial b}(b^*(q), q') \frac{db^*}{dq}(q) dq \\ &\leq \int_{q'}^{\hat{q}} \frac{\partial \pi}{\partial b}(b^*(q), q) \frac{db^*}{dq}(q) dq \\ &= 0.\end{aligned}$$

Formalizing this argument is a little tedious, so we have left it to the appendix.

The expression for the uniform distribution is obtained by substituting the expression for $G(\bar{q}_k)$ from Lemma 1 into the general formula and simplifying:

$$\frac{G(\bar{q}_k)}{G(\bar{q}_{k-1})} = \frac{(N+1-k)/(N+k)}{(N+1-(k-1))/(N+(k-1))} = 1 - \frac{2N+1}{(N+1)^2 - (k-1)^2}.$$

QED

Remarks

1. In this equilibrium firms start out bidding up to their true value until they make it onto the list. Once they make it onto the list they start shading their bids. If q is close to one, then the bid shading is very small. When q is small, in contrast, bids increase slowly with increases in a firm's quality because there isn't much gain from outbidding one more bidder.
2. The strategies have the property that when a firm drops out of the final M it is common knowledge that no other firm will drop out for a nonzero period of time.
3. Bidders shade their bids more when bidding for higher positions, i.e. $b^*(k, b'; q)$ is increasing in k with b' and q fixed, if and only if $\frac{G(\bar{q}_{k-1})}{G(\bar{q}_k)}$ is increasing in k . The expression in the proposition makes clear that this is true when the distribution of q is uniform. One may get some intuition for whether the condition is likely to hold in practice by examining the growth in click-through rates as a firm moves from position k to position $k-1$. In the model this is $\frac{G(\bar{q}_{k-1})}{(1-q^{k-1:N})G(\bar{q}_k)}$. BCPP report that this is "about 1.5 for the top positions." If this ratio is roughly position-independent despite $q^{k:N}$ is declining in k , then $\frac{G(\bar{q}_{k-1})}{G(\bar{q}_k)}$ would have to be increasing in k .
4. EOS show that the similar equilibrium of their model is unique among equilibria in strategies that are continuous in types and note that there are other equilibria that are discontinuous in types. The indifference condition we derive should imply that

equilibrium is also unique in our model if one restricts attention to an appropriate class of strategies with continuous strictly monotone bidding functions. As Chen and He (2006) also note, however, there are also other equilibria. For example, if consumers believe that the links are sorted randomly (and therefore search in a random order), then there will be an equilibrium in which all firms drop out as soon as M firms remain.

4 Reserve Prices

In this section we note that the profit-welfare tradeoffs that arise when considering reserve prices are very different in our model as compared to standard auction models. In a standard auction model reserve prices increase the auctioneer's expected revenues. At the same time, however, they reduce social welfare. Hence, they could inhibit seller or buyer entry in a model in which these were endogenous, and might not be optimal in such models.¹³ Edelman and Schwarz (2006) provide theoretical and numerical analyses of the impact of reserve prices in the position auction model and show that the GSP auction with an optimally chosen reserve price is an optimal mechanism. Here, we show that the considerations are somewhat different in our model: reserve prices can increase both the profits of the auctioneer and social welfare. The reason for this difference is that consumers incur search costs on the basis of their expectation of firm quality. When the quality of a firm's product is low relative to this expectation, the search costs consumers incur are inefficient. By instituting a reserve price, the auctioneer commits not to list products of sufficiently low quality and can reduce this source of welfare loss. This in turn, can increase the number of searches that consumers are willing to carry out. Increases in the volume of trade are another channel through which welfare can increase.

We first discuss the special case of uniformly distributed search costs, which yields neat results. We then present some results for the more general case.

4.1 Reserve prices when search costs are uniformly distributed

In this subsection, we examine the special case of our model in which the search costs distribution G is uniform on $[0, 1]$. We restrict our analyses to equilibria like those described in the previous section in which firms use strictly monotone bidding strategies and bid their

¹³See Ellison, Fudenberg and Möbius (2004) for more on a competing auction model in which this effect would be important.

true value when they will not be on the sponsored-link list. We initially consider a general quality distribution F , and later give some results for F uniform.

4.1.1 The optimal reserve price for consumers also maximizes social welfare

In this section we present a striking result on the alignment of consumer and advertiser/search engine preferences: the welfare maximizing and consumer surplus maximizing policies coincide. Moreover, for any reserve price, the sum of advertiser profit and search engine profit is twice the consumer surplus. This occurs because producer surplus is directly related to the probability that consumers have their needs satisfied and because consumers search optimally and have uniformly distributed search costs.

Proposition 6 *Suppose the distribution of search costs is uniform. Consumer surplus and social welfare are maximized for the same reserve price. Given any bidding behavior by advertisers and any reserve price policy of the search engine, equilibrium behavior by consumers implies $E(W) = 3E(CS)$.*

Proof: Write GCS for the gross consumer surplus in the model: $GCS = CS + \text{Search Costs}$. Write GPS for the gross producer surplus: $GPS = \text{Advertiser Profit} + \text{Search-engine fees}$. Because a search produces one unit of GCS and one unit of GPS if a consumer need is met and zero units of each otherwise we have $E(GCS) = E(GPS)$.

Welfare is given by $W = GCS + GPS - \text{Search Costs}$. Hence, to prove the theorem we only need to show that $E(\text{Search Costs}) = \frac{1}{2}E(GCS)$. This is an immediate consequence of the optimality of consumer search and the uniform distribution of search costs: each ad is clicked on by all consumers with $s \in [0, E(q|X)]$ who have not yet had their needs met, where X is the information available to consumers at the time the ad is presented. Hence, the average search costs expended are exactly equal to one-half of the expected GCS from each click.

QED

Remark

1. Note that the alignment result is fully general in the dimension of not requiring any assumptions on the distribution F of firm qualities.

Our next result is a corollary that follows neatly from the above result: we show that the socially optimal reserve price coincides with the reserve price that would be chosen by a consumer-surplus maximizing search engine both when commitment to a reserve price is

possible and when it is not. This result is of interest for two distinct reasons. First, it is plausible to think that real-world search engines may use consumer-surplus maximization as an objective function. Search engines are in dynamic competition to attract consumers and to the extent that the future is very important (and foregone profits small), designing the search engine to maximize consumer surplus may be a good rule-of-thumb approximation to the optimal dynamic policy. Second, the observation about the no-commitment model turns out to be a useful computational tool – it can be easier to find the equilibrium the no commitment model than to find the maximizer of the social welfare function.

Corollary 2 *Suppose the distribution of search costs is uniform. Suppose that reserve price r^W maximizes social welfare when the search engine has the ability to commit to a reserve price. Then, r^W is an equilibrium choice for a consumer-surplus maximizing search engine regardless of whether the search engine has the ability to commit to a reserve price.*

Proof: The coincidence of the two reserve prices with commitment is an immediate consequence of Proposition 6.

To see that the socially optimal reserve price is also an equilibrium outcome when the search engine lacks commitment power, write $CS(q, q')$ for the expected consumer surplus if consumers believe that the search engine displays a sorted list of all advertisers with quality at least q , but the search engine actually displays all advertisers with quality at least q' . The optimality of consumer search behavior implies $CS(q, q') \leq CS(q', q')$. The assumption that advertisers play an equilibrium with strictly monotone strategies for any reserve price and that r^W is the socially optimal reserve price imply that for any q' we have $CS(q', q') \leq CS(r^W, r^W)$. Any deviation by the firm to a different reserve price yields consumer surplus of $CS(r^W, q')$ for some q' . This cannot improve consumer surplus because $CS(r^W, q') \leq CS(q', q') \leq CS(r^W, r^W)$.

QED

Remark

1. Note that the result is that a consumer-surplus maximizing search engine will not suffer from a lack of commitment power. A social-welfare maximizing search engine would suffer if it lacked commitment power. Proposition 6 shows that consumer surplus and welfare are proportional if consumers have correct beliefs. If the search engine deviates from its equilibrium strategy, then consumers will have incorrect beliefs. Hence, such deviations can increase welfare. Indeed, a deviation to a lower

reserve price would typically be expected to increase welfare because consumers do not internalize the profits that advertisers and the search engine get from their clicks. When a deviation to a slightly lower reserve price results in additional links being displayed it will lead to more clicks and raise welfare. In equilibrium, of course, this incentive cannot exist, so the result will be that the equilibrium reserve price is too low.

4.1.2 Optimal reserve prices with one-position lists

To bring out the economics of setting a socially optimal reserve price, we first consider the simplest version of our model: when the position auction lists only a single firm ($M = 1$). In this case, if the auctioneer commits to a reserve price of r then consumers' expectations of the quality of a listed firm is

$$E(q^{1:N} | q^{1:N} > r) = \frac{\int_r^1 x N F(x)^{N-1} f(x) dx}{\int_r^1 N F(x)^{N-1} f(x) dx}$$

Because consumers with $s \in [0, E(q^{1:N} | q^{1:N} > r)]$ will examine a link if it is presented, the average search cost of searching consumers is $\frac{1}{2}E(q^{1:N} | q^{1:N} > r)$. In the no-commitment model the search engine will only display a link if the net benefit is positive. This implies that it will display links with quality at least $\frac{1}{2}E(q^{1:N} | q^{1:N} > r)$. Equilibrium therefore requires that $r = \frac{1}{2}E(q^{1:N} | q^{1:N} > r)$.

