

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 deriving from sales that are made 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 application 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 abstract away from the fact that the “objects” being auctioned are advertisements. They thereby ignore that the values of the objects are potentially endogenous to the auction design – the value of a link is in part due to consumers’ clicking on the links and making purchases and it is natural to assume that consumer behavior will be affected by the process by which links are selected for display. Chen and He (2006) is a noteworthy exception – they develop

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

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. In this paper, we develop further 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, 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 game 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 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. As noted above, 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.² Our model does not fit in their more general framework, 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 closely related to our paper. They had previously introduced a model that is much 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.³ We very 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 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.

²BCPP also contains an empirical analysis which includes methodological innovations and estimates of how value-per-click changes with position in Yahoo! data.

³As in Diamond (1971) the equilibrium turns out to be that all firms charge the monopoly price.

Our model extends their work in several ways. They 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 results are addressing issues that don't come up in their framework. There is no analog to our derivation of optimal consumer search strategies because their consumers simply click on all sponsored links (if necessary). There is 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 our other auction design issues. (Most of these hinge on how auction design affects the information consumers get about firm qualities and thereby influences consumer search, which again is not something that comes up in their paper.)

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 will assume s_j to have a 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, and this is private information of firm i . 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. Throughout the paper, we use the uniform distribution as a leading example of F .

A monopoly search site conducts button-style 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 $k + 1^{st}$ highest bid for each click it gets.⁴

⁴Note 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 4. We present result first for the unweighted auction because the environment is easier to analyze. It should also be a good approximation to real-world auctions in which there are not substantial differences in click-through rates

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.

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 $(1 - \bar{q})$)

across firms, e.g. where the bidders are all retailers with similar business models.

they get $(\bar{q} - s)$ from their second search. 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, $q^{1:N}$.⁵ 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 orders 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 from the first $k - 1$ advertisers:

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

Proposition 2 (i) *If the firms are sorted by quality in equilibrium, then consumers follow a top-down strategy: they start at the top 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*

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

So a firm in position k will receive

$$(1 - q^{1:N}) \dots (1 - q^{k-1:N}) G(\bar{q}_k)$$

clicks. (ii) *When we make the additional assumption that F is uniform, a consumer with search cost s stops clicking when she reaches position $k^{max}(s)$, where*

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

Proof: (i) 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

⁵As 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.

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

$$E(q^{k:N} | z^1 = \dots = z^{k-1} = 0) = \frac{\int_0^1 x f^{k:N}(x) \text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} dx}{\text{Prob}\{z^1 = \dots = z^{k-1} = 0\}},$$

where $f^{k:N}$ is the PDF of the k^{th} order statistic.

(ii) 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 $k < k^{max}(s)$. QED

For the case of uniform quality distributions, expected consumer surplus is easy to compute given the result of Lemma 1. 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 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 $1 - x$ whereas the unordered payoff is approximately $1 - 2x$. 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 do this under the assumption that search costs are uniformly distributed on $[0, 1]$. 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.⁶

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 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. Given our assumption that 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*

$$D(k, q) = \left(\frac{1+q}{2}\right)^{k-1} \frac{N+1-k}{N+k}$$

$$D(k) = \frac{1 \cdot 3 \cdot \dots \cdot (2k-3)}{(N+1)(N+2)\dots(N+k-1)} \frac{N+1-k}{N+k}$$

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

⁶In a model with endogenous search there will, of course, also naturally 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. We have not fully specified our model to allow for simultaneous dropouts because they do not occur in the equilibria we consider.

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 $1 - E(q_j | q_j > q) = (1 + q)/2$.⁷ 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

3.2 Equilibrium in the bidding game

Consider now an ascending auction in which N firms bid for the $M < N$ positions. The highest M bidders are sorted in order of their bids and the each website pays a per-click fee equal to the bid of the next highest bidder.

Note that conditional on being clicked on a website 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 auction model.

Although one can think of our model as being like the Edelman-Ostrovsky-Schwarz position auction with endogenous click-through rates, the auction model does not fit directly 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.⁸ 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.

When more than M firms remain, firms will get zero if they drop out. Hence, for a firm with quality q it is a weakly dominant strategy to remain in the bidding if more than M firms remain and the price is less than q . It is also a weakly dominant strategy to drop out as soon as the price is greater than q . We assume that all bidders behave in this way.⁹

⁷Note that the j in this expression is a generic index and does not denote the j^{th} highest value.

⁸The BCPP paper has a more general setup, but they still assume that click-through rates do not depend on the types of the other bidders.

⁹This can be viewed as another restriction on the subset of equilibria we are considering.

