

POSITION AUCTIONS WITH CONSUMER SEARCH*

SUSAN ATHEY AND GLENN ELLISON

This article 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* Codes: D44, L86, M37.

I. INTRODUCTION

Google, Yahoo!, and Microsoft allocate the small “sponsored links” at the top and on the right side of their search engine results via similar auction mechanisms. Sponsored-link auctions have quickly become one of the more practically important topics in the economics of auctions, as annual revenues now surpass \$30 billion. They have also quickly become an important topic in the economics of advertising: they have driven the recent growth of online advertising, which is having dramatic effects both on products that are now heavily sold online and on the competing media that are suffering in the competition for advertising dollars. In this article, we address issues of relevance to both fields by developing a model of sponsored-link advertising that incorporates both standard auction-theoretic and two-sided market considerations.

There has been a recent burst of academic papers on sponsored-search auctions spurred by the importance of the topic and by some very elegant results. This literature has coined the phrase “position auctions” to describe the particular multigood auction in which per-click bidding is used to auction off n asymmetric objects with unidimensional bids.¹ A striking result, derived in [Aggarwal, Goel, and Motwani \(2006\)](#), [Edelman](#),

*We thank Ben Edelman, Leslie Marx, John Morgan, Ilya Segal, Hal Varian, various seminar audiences, and four anonymous referees for their comments and Eduardo Azevedo, Eric Budish, Stephanie Hurder, Scott Kominers, and Dmitry Taubinsky for exceptional research assistance. This work was supported by National Science Foundation grants SES-0550897 and SES-0351500 and the Toulouse Network for Information Technology.

1. See [Ausubel and Cramton \(2004\)](#) and [Cramton, Shoham, and Steinberg \(2007\)](#), for broader discussions of multigood auctions.

© The Author(s) 2011. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2011) 126, 1213–1270. doi:10.1093/qje/qjr028.

Advance Access publication on August 4, 2011.

Ostrovsky, and Schwarz (2007), and Varian (2007), is that the generalized second price (GSP) position 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 Vickrey–Clarke–Groves (VCG) mechanism and thus does not induce truthful bidding, but nonetheless results in the same outcome as the VCG mechanism in an interesting class of environments. A number of subsequent papers have extended the analyses in various important dimensions, such as allowing for reserve prices, the use of weights to account for asymmetric click-through rates, and considering more general relationships between positions and click-through rates.²

Most of the literature, however, is squarely auction-focused and continues to abstract away from the fact that the “objects” being auctioned are advertisements.³ 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 article, we develop this line of analysis. By incorporating consumers into our model, we are able to answer questions about how the design of the advertising auction marketplace affects overall welfare, as well as the division of surplus between consumers, search engines, and advertisers. Our framework allows us to provide new insights about reserve price policies, click-through weighting, fostering product diversity, advertisers’ incentives to write accurate ad text, and effects of different bidding mechanisms.

Section II 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.

2. See, among many others, Edelman and Schwarz (2010) on reserve prices, Lehaie (2006), Liu and Chen (2006), and Liu, Chen, and Whinston (2010) on click-weights, and Borgers et al. (2006) and Aggarwal et al. (2008) on alternate click-position relationships.

3. As discussed, Chen and He (2006) is a noteworthy early 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.

Section III presents some basic results on search, welfare, and the economic role of sponsored search advertising. We characterize optimal consumer search strategies. We note that a search engine that presents sponsored links should be thought of as an information intermediary that contributes to welfare by providing information (in the form of an ordered list) that allows consumers to search more efficiently, and we present some calculations that quantify the welfare benefits.

Section IV contains our equilibrium analysis of the sponsored-link auction. Because the value of being in any given position on the search results page depends on the qualities of all of the other advertisers, the auction is no longer a private-values model and hence does not fit within the framework of [Edelman, Ostrovsky, and Schwarz \(EOS\) \(2007\)](#). We note, however, that the common value elements of our model are perhaps surprisingly easy to deal with—the analysis of EOS can be adapted with only minor modifications. We are able to provide explicit formulas describing a symmetric perfect Bayesian equilibrium with monotone bidding functions, and we discuss a number of properties of the equilibrium.

We then turn to auction-design questions. These are obviously of practical interest to firms conducting sponsored-link auctions and to policy makers who must interpret the actions being taken in what is a highly concentrated industry. We find them interesting from a theoretical perspective as well, because the fact that Google is auctioning advertisements rather than generic objects brings up a host of new concerns. Any changes to the rules for selecting ads will affect what consumers infer about the quality of each displayed ad, which in turn affects the value of winning each of the prizes being auctioned. Effects of this variety can substantially change the way one thinks about search engine policies.