If the quality distribution is also uniform, then the probability that a link is displayed is $1 - r^N$. In this case expected quality is

$$E(q^{1:N} | q^{1:N} > r) = \frac{N}{N+1} \frac{1 - r^{N+1}}{1 - r^N},$$

and the mass of searching consumers is $E(q^{1:N} | q^{1:N} > r)$. Hence, expected consumer surplus is

$$E(CS) = \frac{1}{2} \left(\frac{N}{N+1} \right)^2 \frac{(1 - r^{N+1})^2}{1 - r^N}.$$

Maximizing this expression gives the second part of the proposition below.

Proposition 7 *Suppose that the list has one position and that the distribution of search costs is uniform.*

(i) *Consumer surplus and welfare are maximized for the same reserve price. The optimal r satisfies*

$$r = \frac{1}{2}E(q^{1:N} | q^{1:N} \geq r). \tag{1}$$

(ii) If in addition the distribution F of firm qualities is uniform, then the welfare maximizing reserve price r is the positive solution to $r + r^2 + \dots + r^N = N/(N + 2)$. The welfare-maximizing reserve price is one-third when $N = 1$. It is increasing in N and converges to one-half and $N \rightarrow \infty$.

Remarks

1. Note that the formula (1) applies for any advertiser-quality distribution, not just when advertiser qualities are uniform. Providing a general result is easier here than in some other places because with lists of length one it is not necessary to consider how consumers Bayesian update when links do not meet their needs.
2. The $N = 1$ result for a uniform F follows easily from (1): when $r = 1/3$, consumer expectations will be that $q \sim U[1/3, 1]$, consumers search if and only if $s \in [0, 2/3]$, so the average search costs is indeed $1/3$. The N large result extends to general distributions: $E(q^{1:N} | q^{1:N} > r) \approx 1$ for N large, so the solution has $r \approx \frac{1}{2}$.
3. In the model with a uniform F it is easy to see that consumer surplus, search engine profits and social welfare are all higher with a small positive reserve price than with no reserve price. For example, profits are

$$\pi(r) = \frac{N}{N+1} \frac{1-r^{N+1}}{1-r^N} \int_r^1 \left((r/x)^{N-1} r + \int_r^x (N-1)(z/x)^{N-2} z dz \right) N x^{N-1} dx$$

This is increasing in r for small r .

4.1.3 Optimal reserve price with M position lists

Thinking about the socially optimal reserve price as the equilibrium outcome with a consumer-surplus maximizing search engine is also useful in the full M position model. Holding consumer expectations about the reserve price fixed, making a small change dr to the search-engine's reserve price makes no difference unless it leads to a change in the number of ads displayed. We can again solve for the socially optimal r by finding the reserve price for which an increase of dr that removes an ad from the list has no impact on consumer surplus.

The calculation, however, is more complicated than in the one-position case because there are two ways in which removing a link form the set of links displayed can affect consumer surplus. First, as before there is a change in consumer surplus from consumers who reach the bottom of the list and would have clicked on the final link with $q = r$ if it

had been displayed, but will not click on it if it is not displayed. The benefit from these clicks would have been r . The cost would have been the search cost, which is one-half of the average of the consumers' conditional expectations of q when considering clicking on the final link on the list. Second, not displaying a link at the bottom of the list will reduce consumer expectations about the quality of all higher-up links, and thereby deter some consumers from clicking on these links. Any changes of this second type are beneficial: when the list contains $m < M$ links, consumer expectations when considering clicking on the k^{th} link, $k < m$ are $E(q^{k:N} | z^1 = \dots = z^{k-1} = 0, q^{m:N} > r, q^{m+1:N} < r)$. If the final link is omitted consumer beliefs will change to $E(q^{k:N} | z^1 = \dots = z^{k-1} = 0, q^{m-1:N} > r, q^{m:N} < r)$. This latter belief coincides with $E(q^{k:N} | z^1 = \dots = z^{k-1} = 0, q^{m-1:N} > r, q^{m:N} = r)$. Hence, by not including the marginal link consumers will be made to behave exactly as they would with correct beliefs about the m^{th} firm's quality.

We write $p_m(r)$ for the probability that the m^{th} highest quality is r conditional on one of the M highest qualities being equal to r . With uniformly distributed qualities this is given by

$$p_m(r) = \frac{N \binom{N-1}{m-1} (1-r)^{m-1} r^{N-m}}{\sum_{k=1}^M N \binom{N-1}{k-1} (1-r)^{k-1} r^{N-k}}.$$

The discussion above shows:

Proposition 8 *Suppose the distributions of search costs and firm qualities are uniform. For any N and M , the welfare-maximizing reserve price r is the solution to the first-order condition $\frac{\partial E(CS)}{\partial r} = 0$ with consumer behavior held constant. This reserve price has*

$$r > \frac{1}{2} \left(p_M E(q^{M:N} | q^{M:N} > r) + \sum_{m=1}^{M-1} p_m E(q^{m:N} | q^{m:N} > r, q^{m+1:N} < r) \right).$$

We conjecture that for uniformly distributed qualities and search costs, the optimal reserve price is decreasing in M and converges to $\frac{1}{2}$ in the limit as $N \rightarrow \infty$. We computed the expected consumer surplus numerically for $M = 2$ and $N \in \{2, 3, 4, 5\}$. For $N = 2$, expected consumer surplus is maximized at $r \approx 0.276968$. For $N = 5$, expected consumer surplus is maximized at $r \approx 0.469221$.

4.1.4 More general policies

In the analysis above we considered policies that involved a single reserve price that applies regardless of the number of links that are displayed. A search engine would obviously be

at least weakly better off if it could commit to a policy in which the reserve price was a function of the position. For example, a search engine could have the policy that no ads will be displayed unless the highest bid is at least r_1 , at most one ad will be displayed unless the second-highest bid is at least r_2 , and so on. A rough intuition for how such reserve prices might be set (from largely ignoring effects of the second type) is that they should be set so that the reserve price for the m^{th} position is approximately (but slightly greater than) one-half of consumers' expectations of quality when they are considering clicking on the m^{th} and final link on the list. This suggests that declining reserve prices may be better than a constant reserve price.

The idea of using more general reserve prices illustrates a more general idea: as long as an equilibrium in which advertisers' qualities are revealed still exists, consumer surplus (and hence welfare) is always improved if consumers are given more information about the advertisers' qualities. In an idealized environment, the search engine could report inferred qualities along with each ad. In practice, different positionings might be used to convey this information graphically. One version of this already exists on the major search engines: sponsored links are displayed both on the top of the search page and on the right side. The top positions are the most desired by advertisers, but they are not always filled even when some sponsored links are being displayed on the right side.

4.2 Reserve prices under general distributions

This section studies reserve prices under general assumptions on the distributions of search costs and quality. The consumer optimal and socially optimal reserve prices will no longer exactly coincide when search costs are not uniformly distributed, but informally one would expect that there will typically be some rough alignment. In this section we present one formal result illustrating robustness: we show that the consumer-optimal reserve price is always positive. We also include a couple examples (one of which is clearly extreme) illustrating how things can change: we show that the socially optimal reserve price may be zero; and that the profit-maximizing reserve price can also be zero.

4.2.1 Consumer optimal reserve prices are positive

Our first result is that consumer-surplus maximization requires a positive reserve price. We prove this by showing that consumer surplus is increased when small positive reserve prices are implemented. The intuition for this is that the main effect of such reserve prices have is to eliminate extremely low-quality firms from the sponsored-link list. These websites

provide almost no gross consumer surplus when consumers click on them. Hence, the benefits are always outweighed by the much larger search costs incurred on such clicks.

Proposition 9 *Consumer surplus is maximized at a strictly positive reserve prices.*

Proof:

Consider the effect on consumer surplus of a small increase in r starting from $r = 0$. We show that consumer surplus is increased via a two step argument. The simple first step is to note that consumer rationality implies that consumer surplus with optimal consumer behavior is greater than the surplus that consumers would receive if they behaved as if $r = 0$.¹⁴ The second step is to show that consumer surplus under this “ $r = 0$ ” behavior is greater when the search engine uses a small positive reserve price dr than when the search engine uses $r = 0$.

If consumers use the $r = 0$ behavior, then consumer surplus is only affected by the institution of a reserve price if the reserve price eliminates links from the list and consumers would have clicked on these links if they were displayed. The gross consumer surplus from each such click is bounded above by dr . The average search costs incurred on each such click are bounded below by $E(s|s \leq \bar{q}_M)$. The cost is independent of dr whereas the benefit is proportional to dr , so the costs dominate for small dr .