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 . (In principle, drop out points could condition on the history of drop out points in other ways.)

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 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 higher demand. The last is lower 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. When more than M firms remain, each firm bids up to its true value, so that $b^*(k, b^k; q) = q$ whenever $k > M$ and $b^k < q$. For $k \leq M$, the dropout point of a firm that*

has quality q when k bidders remain and the $k + 1^{\text{st}}$ highest bid is b^{k+1} is given by

$$b^*(k, b^{k+1}; q) = b^{k+1} + (q - b^{k+1}) \left(1 - \frac{G(\bar{q}_k)}{G(\bar{q}_{k-1})} (1 - q) \right)$$

When both qualities and search costs are uniform, we have

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

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 preceding the proposition therefore establish that firms are indifferent to local deviations from the prescribed strategies. Showing that firm i cannot gain by deviating and dropping out at a lower bid when k others remain is easy: if dropping out earlier does not change the firm's position then the deviation makes no difference; and if it does make a difference then it is immediate from the equation from which the indifference condition was derived shows that the firm would have been better off waiting until the next dropout and then dropping out.

It then remains only to show that a firm i cannot gain by dropping out later. Again, such a deviation makes no difference if the firm i does not have the lowest remaining value, nor does it make a difference if firm i is still the first to drop out. If some other firm does drop out in the interim, then we appeal to the one-stage deviation principle: the equation from which we derived the indifference shows that the firm is worse off if it drops out immediately after the other firm; and because the prescribed strategies are monotone firm i will indeed be the next firm to drop out if it follows the equilibrium strategies from this point forward.

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, bids increase very 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 latter condition is satisfied when the distribution of q is uniform. In that distributions is nondecreasing in k , recalling that $G(\bar{q}_k)$ is the fraction of consumers who have search costs less than the average quality of the k^{th} best firm, so that the percentage change in the click-through rate from moving to a higher position is smaller as the firm moves up the list. Note that in reality, the biggest percentage change in click-through rates comes as firms move into the highest position.

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. 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 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.¹⁰ 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.

We assume throughout this section that consumer search costs and advertiser quality are uniformly distributed on $[0, 1]$, that equilibria with strictly monotone bidding strategies are

¹⁰See Ellison, Fudenberg and Möbius (2004) for more on a competing auction model in which this effect would be important.

always played, and that firms bid their true value when they will not be on the sponsored-link list, including the situation when their value is less than the reserve price.

4.1.1 The optimal reserve price for consumers also maximizes social welfare

In this section we point out an important feature of our model: the interests of producers are aligned in our model. Specifically, we show that the sum of advertiser profit and search engine profit is equal to twice consumer surplus. Hence, welfare maximizing and consumer-surplus maximizing policies coincide. 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.¹¹ Consumer surplus and welfare are maximized for the same reserve price, and given any bidding behavior by advertisers and 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

The alignment of social welfare and consumer surplus has a consequence that will be useful in characterizing socially optimal policies.

Corollary 1 *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.*

¹¹Note that this result does not depend on F being uniform.

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' , so no deviation from r^W is profitable.

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 reserve price is too low and welfare is reduced.

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

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

Consumers with $s < E(q^{1:N}|q^{1:N} > r)$ will examine a link if it is presented. The mass of such consumers is $E(q^{1:N}|q^{1:N} > r)$. Their average search cost is $\frac{1}{2}E(q^{1:N}|q^{1:N} > r)$. The probability that a link is displayed is $1 - r^N$. Hence, expected consumer surplus is

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

By Proposition 6, the socially optimal reserve price is the maximizer of this expression. Taking derivatives with respect to r we find

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*

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

(ii) Suppose F and G are both uniform and that the list has one position. 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 $N \rightarrow \infty$.

One can get a very simple intuition for the optimal reserve prices in the $N = 1$ and $N \rightarrow \infty$ cases from Corollary 1. The optimal reserve price r in the no commitment model must be such that the CS-maximizing search engine is just indifferent between displaying the highest-quality ad and not displaying any ad when the search engine knows that the quality of the top firm is exactly r and consumers only know that the quality is at least r . The average gross benefit to consumers who click on the link in this case is r . Hence, the optimal reserve price r is such that expected consumer search costs in equilibrium are also equal to r . The average search cost will be one-half of consumers' conditional expectations about link quality, which is $E(q^{1:N}|q^{1:N} \geq r)$. Hence, the optimal r satisfies (1). In the $N = 1$ case, the solution to this is $r = 1/3$. When $r = 1/3$, consumer expectations will be that $q \sim U[1/3, 1]$. Hence, consumers search if and only if $s \in [0, 2/3]$ and the average search costs is $1/3$. When N is very large $E(q^{1:N}|q^{1:N} > r) \approx 1$, so the solution has $r \approx \frac{1}{2}$.