QED

4.2.2 Social welfare and revenue need not increase with reserve prices

In this section we show that both social welfare and search engine revenues need not be maximized at a positive reserve price. We do so by presenting examples in which this occurs. The examples are somewhat special, but serve to illustrate mechanisms by which reserve prices can have adverse effects on advertisers and the search engine.

We start with social welfare. The intuition for why this can be reduced by reserve prices is that reserve prices affect social welfare in two ways. First, they directly prevent consumers from clicking on websites with quality less than r . Second, they influence social welfare via their other effects on consumer behavior. Consumers do too little searching from a social perspective because they do not take firm profits into account. If changes to the reserve price policies decrease the number of clicks that occur in equilibrium, then social welfare can decrease.

¹⁴Formally, we suppose that consumers behave exactly as they would if the list had M links and $r = 0$ when deciding whether to click on any link that is displayed and do not click on links that are not displayed.

Proposition 10 *Social welfare can be strictly greater with a zero reserve price than with any positive reserve price.*

Proof:

Consider a model with $M = N = 2$ and the quality distribution F is uniform on $[0, 1]$. Suppose that a fraction γ_1 of consumers have search costs uniformly distributed on $[\frac{2}{3} - \epsilon, \frac{2}{3}]$, a fraction γ_2 have search costs uniformly distributed on $[0, 1]$ and a fraction γ_3 have search costs uniformly distributed on $[0, \epsilon]$.

In the first subpopulation (with $s \approx \frac{2}{3}$), small reserve prices reduce welfare. These consumers click on the first website but not the second when there is no reserve price. Hence, the gain in welfare derived from the search engine not displaying a site they would have clicked on is just $O(r^2)$. Small reserve prices also have an effect that works through changes consumer beliefs: given any small positive r , consumers will not click at all if only one link is displayed. The expected gross surplus from clicking on a single link is $2\frac{1+r}{2}$, whereas the search cost incurred is less than $\frac{2}{3}$, so losing these clicks is socially inefficient. The probability that this will occur is $2r(1 - r)$ so the loss in social welfare is $O(r)$. The appendix contains a formal derivation of this and shows that the per consumer loss in welfare from using any reserve price in $[0, \frac{1}{3} - 2\epsilon]$ is at least $\frac{2}{3}r$.

An example using just the first subpopulation does not suffice to prove the theorem for two reasons: (1) the search cost distribution in this example does not have full support; and (2) although small reserve prices reduce welfare in the first subpopulation it turns out that a larger reserve price ($r > \frac{1}{3} - 2\epsilon$) will increase welfare. (The argument above no longer applies when r is sufficiently large so that consumers will click on the link when a single link is displayed.)

The first problem is easily overcome by adding a very small fraction γ_2 of consumers with search costs uniformly distributed on $[0, 1]$. Welfare gains in this group are first-order in r when r is small and bounded when r is large, so adding a sufficiently small fraction of such consumers won't affect the calculations.

The second problem is also easily overcome by adding a subpopulation of consumers with search costs in $[0, \epsilon]$. Welfare is improved in this subpopulation when a small positive reserve price is implemented, but the effect is so weak that we can add a large mass of these consumers without overturning the small r result from the first subpopulation.¹⁵ There is a substantial welfare loss in this subpopulation if the search engine uses a large reserve

¹⁵The per consumer welfare benefit from a small reserve price is bounded above by total search costs incurred in clicking on ads with quality less than $\text{Min}\{r, \epsilon\}$, which is less than $2\text{Min}\{r, \epsilon\}\epsilon/2$.

price. Hence, adding an appropriate mass of these consumers makes the net effect of using a reserve price of $\frac{1}{3} - 2\epsilon$ or greater also negative.

QED

Our second result is a demonstration that using a small positive reserve prices can also reduce search engine revenue. The example we use to demonstrate this highlights another difference between our model and standard auction models. The crucial property of these models that helps drive our examples is that increasing r increases consumer expectations of the quality of all links, including the bottom one. This makes the M^{th} position more attractive, which can reduce bids for the $M - 1^{\text{st}}$ position. Because bids depend recursively on lower bids, this can reduce reduce bids on higher positions as well.

Proposition 11 *There exist distributions F and G for which search engine revenue is decreasing in the reserve price in a neighborhood of $r = 0$.*

To see why this can be true, consider the following stark example. Suppose $M = N = 2$, and all consumers have search costs of exactly \bar{q}_2 . Assume that with no reserve price consumers click only on the top link. Hence, firms will bid up to their true value to be in the top position and search engine revenue is $E(q^{2:N})$. Given any positive reserve price r , consumers will click on both links. The increased attractiveness of the second position leads to a jump down in bids for the first position. This, of course, leads to a jump down in revenue. To see this formally, bids for the first position (when two firms have $q > r$) will satisfy the indifference condition:

$$(q - b^*(q)) = (1 - q)(q - r).$$

This gives $b^*(q) = r + q(q - r)$. When $r \approx 0$ expected revenues are approximately $E((q^{2:N})^2)$. This is a discrete jump down from $E(q^{2:N})$.

Again, the example is not a formal proof of the Proposition for two reasons: (1) the search cost distribution does not have full support; and (2) we've assumed the search cost distribution has a mass point at \bar{q}_2 . One could easily modify the example to make it fit within our model. Problem (1) could be overcome by adding a small mass γ_2 of consumers with search costs uniformly distributed on $[0, 1]$. And problem (2) could be overcome by spreading out the first population to have search costs uniformly distributed on $[\bar{q}_2, \bar{q}_2 + \epsilon]$. We did not feel, however, that the calculations would be sufficiently enlightening to make it worth doing this for a uniform F or some other such example.

This example is obviously quite special. We include it to illustrate the mechanism that makes it possible, and do not mean to suggest that declining reserve prices are likely to be seen in real-world auctions.

5 Click-weighted Auctions

In Google’s (and more recently in others’) ad auctions the winning bidders are not the firms with the highest per-click bids: advertisers are ranked on the basis of the product of their bid and a factor that is something like an estimated clickthrough rate. The rough motivation for this is straightforward: weighting bids by the click-through rate is akin to ranking them on their contribution to search-engine revenues (as opposed to per-click revenues which is a less natural objective). In this section we develop an extension of our model with observably heterogeneous firms and use it to examine the implications of click-through weighting.

Formally, we consider a model in which each firm has a two dimensional type (δ, q) . A firm of type (δ, q) is able to meet the needs of a fraction δq of consumers. Whether it can meet the need is partially observable. A fraction δ of consumers know from reading the advertisement that the firm can meet their need with probability q (but still don’t know the value of q) and a fraction $1 - \delta$ know that the firm cannot meet their need. Whether a firm can potentially help a consumer is independent across firms. We further assume that the δ parameters are known to the search site and to consumers.

We assume that there are no costs incurred in reading the ads and learning whether a firm is a potential match. Consumers do, however, still need to pay s if they want to investigate a site further and learn whether it does meet their need. Again, this happens with probability q if the firm is a potential match.

5.1 A standard argument for click-weighting auctions

A model of the click-weighted auction is that the firms submit per-click bids b_1, \dots, b_N .¹⁶ The winning bidders are the M bidders for which $\delta_i b_i$ is largest. They are ranked in order of $\delta_i b_i$. If firm i is in the k^{th} position, its per-click payment is the lowest bid that would have placed it in this position, $\delta^{k+1} b^{k+1} / \delta_i$.

Proposition 12 *In equilibrium, the winners of the click-weighted auction are the M firms for which $\delta_i q_i$ is largest. In the limit as $s \rightarrow 0$, social surplus converges to the first-best.*

¹⁶ Again, we can think of this informally as an oral ascending bid auction, but our formalization will be as a multistage game as in our base model.

Proof: Each firm gets zero payoff if it is not on the list. Hence, as long as more than M firms remain, each firm i will want to increase its bid until it reaches q_i . This ensures that the firms for which $\delta_i q_i$ is largest are the winners.

When s is small consumers will search all listed firms that are potential matches until finding a match. The probability of finding a match is $1 - \prod_{k=1}^M (1 - \delta^k q^k)$. This is maximized when the listed firms are those for which $\delta_i q_i$ is largest.

QED

5.2 Inefficiencies of click-weighting

The above proposition is only a partial efficiency theorem for a two reasons, however.

5.2.1 Inefficiency in the set of listed firms

First, when s is not extremely close to zero, utility is not necessarily maximized by choosing the firms for which $\delta_i q_i$ is largest. The reason is that consumers' search costs are reduced if we include firms for which q_i is larger even if $\delta_i q_i$ is lower. For example, if M is a large, then a list of the sites with the largest q 's would be almost sure to contain several sites that were potential matches for each consumer, even if the δ 's for these sites are small. By searching through the sites that are potential matches a consumer would meet his or her need with high probability and incur minimal search costs.

One practical implication of this observation is that click-weighted auctions may allow firms like eBay and Nextag to win more sponsored-link slots than would be socially optimal. The breadth of these sites may allow them to meet more consumers' needs than would a more specialized site, but the extra revenues may be more than fully offset by additional consumer search costs.