Remarks

1. The formula (1) applies for an arbitrary advertiser-quality distribution, not just when advertiser qualities are uniform on $[0, 1]$. The generalization is easier here than for several other results because with lists of length one it is not necessary to consider how consumers Bayesian update when links do not meet their needs.

2. A range of positive reserve prices will increase both search-engine profit and social welfare relative to the no-reserve benchmark. Search engine profits can be written as the product of the click-through-rate, which is $E(q^{1:N} | q^{1:N} \geq r)$, and the expected payment per click, which can be derived by integrating over the possible values of the first-order statistic.

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

Evaluating the integrals shows that profits are increasing in r for small r , so the profit-maximizing reserve price is also positive.

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 welfare.

The calculation, however, is more complicated than in the one-position case because there are two ways in which removing a link from 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 additional 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 main points still hold qualitatively, but some differences arise. It is still true that the consumer optimal and socially optimal reserve prices are positive. A small reserve price eliminates low quality websites from the list, reducing search costs and improving welfare. Also, using numerical examples, we show that for a range of distributions the search engine’s revenue maximizing reserve price is also positive. So, typically, using a reserve price can increase search engine revenues while enhancing consumer welfare.

Yet, in the general case, it is no longer true that the welfare maximizing and consumer-optimal reserve prices are the same. This result holds only when the distribution of search costs G is uniform. Despite the fact that the interests of the search engine, consumers, and advertisers are no longer perfectly aligned, setting a reserve price can usually improve both consumer surplus and revenues.

Also, for some pathological distributions, it may happen that revenue falls if the reserve price rises from zero to a small amount. This is a quite unusual result in auction settings, and somewhat unintuitive. In most auction models, revenue is increasing with the reserve price when the latter is close to zero. The reason this may happen, is that a reserve price can raise the inferred quality of the M^{th} link, making this position more attractive. This causes advertisers to bid less for the $M - 1^{st}$ spot. In turn, this reduces bids for the $M - 2^{nd}$ spot, and so on, reducing all higher bids. Still, it takes very special distributions F and G for this effect to dominate the usual revenue increasing effects of auction theory. Using a series of numerical examples, we show that this effect is quite rare, and more of a theoretical curiosity than a practical concern.

When considering reserve prices, the actual number of ads displayed may vary according to how many of the realized advertiser’s qualities are above r . One may argue whether it is reasonable for consumers to incorporate this information when inferring the quality of websites. For small lists it seems more natural that, when search results are displayed, consumers observe this number just by glancing on the screen. In this case, the \bar{q}_k will

depend on the number of links displayed. Yet, if the list is big enough that they have to scroll down to obtain this information, it is reasonable to assume that consumers do not observe this number. As it turns out, results are similar for both specifications. For simplicity, we assume the latter, so that the \bar{q}_k do not vary with realizations of q , and are given by

$$\begin{aligned}\bar{q}_k &= E(q^{k:N} | z^1 = \dots = z^{k-1} = 0, q^{k:N} \geq r) \\ &= \frac{\int_r^1 x f^{k:N}(x) \text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} dx}{\int_r^1 f^{k:N}(x) \text{Prob}\{z^1 = \dots = z^{k-1} = 0 | q^{k:N} = x\} dx}\end{aligned}\tag{2}$$

4.2.1 Consumer optimal and socially optimal reserve prices are positive

Reserve prices affect welfare in two ways. First, they exclude bidders with low quality from the list. Second, they change consumer behavior through the \bar{q}_k . We will show that, for a small change near $r = 0$, the former effect dominates, and raising reserve prices increases both consumer surplus and total welfare. The intuition is that, since the consumers whose behavior changes were indifferent between clicking or not on the next website, the latter effect is much smaller.

Proposition 9 *Both consumer surplus and total welfare are maximized by strictly positive reserve prices.*

Proof: First consider the effect on consumer surplus. We will sign the effect of a small increase in r starting from $r = 0$. Assuming, for now, consumer behavior is fixed, we show that the effect of changes in the list of displayed links is beneficial. Supposing the M^{th} advertiser's quality was lower than r (which happens with probability of the order of r^{N-M+1}), he gets excluded from the list. This generates expected benefits of at least the order of r^{N-M+1} .