5.2.2 Inefficiency in the ordering of listed firms

Second, the click-weighted auction may provide consumers with less than ideal information about the relative q 's of the different websites.

To illustrate this we consider what happens in our model when search costs are small. We do this not because we think the small search costs are important to the argument, but because the equations describing the model are simpler when demand is affected by rank only because consumers first try the top websites and not also because some consumers stop searching before they come to the bottom of the list (which is what requires us to

consider Bayesian updating). If the δ 's are bounded away from zero, this will be the case when $s < \bar{s}$ for some positive constant \bar{s} (which depends on M , N , and the δ 's.)¹⁷

Proposition 13 *Suppose that $M = 2$ and $N > 2$. Suppose all consumers have $s < \bar{s}$. Then the click-weighted auction has an equilibrium in which both firms drop out immediately as soon as just two firms remain.*

To see that the model has an equilibrium in which there is no competition for position, let δ_1 and δ_2 be the weights of the two remaining firms. Assume $\delta_1 < \delta_2$. When players follow these strategies, and the third-to-last firm drops out at b^3 , consumers posteriors would be that $q_1 \sim U[b^3\delta^3/\delta_1, 1]$ and $q_2 \sim U[b^3\delta^3/\delta_2, 1]$. Hence, all consumers would ignore the ordering of the firms and first examine website 1 regardless of its position on the list (provided that it can potentially meet their needs). Given this, there is no incentive for either firm to deviate and bid for a higher position.

The lack of sorting by q means this auction loses the welfare gain from sorting discussed in Section 3. Such immediate dropout equilibria existed in the unweighted auction model, but are more robust here. In the EOS model all “envy-free” equilibria were at least as good for the auctioneer as the VCG-equivalent equilibrium with complete sorting. The envy-free refinement does not apply here.

Although we think these incomplete-sorting equilibria are natural, it should be noted that greater information revelation is also possible. In fact, the model also has an equilibrium with full sorting when $s < \bar{s}$ for all consumers in one special case.

Proposition 14 *Suppose that $N = M = 2$ and $s < \bar{s}$ for all consumers. Then, the click weighted auction has an equilibrium in which the two firms bid according to $b_i^*(q) = \delta_j q_i^2$. In this equilibrium the firm with the highest q is always in the first position on the list.*

Proof: Note that the strategies are monotone and satisfy $\delta_1 b_1^*(q) = \delta_2 b_2^*(q)$. Hence, if firms follow these strategies the winner in a click-weighted auction is the firm with the highest q . Because all consumers search both firms, firm i 's demand is δ_i if it is first on the list and its expected demand from the second position (condition on the other firm being about to drop out) is $\delta_i(1 - \delta_j q)$. Firm i 's indifference condition becomes

$$\delta_i(q - b_i^*(q)) = \delta_i(1 - \delta_j q)(q - 0).$$

This condition is satisfied for the given bidding function.

¹⁷It suffices to set $\bar{s} = E(q^M | \delta^M = 1, \delta^1 = \dots = \delta^{M-1} = \underline{\delta}, z^1 = z^2 = \dots = z^{M-1} = 0)$.

QED

This example uses several special assumptions: the $s < \bar{s}$ assumption eliminates the quality terms from the equation; the third-highest bid is assumed to be zero; and there are only two firms on the list. We believe that the example is nonrobust and does not generalize far beyond this.

5.3 A new auction design: two-stage auctions and efficient sorting

To eliminate the welfare loss due to imperfect sorting one could use a two step procedure. First, have the firms bid as in the standard click-weighted auction until only M bidders remain. Then, continue the auction allowing bidders to raise bids further, but use a different weighting scheme so that the equilibrium will have the firm with the highest q winning.

In theory, this is not hard to do. For example, suppose $M = 2$ and $N > 2$ and $s < \bar{s}$. Then the indifference conditions for an equilibrium in which the high q firm always wins are:

$$\begin{aligned} (1 - \delta_2 q)(q - b^3 \delta^3 / \delta_1) &= q - b_1^*(q) \\ (1 - \delta_1 q)(q - b^3 \delta^3 / \delta_2) &= q - b_2^*(q) \end{aligned}$$

Hence, the equilibrium bids must be

$$\begin{aligned} b_1^*(q) &= b^3 \delta^3 / \delta_1 + \delta_2 q (q - b^3 \delta^3 / \delta_1) \\ b_2^*(q) &= b^3 \delta^3 / \delta_2 + \delta_1 q (q - b^3 \delta^3 / \delta_2) \end{aligned}$$

These will give an equilibrium with the highest firm winning if the rules of the auction are that bidder 1 wins if $b_1^{*-1}(b_1) > b_2^{*-1}(b_2)$ where the b_i^{*-1} are the inverses of the functions given in the last pair of equations.

In this setup, if $\delta_1 < \delta_2$, then the bids entering the second stage satisfy $b_2 < b_1$. Looking at the bidding functions we see that firm 2 continues to be favored at low bid levels, in the sense that if firm 2 raises its bid to b_1 and firm 1 does not raise its bid then firm 2 wins. However, it is possible that at high bids the bid preference is going in the opposite direction: at high bid levels firm 2 may need to bid a higher per-click amount than firm 1 to win.

5.4 Product variety

In our click-weighted model, each site was assumed to have an independent chance of meeting each consumer's needs. In practice these probabilities are unlikely to be independent.

For example, among the sponsored links provided on a recent search for “shorts” were An-nTaylorLoft, ShopAdidas, and RalphLauren. A consumer who clicks on the Adidas site will be more likely to be interested in other sites selling athletic shorts than in other sites selling fashion shorts.

To consider this issue in the simplest extension of our model suppose that there are three sites: site 1A, site 1B, and site 2. Suppose that a fraction $\delta_1 > 1/2$ of consumers are type 1 consumers and can potentially have their needs met by both site 1A and site 1B. The remaining $\delta_2 = 1 - \delta_1$ consumers are type 2 consumers and can potentially have their needs met only by site 2. Suppose that the sponsored link list contains two firms ($M = 2$). Assume that the qualities are independent draws from a uniform distribution on $[0, 1]$.

Intuitively, in such a model conducting a weighted auction instead of an unweighted auction has two effects. First, regardless of whether the weights favor sites 1A and 1B or site 2, weighting reduces the average quality of the listed sites. Second, increasing (decreasing) the weight of site 2 makes it more (less) likely that site 2 will appear on the list. When δ_1 is very large, it is advantageous to have both sites 1A and 1B on the list for same reason as in the standard click-weighting model. When δ_1 is only slightly larger than δ_2 , however, there is a benefit to having a diverse list: when site 1A is first on the list and sites 1B and site 2 are of equal quality, the incremental benefit from including site 1B is smaller than the incremental benefit including site 2 because some of site 1B’s potential consumers will have had their needs met by site 1A.

This model gets complicated, so we again simplify the analysis by considering the special case in which consumers have $s \approx 0$, so that clicks decline at lower positions only because needs are being met and not also because of quality-inferences.

In this environment, consider a weighted $k + 1^{st}$ price ascending bid auction in which winning bidders are chosen by comparing b_{1A} , b_{1B} , and wb_2 .¹⁸ We focus on the case of $w \geq 1$ to discuss when favoring firm 2 is better than equal weighting.

Again, each firm i will bid up to q_i to be included on the two-firm list. Once the bidding is down to two firms, however, equilibrium bidding will produce a slightly different outcome from the unweighted model of section 3. If firms 1A and 1B are on the list, there will again be an equilibrium with full sorting. When firms 1x and 2 are on the list, however, there cannot be an equilibrium with full sorting. Because demand is independent of the expected quality of each site (due to the simplifying assumption that $s \approx 0$ for all consumers and

¹⁸As in section 5.1, the b_i are per-click bids and the per-click payment of firm k is the $k + 1^{st}$ highest bid adjusted for the weight difference (if a difference exists).

the fact that customers served by the two sites are distinct), both firms will drop out immediately.

Given these bidding strategies, suppose that firm 1A is first on the list and the weight w is pivotal in determining which other firm appears, i.e. $q_{1B} = wq_2$. Having firm 1B also on the list provides incremental utility only to type 1 buyers whose needs were not met by firm 1A. Hence, the expected incremental value of including firm 1B (conditional on q_{1A}) is $\delta_1(1 - q_{1A})E(q_{1B}|q_{1B} < q_{1A}, q_{1B} = wq_2) = \delta_1(1 - q_{1A})q_{1A}/2$.¹⁹ Including firm 2 can provide incremental utility to any type 2 buyer: the incremental benefit is $(1 - \delta_1)E(q_2|q_{1B} < q_{1A}, q_{1B} = wq_2) = (1 - \delta_1)q_{1A}/2w$. Using $w > 1$ will provide greater consumer surplus than $w = 1$ if the second term is greater than the first (in expectation) when $w = 1$. The distribution of q_{1A} conditional on q_{1A} being the largest of the three and the other two satisfying $q_{1B} = wq_2$ is just the distribution of the larger of *two* uniform $[0, 1]$ random variables. This implies that the conditional expectation of q_{1A} is $2/3$ and the conditional expectation of q_{1A}^2 is $1/2$. Hence, there is a gain in consumer surplus from choosing $w > 1$ if $\delta_1(1/3 - 1/4) < (1 - \delta_1)1/3$. We have

Proposition 15 *The consumer-surplus maximizing weighted auction is one that favors diversity of the listings if $\delta_1 < 4/5$.*

Proof: A formal proof is given in the Appendix.