As for the effect from the change of consumer behavior, consider the expression 2. The integrand on the numerator has order $N - M + 1$, and on the denominator order $N - M$. So changes on \bar{q}_k are of the order of r^{N-k+1} . As consumers with $s = \bar{q}_k$ were initially indifferent between clicking or not on the k^{th} link, consumer welfare changes by at most the order of $r^{2(N-M+1)}$. Therefore the net effect of a small increase in r is always beneficial to consumers.

Now consider the effect on gross producer surplus $GPS = \text{advertiser profits} + \text{search-engine fees}$. This is equal to the total amount of needs met. Again, a small increment to r will have two effects. First, the \bar{q}_k will increase by the order of r^{N-k+1} . This will increase

the number of clicks performed, and therefore the number of needs met, which can only increase GPS .

The downside (for GPS) of reserve prices is that they exclude advertisers with very low q . This happens with a probability of the order of r^{N-M+1} . As this advertisers have a probability of less than r of meeting a need, this latter effect has order r^{N-M+2} . So GPS decreases by at most the order of r^{N-M+2} . As consumer surplus increases by at least the order of r^{N-M+1} , total welfare also increases.

QED

4.2.2 Reserve prices may decrease revenue

Unlike welfare, revenue may actually decrease by increasing reserve prices from 0 to some small value. The reason this may happen, is because increasing r increases \bar{q}_M . This will make the M^{th} position more attractive, reducing bids for the $M-1^{st}$ position. Because bids depend recursively on lower bids, this will cause a reduction on all higher bids. There are, of course, many other effects in play. The most important ones are that reserve prices may increase a bid if the next lowest bidder had a value below r , and that total clicks increase with the \bar{q}_k . Indeed, these two positive effects will usually dominate. As the numerical examples show, it is quite rare for revenue to decrease with the reserve price near 0.

We now consider an informal example that illustrates this effect. Suppose $M = N$, and G is concentrated just above \bar{q}_N . That is, $G(q)$ is 0 for $q \leq \bar{q}_N$ and $G(q) = 1$ for $q > \bar{q}_N$. Formally, this is not a distribution function, but it will be clear from this discussion how to obtain a rigorous example. Suppose, at first, that there is no reserve price. As $G(\bar{q}_N) = 0$, an advertiser on the last spot N will get no clicks, being effectively out of the list. So firms will bid up to their true qualities $b^N = q^N$ to get into slot $N-1$. Now, suppose the reserve price is raised by a small amount. This increases \bar{q}_N , so that now every consumer clicks on every link, until he is served. In particular, this greatly reduces the incentives to bid for position $N-1$, and instead of their true values firms shade their bids to $b^N = q^N - (1 - q^N)(q^N - r)$. Since the recursion defining the other bids as a function of b_N is not changed, this will decrease all higher bids, greatly reducing revenue.

Although this example is not rigorous, its basic mechanics work whenever g is very large around \bar{q}_N . For an actual atomless distribution G , revenues will not jump down discontinuously as in the informal example. But, if $g(\bar{q}_N)$ is large enough, revenues will decrease with r near $r = 0$. We therefore have the following proposition:

Proposition 10 *There exist distributions F and G for which revenue is decreasing in*

respect to reserve prices in a neighborhood of $r = 0$.

Still, as g is a density, it averages 1 over the unit interval. So it may only take high values in a very small area. Because \bar{q}_M depends only on F , very special combinations of F and G are required for revenue to fall with the use of a small reserve price. This observation, and the many numerical examples, suggest that this possibility is quite rare in practice. Search diversion

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 the 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 each firm i bids the maximum per-click fee b_i it is willing to pay. 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 11 *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.*

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

to 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.)¹²

Proposition 12 *Suppose that $M = 2$ and $N > 2$. Suppose all consumers have $s < \bar{s}$. Then the click-weighted auction has an equilibrium in which the firm with the lower δ drops 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 even though it is on the bottom of the list (provided that it can potentially meet their needs). Given this, there is no incentive for firm 1 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 13 *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

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

¹²It suffices to set $\bar{s} = E(q^k | \delta^k = 1, \delta^1 = \dots = \delta^{k-1} = \delta)$.

This condition is satisfied for the given bidding function.

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

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

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 ($M = 2$).

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 .¹³ 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 the fact that customers served by the two sites are distinct), both firms will drop out

¹³As in section 5.1, the b_i are per-click bids and per-click payments of firm k is the $k + 1^{st}$ highest bids adjusted for the weight difference (if a difference exists).