The proof in the appendix includes an explicit formula for consumer surplus that could be maximized over w to find the optimal weight for particular values of δ_1 .

Note that the sense in which diversity is favored in this proposition is quite strong. The diversity-providing link is favored in an absolute sense, not just relative to the fraction of consumers for which it is of interest.

To implement diversity-favoring weights, a search engine would need to infer which sponsored links contributed to the diversity of a set of offerings. One way to do this might be to estimate contributions to diversity by looking at whether the likelihood that a particular consumer clicks on a particular site are positively or negatively correlated with whether that consumer clicked on each other site.

What is meant by “standard” click-weighting is not obvious in models like this. One description of the click-weighted auction one sees in the literature is the weight used is the estimated CTR conditional on the firm being first on the list. In the example above, the CTR’s for firms 1A, 1B, and 2 conditional on being first on the list are δ_1 , δ_1 , and

¹⁹Conditioning on $q_{1B} = wq_2$ is irrelevant because conditional on $wq_2 < q_{1A}$, wq_2 is uniform on $[0, q_{1A}]$.

δ_2 , respectively so these standard weights would favor firms 1A and 1B for any $\delta_1 > 1/2$. CTR's could also be estimated using an average of observed CTR's from when a firm is in the first and second positions. This would still favor firm 2 for a smaller range of δ_1 than is optimal, however, because the optimal weights are entirely based on CTR's when firms are in the second position.

5.5 What does click-weighting mean?

The question of what is meant by the “standard” click-weight is of broader importance. In the model of section 5.1, the click-weights were assumed to be the (known) parameters δ . In practice, click weights will be estimated from data on click-throughs as a function of rankings. When the relationship between clicks and rankings is not a known function independent of other website attributes it is not clear what these will mean.

One interesting example is our base model. In this model, suppose that click-through rates are estimated via some regression estimated on data obtained when different subsets of firms randomly choose to compete on different days. Suppose that each website has the same q across days. In this situation, the the clicks that a given site gets when it is in the k^{th} position is a decreasing function of its quality. Conditional on k , the quality of sites $1, 2, \dots, k - 1$ is higher when $q^{k:N}$ is higher. Hence, the likelihood that consumers will get down to the k^{th} position without satisfying their need is lower.

Using click-weights like this will tend to disadvantage higher-quality sites reducing both the average quality of the set of sites presented and eliminating the sorting property of our base model.

6 More auction design

6.1 Obfuscation

In this section we consider advertisers' decisions on how much information to convey in ad text. Consumers benefit from transparent ads that make it apparent whether a link will meet their needs, because such ads let them to avoid unproductive clicks. In our base model firms will also be happy to make ads completely transparent – they only receive a benefit if they can meet a consumer's need. In practice, firms may also receive a lesser benefit from a click even if they cannot meet the consumer's need. We consider such an extension here. We note that the simplest pay per click auction works fairly well, but that both click-weighted

pay-per-click and pay-per-action auctions create incentives for obfuscation.²⁰

We augment our base model in two ways. First, we assume that each firm i receives some benefit z from each click it receives independent of whether it meets the consumer's need. There are several motivations for including such a benefit: firms could earn immediate profits from sales of unrelated products; they could get earn profits on future sales; or they could earn advertising revenues. Second, we assume that each firm choose an obfuscation level $\lambda_i \in \Lambda \subset [0, 1]$. If firm i chooses obfuscation level λ_i then a fraction $1 - \lambda_i$ of the consumers whose needs will not be met by the website will realize this just by reading the text of the firm's ad (without incurring any search costs). We define $\delta_i \equiv q_i + \lambda_i(1 - q_i)$ to be fraction of consumers who cannot tell whether site j will meet their need. Note that our base model can be thought of as a special case of this model with $z = 0$ and no option other than full obfuscation, $\Lambda = \{1\}$.

It is also necessary to specify how consumer inferences change with the possibility of obfuscation. One possibility would be to assume that consumers understand the equilibrium choice of $\lambda_i(q_i)$, observe λ_i for each website, and use this information to update their beliefs about q_i . We feel that this is unrealistic, however, and it leads to an unreasonable and uninteresting degree of equilibrium multiplicity driven by the beliefs assigned when out-of-equilibrium obfuscation levels are observed. We assume instead that consumers cannot detect the obfuscation level chosen by any individual firm. We restrict our analysis to equilibria in which firms are sorted on quality and consumers search in a top-down manner.

Let γ_k be the fraction of consumers who will click on link k if the first $k - 1$ links do not meet their needs and they are in the group that cannot tell whether the k^{th} link meets their needs. This will be a function of consumer beliefs about the quality of the k^{th} website and the equilibrium obfuscation strategies.²¹ One thing that simplifies our analysis is that γ_k does not depend on the actual obfuscation level of the firm in position k .

6.1.1 Pay per click

Consider first the simplest unweighted pay-per-click auction. Conditional on having dropped out of the auction at a bid that places firm i in the k^{th} position ($k \leq M$), firm i 's payoff is

$$\Pi(k, \lambda, b^{k+1}; q_i) = X\gamma_k(q_i + \delta_i z - \delta_i b^{k+1}),$$

²⁰See Ellison and Ellison (2004) for a discussion of obfuscation including a number of examples involving e-retailers.

²¹Note that beliefs about the quality of the k^{th} firm will no longer be independent of the realized qualities because consumers will get some information about the qualities of lower-ranked firms by observing whether these firms can also potentially meet their needs.

where X is the number of consumers who reach position k without having their needs met. In equilibrium, a small change in λ_i that does not affect firm i 's position on the list cannot increase its profits. Note that

$$\frac{\partial \Pi}{\partial \lambda_i} = \frac{\partial \delta_i}{\partial \lambda_i} (z - b^{k+1}) = (1 - q_i)(z - b^{k+1}).$$

In equilibrium b^{M+1} will be at least $z + q^{M+1:N}$, so this is negative and no obfuscation occurs in equilibrium. The intuition is simple. Firms must pay for any clicks they receive. The willingness to pay of losing bidders puts a floor on how cheap clicks can be. And this gives firms an incentive to avoid clicks that provide only incidental benefits.

If there was heterogeneity in the benefits z_i that firms receive from clicks that do not meet consumers' needs, then it is possible that firms with large z_i could engage in obfuscation. But note that it would still be necessary for z_i to be larger than the bid of the firm in the next highest position, which suggests that obfuscation is unlikely to occur except perhaps at very low positions on the list.

Consider now a click-weighted pay-per-click auction in which the search engine uses click-through weights proportional to the δ_i .²² Conditional on being in the k^{th} position ($k \leq M$), firm i 's payoff is

$$\Pi(k, \lambda, b^{k+1}; q_i) = X \gamma_k (q_i + \delta_i z - \delta_i \frac{\delta^{k+1} b^{k+1}}{\delta_i}).$$

This expression is monotone increasing in δ_i . Hence, in equilibrium we get full obfuscation: all firms chose $\lambda_i = 1$. The intuition is that in the click-weighted auction each firm's total payment to the search engine is equal to the total payments that the firm below it would make. Importantly, this quantity is independent of the number of clicks firm i receives. Hence, firm i will design its ad to maximize the total revenue it can generate, and has no incentive to avoid unproductive clicks. In thinking about practical applications, it is noteworthy that we get full obfuscation even if the benefit z is quite small.

In summary, the simplest pay-per-click auction first used by Overture/Yahoo! minimized obfuscation, but the now-standard click-weighted auction has an adverse side effect: it creates incentives for obfuscation, because firms are not penalized for unnecessary clicks. Search engines may attempt to combat obfuscation in various ways. Real-world search engines have rules forbidding misleading ad text. They can enforce these via manual reviews and by using the technological capabilities they have because they are search engine

²²Note that we are implicitly assuming here that in equilibrium the search-engine has learned firm i 's click-through rate and uses it in determining the rankings and the per-click price firm i must pay.

operators: advertisers' landing pages can be examined for relevance to the query and advertisements with irrelevant landing pages can be declined. Other approaches are also possible. The pricing formula could be adjusted so that firm i 's total payment is no longer completely independent of its click-through rate, or continuous relevance measures (based on textual analysis or the number of consumers who immediately return to the search page) could be added to the pricing formula.

6.1.2 Pay per action

Search engines have been developing the capability to track how sales are made by their advertisers. This leads to the question of whether it is advantageous to charge producers by sale made, instead of charging them per click.²³ Here, we discuss what impact this might have on obfuscation.

To model pay-per action auctions we suppose that firms submit bids b_i which represent payments to be made to the search engine every time a consumer clicks on their link *and meets their need*. We assume the search engine tracks these conversions and uses the conversion rates as an additional weighting factor just as click-through-rates are used in the click-weighted auctions: if the number of clicks that each website j will receive in any position k is proportional to δ_j and the fraction of clicking consumers who will have their need met is y_j , then the search engine ranks the firms on the basis of $\delta_j y_j b_j$ and firm i will make a payment of $\frac{\delta^{k+1} y^{k+1} b^{k+1}}{\delta_i y_i}$ every time it meets a need if its ad is displayed in position k .

Conditional on being in the k^{th} position ($k \leq M$), firm i 's payoff is

$$\Pi(k, \lambda, b^{k+1}; q_i) = X\gamma_k(q_i + \delta_i z - \delta_i y_i \frac{\delta^{k+1} y^{k+1} b^{k+1}}{\delta_i y_i}).$$

This expression is virtually identical to the expression for the standard click-weighted auction. The result on obfuscation carries over: all firms chose full obfuscation. The intuition is identical to that for the click-weighted auctions: each firm's total payment to the search engine is equal to the total payments that the firm below it would make. This is independent of the number of clicks firm i receives, so firm i has no incentive to avoid unproductive clicks.

²³There are many reasons why pay-per-action pricing might be desirable for advertisers; for example, it decreases risk for them if their ads are shown on a large network, where click fraud may be a problem.

6.2 Search-diverting Sites

In the U.S. market three main firms provide search services and sell keyword advertising: Google, Yahoo!, MSN.²⁴ Google is substantially larger than the others and earns higher revenue per search. When keyword searches are performed on each of these search engines, particularly on MSN and Yahoo!, sponsored link slots are not infrequently occupied by sites like Nextag, Shopzilla, Bizrate, Smarter, Shopping, Cataloglink, Coupon Mountain. These sites provide some search and shopping comparison services, but also earn revenues simply by prominently posting lists of sponsored links provided by Google. We refer to them collectively as *search-diverting* sites.

Under particular assumptions, this could be an efficient way for MSN and Yahoo! to monetize their searches in light of Google’s technology and scale advantages. Suppose that Google pays 100% of the revenues generated by these sponsored links to the search-diverting sites. Suppose that all consumers who click on a search-diverting site continue their search on the list of sponsored links provided there exactly as they would on Google itself. Suppose also that consumers recognize that all search-diverting sites provide exactly the same list, and hence won’t click on a second site of this variety. Then, each of the search-diverting sites would be willing to bid up to Google’s total expected revenue for opportunity to be in the first position. Even though consumers usually just click on the first search diverting ad and never return to the search engine they started from, MSN and Yahoo! would receive the full Google revenues from this single click.

Under other assumptions, however, the presence of these search-diverting sites could reduce search-engine profits and welfare. Here is a simple example. Suppose that all sponsored link lists contain M places and there are $N > M$ advertisers bidding for a particular keyword on Google. Suppose that $N + 1$ firms are bidding for the same keyword on Yahoo!/MSN: the same N firms bidding on Google plus an $N + 1^{st}$ “search-diverting” firm that displays Google’s sponsored link list and receives a share $\phi < 1$ of all revenues that Google receives from clicks on this list. Suppose that consumers are unaware of this asymmetry and therefore click on any ads in Yahoo!/MSN in the order in which they are listed.²⁵ Once they click on the search-diverting ad, however, assume that they realize that it is providing a sponsored link list generated by N firms bidding against each other.

As long as the revenue share ϕ is not too small, the equilibrium of this model will have

²⁴The fourth largest search engine, Ask.com, displays ads provided by Google.

²⁵Suppose also that their beliefs about quality before clicking on the links are as in our base model with $N + 1$ firms.

the search diverting firm in the first position. It gets a fraction ϕ of Google’s revenues from all clicks that take place, whereas the benefit to the highest q regular advertiser from being listed first ahead of the search-diverting firm is just the increment in clicks that derives from being thought to be first of $N + 1$ firms (which is the belief if it comes top on the Yahoo!/MSN list) relative to being first of N firms (which is the belief that consumers have after seeing it on the Google list that the search-diverting firm presents).

Indeed, the model has an equilibrium in which all regular advertisers drop out at the reservation price because consumers will see their ad on the Google list and hence they get no incremental sales from clicks on their Yahoo!/MSN list.

Consumer surplus is reduced slightly by the presence of the search-diverting firm because consumers have to make one extra click to begin the search process. Search engine profit is reduced because because the search-engine only gets the search-diverting firm’s bid, which is less than its revenue, which is a fraction ϕ of the revenue that would have been obtained if the N regular advertisers were the only bidders.

6.3 Ad relevance and consumer inference about search engine quality

Another dimension in which search engines differ is in their tolerance for irrelevant ads. Google often presents no ads for certain types of queries (such as the names of ordinary people), whereas MSN presents more non-specific ads such as ads for ringtones, eBay, or Amazon. One way to think of such policies would be to regard them as similar to reserve prices. In this section, we note that an additional way would be to add consumer uncertainty as to the relevance of links to the model.

To model this, suppose the distribution of q_i is either $U[0, 1]$ with probability ρ (relevance probability) or degenerate at zero with probability $1 - \rho$. So either there is a distribution of probabilities of matches, or it is certain that no firm offers a match.²⁶ Consumers will then update both with respect to ad quality and with respect to whether the set of ads presented is relevant or irrelevant as they move down the list.

²⁶This formulation does not provide a motivation for advertisers to provide non-specific advertisements. One interpretation is that the advertisers do not in fact choose to do so, but they select a “broad” match to a class of search terms, and the search engine technology sometimes makes mistakes in determining what is relevant. One could also add a benefit independent of meeting consumers’ needs as in our model of obfuscation.

Given this model, we have

$$\begin{aligned} \Pr\left(q_i \sim U[0, 1] \mid z^1 = \dots = z^{k-1} = 0\right) &= \frac{\Pr\left(z^1 = \dots = z^{k-1} = 0 \mid q_i \sim U[0, 1]\right) \rho}{\Pr\left(z^1 = \dots = z^{k-1} = 0 \mid q_i \sim U[0, 1]\right) \rho + (1 - \rho)} \\ &= \frac{\rho \prod_{j=1}^{k-1} \frac{2j-1}{N+j}}{\rho \prod_{j=1}^{k-1} \frac{2j-1}{N+j} + (1 - \rho)}. \end{aligned}$$

Consumers' expectations about the quality of the k^{th} sponsored link conditional on reaching it are:

$$\begin{aligned} E(q^{k:N} \mid z^1 = \dots = z^{k-1} = 0, q_i \sim U[0, 1]) \Pr\left(q_i \sim U[0, 1] \mid z^1 = \dots = z^{k-1} = 0\right) \\ = \frac{N + 1 - k}{N + k} \frac{\rho \prod_{j=1}^{k-1} \frac{2j-1}{N+j}}{\rho \prod_{j=1}^{k-1} \frac{2j-1}{N+j} + (1 - \rho)}. \end{aligned}$$

Some observations from these formula are:

1. Holding N and k fixed, $\Pr\left(q_i \sim U[0, 1] \mid z^1 = \dots = z^{k-1} = 0\right)$ is increasing and convex in ρ . When k is large, even a small decrease in ρ away from 1 can lead to a large decrease in beliefs about the quality of the k^{th} site. If ρ is not close to one than most consumers will only click on a small number of ads.
2. Holding N and ρ fixed, $\Pr\left(q_i \sim U[0, 1] \mid z^1 = \dots = z^{k-1} = 0\right)$ is decreasing and convex in k . When N is large, a consumer's posterior belief in relevance drops dramatically after failing to have his or her needs met at the first site.

7 Conclusions

In this paper we have integrated a model of consumer search into a model of auctions for sponsored-link advertising slots. General observations from previous papers about the form of the auction equilibrium are not much affected by this extension: advertisers bid up to their true value to be included in the sponsored-link list and then shade their bids when competing for a higher rank.

The differences in the auction environment does, however, have a number of different implications for auction design. One of these is that reserve prices can increase both search-engine revenues and consumer surplus. The rationality of consumer search creates a strong alignment between consumer surplus and social welfare in our model and a consumer-surplus maximizing search engine will have a strong incentive to screen out ads so that consumers don't lose utility clicking on them. Another set of different implications arise

when we consider click-through weighting. Here, the auction that is efficient with no search costs ceases to be efficient for two reasons: it may select the wrong firms and it may provide consumers with little information to guide their searches. The informational inefficiency can be avoided with an alternate auction mechanism. An additional worry about click-weighted auctions is ofuscation – advertisers are not actually paying per click, so they have no incentive to design ad text to help consumers avoid unnecessary clicks.

A more basic theme of our paper is that sponsored link auctions create surplus by providing consumers with information about the quality of sponsored links which allows consumers to search more efficiently. Sorting firms on the basis of their click-weighted bids is part of this, but in principle one could imagine many other search engine designs that produce different sets of sponsored links and/or present consumers with information in other ways. This should be an interesting area for pure and applied research.