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 $\delta_1(1 - q_{1A})E(q_{1B}|q_{1B} < q_{1A}, q_{1B} = wq_2) = \delta_1(1 - q_{1A})q_{1A}/2$.¹⁴ 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 14 *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 , e.g. for $\delta_1 = 1/2$.

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 δ_2 , respectively so these standard weights would favor firms 1A and 1B for any $\delta_1 > 1/2$.

¹⁴Conditioning on $q_{1B} = wq_2$ is irrelevant because conditional on $wq_2 < q_{1A}$, wq_2 is uniform on $[0, q_{1A}]$.

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.

5.6 Equivalence with pay per action auctions

New technology being developed now allows search engines to track how many sales are made by advertisers. This leads to the question of whether it is advantageous to charge producers by sale made, instead of charging them per click. There are many reasons why pay per action pricing might be desirable for advertisers; for example, it might decrease risk for them if their ads are shown on a large network, where click fraud may be a problem. We do not model those benefits here.

Instead, we will show that charging per action is equivalent to charging per click, if the per-action bid is weighted by the conversion rate coherently. We mean equivalent in the sense that both auctions have equilibria with the same link ordering, payments, and consumer behavior. In the pay per action model, the profit of a firm of type (q, δ) in position

k is

$$(1 - q_1 \delta_1) \cdots (1 - q_{k-1} \delta_{k-1}) \delta G(\bar{q}_k) \left(q - q \frac{b_{k+1}}{\tau} \right),$$

where τ is the estimated conversion rate, so that $\tau = q\delta$. Notice that, defining τ this way, the payoff function becomes exactly the same in click-weighted pay per click auctions and in pay per action auctions. Also, non weighted auctions are only a particular case of weighted auctions when all $\delta = 1$, so this observation is also valid for non weighted auctions. We note this observation in the following proposition.

Proposition 15 *Equilibria of the pay per action model have the same ordering, payments, and consumer behavior as in the pay per click model.*

Still, this equivalence result should not hold in a more complete model. For example, if the τ were allowed to be arbitrary, this equivalence would no longer hold. By choosing $\tau = q\delta$, we are thinking of an stationary situation, where the estimated conversion rate is consistent with the conversion rate received in equilibrium. But in a richer dynamic model, this may not be the case. In the next section we extend the basic model to allow firm to choose the informativeness of their advertisements. With this modification, pay per click and pay per action auctions give different incentives for firms to shroud their ads, and the two formats are no longer equivalent.

6 More auction design

6.1 Search-diverting Sites

In the U.S. market three main firms provide search services and sell keyword advertising: Google, Yahoo!, MSN.¹⁵ Google is substantially larger than the others and earns higher revenue per search. When keyword searches 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

¹⁵The fourth largest search engine, Ask.com, displays ads provided by Google.

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 would usually only 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.¹⁶ 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 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

¹⁶Suppose also that their beliefs about quality before clicking on the links are as in our base model with $N + 1$ firms.

if the N regular advertisers were the only bidders.

6.2 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.¹⁷ 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.

Given this model, we have

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

Some observations from this 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 . Starting from $\rho = 1$, even a small decrease in the probability of relevant searches causes a large decrease in a consumer's posterior belief in relevance after failing to have his or her needs met at the first k sites.
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 . This implies that a consumer's posterior belief in relevance falls very quickly after failing to have his or her needs met at the first site.

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

$$E(q^{k:N} \mid z^1 = \dots = z^{k-1} = 0) = \tag{3}$$

¹⁷This 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.

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

This decreases more quickly with increases in k when $\rho < 1$ as compared to $\rho = 1$. If ρ is far from one, then consumers will only search a small number of ads.

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.

A more basic theme of our paper is that an important role of sponsored link auctions is to provide information about the sponsored links which allows consumers to search more efficiently. Sorting firms on the basis of their probability of meeting consumers' needs is part of this, but in principle one could imagine many other search engine designs that present consumers with much more information. This could be an interesting area for pure and applied research.

Appendix

Proof of Proposition 14

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,

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

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

$$E(CS(w)) = \alpha \left(\left(1 - \frac{1}{3w}\right) \frac{8w - 3}{12w - 4} + \frac{1}{3w} \left(\frac{3}{4} + \frac{11}{42}\right) \right) + (1 - \alpha) \left(\left(1 - \frac{1}{3w}\right) \frac{6w^2 - 1}{12w^2 - 4w} + \frac{1}{3w} \cdot 0 \right)$$

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

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

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.
- Edelman, Benjamin and Michael Ostrovsky (2005): “Strategic Bidder Behavior in Sponsored Search Auctions,” mimeo, Harvard University and Stanford University.
- 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.
- 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.