Appendix

Proof of Proposition 5

First, we show by induction on k that the specified strategies are differentiable and strictly monotone increasing in q and satisfy $b^*(k, b^{k+1}, q) \leq q$ on the equilibrium path. For $k = M + 1$ this is immediate from $b^*(M + 1, 0, q) = q$. If it holds for some $K > 2$ then for any b^K faced by a type q bidder on the equilibrium path we have

$$\begin{aligned} b^*(K - 1, b^K, q) &= b^K + (q - b^K) \left(1 - \frac{G(\bar{q}_{K-1})}{G(\bar{q}_{K-2})} (1 - q) \right) \\ &\leq b^K + (q - b^K) = q. \end{aligned}$$

The inequality here follows from two observations: $q - b^K > 0$; and the term in parentheses is between 0 and 1. (The first of these follows from the inductive hypothesis via $q - b^K \geq q^{K:N} - b^K \geq 0$ and the second comes from $1 - q < 1$, $\bar{q}_{K-1} < \bar{q}_{K-2}$, and G strictly monotone.) To see that the bidding function is differentiable and strictly monotone increasing in q , one can compute the derivative and see that it is positive. (The inductive hypothesis is again used here via $q - b^K \geq 0$.)

We now show that the bidding functions are a perfect Bayesian equilibrium. By the single-stage deviation principle, it suffices to show that no single-stage deviation can increase the profit of a player i of type q_i . We do this by another inductive argument. We first show that this is true of deviations in the final stage ($k = 2$). And we then show that the nonexistence of profitable deviations at all later stages (all $k' < k$) implies that there is also no profitable single stage deviation at stage k .

Consider the final stage of the game. Suppose firm i has quality q_i and that $b^3 = b^*(3, b^4, \underline{q})$ so that firm i 's belief is that the other active firm has $q_j \sim F|_{q > \underline{q}}$. Firm i 's expected payoff as a function of its dropout point \hat{q} can be written as $\frac{1}{1-F(\underline{q})} \pi(q_i, \hat{q})$ where

$$\pi(q_i, \hat{q}) = \left(\int_{\underline{q}}^{\hat{q}} G(\bar{q}_1)(q_i - b^*(3, b^3, q)) f(q) dq + \int_{\hat{q}}^1 G(\bar{q}_2)(1 - q)(q_i - b^3) f(q) dq \right).$$

To show that this is maximized at $\hat{q} = q_i$ it suffices to show that $\pi(q_i, q_i) - \pi(q_i, \hat{q}) \geq 0$ for all \hat{q} .

For $\hat{q} \leq q_i$ we have

$$\pi(q_i, q_i) - \pi(q_i, \hat{q}) = \int_{\hat{q}}^{q_i} (G(\bar{q}_1)(q_i - b^*(3, b^3, q)) - G(\bar{q}_2)(1 - q)(q_i - b^3)) f(q) dq.$$

To show that this is nonnegative it suffices to show that

$$G(\bar{q}_1)(q_i - b^*(3, b^3, q)) \geq G(\bar{q}_2)(1 - q)(q_i - b^3)$$

for all $q \in [\hat{q}, q_i]$. Because the bidding functions are differentiable and strictly monotone increasing in q , the argument in the text before the proposition applies and therefore for each q in this interval the local indifference condition holds:

$$G(\bar{q}_1)(q - b^*(3, b^3, q)) = G(\bar{q}_2)(1 - q)(q - b^3)$$

Subtracting the two equations we find that it suffices to show

$$G(\bar{q}_1)(q_i - q) \geq G(\bar{q}_2)(1 - q)(q_i - q).$$

This is indeed satisfied for all $q \in [\hat{q}, q_i]$ because $G(\bar{q}_1) > G(\bar{q}_2)$ and $(1 - q) < 1$. The argument for $\hat{q} > q_i$ is virtually identical. Together, these two cases establish that there is no profitable single-stage deviation in the final stage.

Suppose now that there are no profitable deviations from the given strategies in stages $2, 3, \dots, k - 1$ and consider a stage k history with $b^{k+1} = b^*(k + 1, b^{k+2}, q)$. To show that there is no profitable single stage deviation, we'll consider separately deviations to $\hat{b} > b^*(k, b^{k+1}, q_i)$ and deviations to $\hat{b} < b^*(k, b^{k+1}, q_i)$.

The first case is quite similar to the argument for $k = 2$. Deviating to $\hat{b} > b^*(k, b^{k+1}, q_i)$ makes no difference unless player i is eliminated in stage k when he bids $b^*(k, b^{k+1}, q_i)$ and is not eliminated when he bids \hat{b} . Hence for all relevant realizations of the $k - 1^{\text{st}}$ highest quality, player i will be the first to drop out in stage $k - 1$ if he then follows the equilibrium strategy. Hence, writing the change in payoff is proportional to

$$\begin{aligned} & \int_{q_i}^{\hat{q}} E \left((1 - q^{1:N})(1 - q^{2:N}) \dots (1 - q^{k-2:N})(1 - q) | q^{k-1:N} = q \right) \cdot G(\bar{q}_k) \cdot (q_i - b^{k+1}) f(q) dq \\ & - \int_{q_i}^{\hat{q}} E \left((1 - q^{1:N})(1 - q^{2:N}) \dots (1 - q^{k-2:N}) | q^{k-1:N} = q \right) \cdot G(\bar{q}_{k-1}) \cdot (q_i - b^*(k, b^{k+1}, q)) f(q) dq, \end{aligned}$$

where \hat{q} is the solution to $b^*(k, b^{k+1}, \hat{q}) = \hat{b}$. (A solution to this exists because the bidding functions are differentiable and approach 1 in the limit as $q \rightarrow 1$.) As above, this will be nonnegative if

$$(1 - q)G(\bar{q}_k)(q_i - b^{k+1}) \geq G(\bar{q}_{k-1})(q_i - b^*(k, b^{k+1}, q))$$

for all $q \in [q_i, \hat{q}]$. Subtracting the local indifference condition from the two sides of this equation we again obtain that a sufficient condition is

$$(1 - q)G(\bar{q}_k)(q_i - q) \geq G(\bar{q}_{k-1})(q_i - q) \quad \forall q \in [q_i, \hat{q}].$$

This will hold because $q_i - q < 0$ and $0 < (1 - q)G(\bar{q}_k) < G(\bar{q}_{k-1})$. The argument for deviations to $\hat{b} < b^*(k, b^{k+1}, q_i)$ is just a little more complicated. In this case, the deviation makes no difference unless player i is eliminated in stage k when he bids \hat{b} and is not eliminated at this stage when he bids $b^*(k, b^{k+1}, q_i)$. We show that the change in payoff is not positive by a two-step argument: we show that that the payoff from dropping out at \hat{b} is worse than the payoff from bidding $b^*(k, b^{k+1}, q_i)$ at stage k and then dropping out immediately in stage $k - 1$; and that this in turn is less than the payoff from bidding $b^*(k, b^{k+1}, q_i)$ at stage k and then following the given strategies. The latter comparison is immediate from the inductive hypothesis. Hence, it only remains to show that

$$\begin{aligned} & \int_{\hat{q}}^{q_i} E \left((1 - q^{1:N})(1 - q^{2:N}) \dots (1 - q^{k-2:N}) | q^{k-1:N} = q \right) \cdot G(\bar{q}_{k-1}) \cdot (q_i - b^*(k, b^{k+1}, q)) f(q) dq, \\ & - \int_{q_h}^{q_i} E \left((1 - q^{1:N})(1 - q^{2:N}) \dots (1 - q^{k-2:N})(1 - q) | q^{k-1:N} = q \right) \cdot G(\bar{q}_k) \cdot (q_i - b^{k+1}) f(q) dq \end{aligned}$$

is nonnegative where $\hat{q} < q_i$ is the solution to $b^*(k, b^{k+1}, \hat{q}) = \hat{b}$. This is just like the argument for the $\hat{q} > q_i$ case above. The expression is nonnegative if

$$G(\bar{q}_{k-1})(q_i - b^*(k, b^{k+1}, q)) \geq (1 - q)G(\bar{q}_k)(q_i - b^{k+1})$$

for all $q \in [\hat{q}, q_i]$. Subtracting the local indifference condition from the two sides of this equation we again obtain that a sufficient condition is

$$G(\bar{q}_{k-1})(q_i - q) \geq (1 - q)G(\bar{q}_k)(q_i - q) \quad \forall q \in [\hat{q}, q_i].$$

This will hold because $q_i - q > 0$ and $0 < (1 - q)G(\bar{q}_k) < G(\bar{q}_{k-1})$.

This completes the proof that there is no profitable deviation at stage k and the result follows by induction.

QED

Additional Details on the Proof of Proposition 10

With no reserve price, consumers with search costs in $[\frac{2}{3} - \epsilon, \frac{2}{3}]$ will click only on the first link. Per consumer social welfare is

$$\begin{aligned} W &= 2E(q^{1:N} - \frac{2}{3} - \frac{\epsilon}{2}) \\ &= \frac{2}{3} + \frac{\epsilon}{2}. \end{aligned}$$

Suppose now that the search engine uses a small positive reserve price r . (More precisely assume $r \in (0, \frac{1}{3} - 2\epsilon)$). These consumers now click on the first link only if two links are displayed. Per consumer social welfare becomes

$$\begin{aligned} W &= (2E(q^{1:N} | q^{2:N} \geq r) - s)(1 - r)^2 \\ &= (2(\frac{2}{3} + \frac{1}{3}r) - (\frac{2}{3} - \epsilon/2))(1 - r)^2 \\ &= \frac{2}{3} + \frac{\epsilon}{2} - \frac{2}{3}r - \frac{2}{3}(r^2 - r^3) - \frac{\epsilon}{2}(2r - r^2) \\ &< \frac{2}{3} + \frac{\epsilon}{2} - \frac{2}{3}r. \end{aligned}$$

For somewhat larger r , specifically $r \in [\frac{1}{3} - 2\epsilon, \frac{5}{9} - \frac{4}{3}\epsilon]$, consumers in the high search cost group will click on the top link even if only one link is displayed. In the high search cost population per-consumer welfare is now

$$W = (2E(q^{1:N} | q^{2:N} \geq r) - s)(1 - r)^2 + (2E(q^{1:N} | q^{1:N} > r, q^{2:N} < r)2r)(1 - r).$$

Using this, we one can show that the per consumer welfare gain in the high-search cost subpopulation is at most $\frac{2}{3}r^2(1 - 2r)$. This is negative for $r > \frac{1}{2}$ and is uniformly bounded above by $\frac{2}{81}$. Computing the mass of needs that go unmet because of the reserve price that would have been met without a reserve price we find that the per consumer loss in welfare in the low-search cost population is at least $2r(1 - r)\frac{1-r}{2}2\frac{r}{2} + r^2(4\frac{r}{2} - 2\frac{r^2}{4}) - 2\epsilon^2$. It is easy to choose γ_1 and γ_3 so that this outweighs any gains in the high-search costs population whenever $r \geq \frac{1}{3} - 2\epsilon$.

For even larger r the high search cost consumers will be willing to search both sites when two are listed. But again, one can show that the welfare losses in the low search cost population will outweigh this.

QED

Proof of Proposition 15

To compute expected consumer surplus we compute the probability that each subset of firms is listed and the expected quality of the listed firms conditional on that subset being selected. Write L for the set of firms listed. The main probability fact we need is easy:

$$\text{Prob}\{L = \{1A, 1B\}\} = 1/3w$$

To see this, note that $L = \{1A, 1B\}$ is possible only if $q_2 \in [0, 1/w]$. This happens with probability $1/w$ conditional on q_2 being in this range, $L = \{1A, 1B\}$ occurs with probability $1/3$ (because wq_2 is then uniformly distributed on $[0, 1]$).

The expected qualities are

$$\begin{aligned} E(q_{1x}|L = \{1A, 1B\}, q_{1x} > q_{1y}) &= \frac{3}{4} \\ E(q_{1x}|L = \{1A, 1B\}, q_{1x} < q_{1y}) &= \frac{1}{2} \\ E(q_{1x}|L = \{1x, 2\}) &= \frac{8w - 3}{12w - 4} \\ E(q_2|L = \{1x, 2\}) &= \frac{6w^2 - 1}{12w^2 - 4w} \end{aligned}$$

The first two are again identical to the formulas for the unweighted case because this L only arises when $q_2 \in [0, 1/w]$ and in this event wq_2 is uniformly distributed on $[0, 1]$. The latter two formulas can be derived fairly easily by conditioning separately on values with $q_2 \in [0, 1/w]$ and values with $q_2 \in [1/w, 1]$. For example,

$$\begin{aligned} E(q_{1x}|L = \{1x, 2\}) &= \\ \frac{\text{Pr}\{q_2 \in [1/w, 1]\}\text{Pr}\{L = \{1x, 2\}|q_2 \in [1/w, 1]\}E(q_{1x}|L = \{1x, 2\}, q_2 \in [1/w, 1]) + \text{Pr}\{q_2 \in [0, 1/w]\}\text{Pr}\{L = \{1x, 2\}|q_2 \in [0, 1/w]\}E(q_{1x}|L = \{1x, 2\}, q_2 \in [0, 1/w])}{\text{Pr}\{q_2 \in [1/w, 1]\}\text{Pr}\{L = \{1x, 2\}|q_2 \in [1/w, 1]\} + \text{Pr}\{q_2 \in [0, 1/w]\}\text{Pr}\{L = \{1x, 2\}|q_2 \in [0, 1/w]\}} \\ &= \frac{(1 - 1/w)(1/2)(2/3) + (1/w)(1/3)(5/8)}{(1 - 1/w)(1/2) + (1/w)(1/3)} \end{aligned}$$

Expected consumer surplus when weight w is used is then given by

$$E(CS(w)) = \alpha \left(\left(1 - \frac{1}{3w}\right) \frac{8w - 3}{12w - 4} + \frac{1}{3w} \left(\frac{3}{4} + \frac{1}{4} \cdot \frac{1}{2}\right) \right) + (1 - \alpha) \left(\left(1 - \frac{1}{3w}\right) \frac{6w^2 - 1}{12w^2 - 4w} + \frac{1}{3w} \cdot 0 \right)$$

The difference between this expression and the expected consumer surplus from an unweighted auction can be put in a relatively simple form by grouping terms corresponding

to cases when the list is unaffected by the changes in weights and cases when it is affected. We find

$$\begin{aligned}
E(CS(w)) - E(CS(1)) &= \frac{2}{3} \left(\alpha \frac{8w-3}{12w-4} + (1-\alpha) \frac{6w^2-1}{12w^2-4w} - 5/8 \right) \\
&+ \frac{1}{3w} \left(\alpha \left(\frac{3}{4} + \frac{11}{42} \right) - \alpha \frac{7}{8} \right) \\
&\left(\frac{1}{3} - \frac{1}{3w} \right) \left(\alpha \frac{8w-3}{12w-4} + (1-\alpha) \frac{6w^2-1}{12w^2-4w} - \alpha \frac{7}{8} \right)
\end{aligned}$$

Writing $f_1(w)$, $f_2(w)$ and $g_3(w)h_3(w)$ for the three lines of this expression note that all three terms are equal to zero at $w = 1$. $f_2(w)$ is identically zero. The derivative of the third evaluated at $w = 1$ is just $dg_3/dw|_{w=1}h_3(1)$. After these simplifications it takes just a little algebra to show

$$\frac{d(E(CS(w)) - E(CS(1)))}{dw} = \frac{1}{24}(4 - 5\alpha).$$

This implies that some $w > 1$ provides greater consumer surplus than $w = 1$ provided that $\alpha < 4/5$. To complete the proof, we should also work out the equations for consumer surplus when $w < 1$ and show that these do not also provide an increase in consumer surplus.

QED

References

- Aggarwal, Gagan, John Feldman, and S. Muthukrishnan (2006): “Bidding to the Top: VCG and Equilibria of Position-Based Auction,” arXiv: cs.GT/0607117 v1.
- Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani (2006): “Truthful Auctions for Pricing Search Keywords,” ACM Conference on Electronic Commerce.
- Tilman Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, Vaclav Petricek (2006): “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” Mimeo.
- Borgs, Christian, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi (2006): “Multi-unit Auctions with Budget Constrained Bidders,” mimeo, Microsoft Research.
- Chen, Yongmin and Chuan He (2006): “Paid Placement: Advertising and Search on the Internet,” mimeo, NET Institute.
- Demange, Gabrielle, David Gale and Marilda Sotomayor (1986): “Multi-Item Auctions,” *Journal of Political Economy* 94(4), 863–872.
- Edelman, Benjamin and Michael Ostrovsky (2007): “Strategic Bidder Behavior in Sponsored Search Auctions,” *Decision Support Systems* 43(1), 192–198.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz (2007): “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review* 97(1), 242–259.
- Edelman, Benjamin and Michael Schwarz (2006): “Optimal Auction Design in a Multi-unit Environment: The Case of Sponsored Search Auctions,” mimeo, Harvard University and Yahoo! Research.
- Ellison, Glenn and Sara Fisher Ellison (2004): “Search, Obfuscation, and Price Elasticities on the Internet,” NBER working paper 10570.
- Ellison, Glenn, Drew Fudenberg, and Markus Möbius (2004): “Competing Auctions,” *Journal of the European Economic Association* 2(1), 30–66.
- Immorlica, Nicole, Kamal Jain, Mohammad Mahdian, and Kunal Talwar (2006): “Click Fraud Resistant Methods for Learning Click-Through Rates,” mimeo, Microsoft Research.
- Hu, Yu (2004): “Performance-Based Pricing Models in Online Advertising,” mimeo, MIT Sloan School of Management.
- Lehaie, Sebastien (2006): “An Analysis of Alternative Slot Auction Designs for Sponsored Search,” EC ’06.
- Varian, Hal R. (2007): “Position Auctions,” *International Journal of Industrial Organization* 25(6), 1163–1178.