Our first auction design section, Section V, focuses on reserve-price policies. Recall that in standard auctions with exogenous values, reserve prices raise revenue for the auctioneer, but this comes at a welfare cost—some potential gains from trade are not realized. In contrast, in our model, reserve prices can enhance total social surplus, and in some cases can even be good for advertisers. The reason is that reserve prices 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 they can increase the number of links that are examined in

equilibrium.⁴ The section also focuses on conflicts (or the lack thereof) between the preferences of consumers, advertisers, and the search engine. Indeed, our analysis begins with an observation that when consumer search costs are uniformly distributed (but other aspects of the model are left quite general) there is a perfect alignment of consumer-optimal and socially optimal policies. This observation turns out to be a nice way to bring out several insights: we derive results on consumer and social welfare; we use it as a computational tool; and we note that it also implies that there is an inherent conflict between the search engine and its advertisers—any departure from the socially optimal policy that increases search engine profits must do so by reducing advertiser surplus. We also present a number of results concerning what does and does not generalize to the case when consumer search costs are drawn from general distributions.

Section VI examines click-weighted auctions similar to those used by Google, Yahoo!, and Microsoft. Google's introduction of click-through weighting in 2002 is regarded as an important competitive advantage and Yahoo!'s introduction of click-through weights into its ranking algorithm in early 2007 ("Panama") was highly publicized as a critical improvement.⁵ It is intuitive that weighting bids by click-through rates should improve efficiency—surplus is only generated when consumers click on links. EOS note briefly (at the end of Section III) that their efficiency result extends to establish the efficiency of click-weighted auctions when click-through rates are the product of a position effect and an advertiser effect.⁶ Our analysis places some caveats on this conventional wisdom about efficiency. In the presence of search costs, we show that the click-weighted auction does not necessarily generate the right selection of ads—general ads may be displayed when it would be more efficient to display ads that serve a narrower population segment well. There can also be welfare losses when asymmetries in the click-through weights

4. In this respect, our model is related to that of Kamenica (2008) which develops a rational alternative to "behavioral" explanations for why demand sometimes rises when consumers are presented with a smaller choice set.

5. Eisenmann and Hermann (2006) report that Google's move was in part motivated by a desire for improved ad relevance: "according to Google, this method ensured that users saw the most relevant ads first."

6. Lehaie (2006), Liu and Chen (2006), and Liu, Chen, and Whinston (2010), note that although click-weighting is efficient, a profit-maximizing search engine will typically want to choose different weights.

make the ordering of the ads less informative about quality. Finally, we note that the introduction of click-weighting can create incentives for firms to write misleading and overly broad text. The intuition for the latter result is that even though firms pay per click, in a click-weighted auction, firms that generate more clicks on average must pay less per click to maintain their position, and so they have no incentive to economize on consumer clicks. The result is also robust to the use of pay-per-action pricing models, so long as the auction is action-weighted.

A number of results are described only informally in the text with the formal models and analysis left to the appendixes. Appendix A examines a special case in which we can derive more explicit versions of several results. Appendixes B and C develop the models discussed in Sections V.F and VI, respectively. Appendix D contains some omitted proofs. We direct readers who prefer an integrated discussion to the working paper version of this article to [Athey and Ellison \(2009\)](#).

As noted, our article contributes to a rapidly 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 that 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 the results would carry over to click-weighted auctions.

We have already mentioned a number of publications that extend the analysis in various directions. [Edelman and Schwarz \(2010\)](#) were the first to analyze optimal reserve prices. They present both theoretical and numerical analyses, including a demonstration that the GSP auction with a single, optimally chosen reserve price is an optimal mechanism. Our work departs from theirs in our consideration of the feedbacks between auction rules, consumer expectations, and the value of advertising slots, and our model provides a motivation for reserve prices that vary with position. Two other papers are noteworthy for considering more general click-through processes and presenting empirical results. [Borgers, Cox, Pesendorfer, and Petricek \(BCPP\) \(2007\)](#)

extend the standard 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.⁷ Jerzierski and Segal (2009) develop a model in which consumers have more general preferences across bundles of ads and provide both reduced form empirical results relevant to our article, for example, noting that consumers do not always search in a top-down manner and that clicks on lower ads are affected by the quality of higher ads, and structural estimates of consumer utility parameters.

Chen and He (2006) previously developed a model that introduced several of the key elements of our model. They assumed that consumers have needs, that advertisers have different valuations because they have different probabilities of meeting consumers' needs, and that consumers search optimally until their need is satisfied. They also included some desirable elements that we do not include: they endogenize the prices advertisers charge consumers and allow firms to have different production costs.⁸ Our primary departure from their framework is our consideration of incomplete information: we assume that advertisers' qualities are drawn from a distribution and not known to consumers (and other advertisers). This assumption plays a central role in many of our analyses. For example, most of our auction design analyses hinge on how the design affects the information consumers get about firm qualities and thereby influences consumer search. Our article also differs from theirs in that much of our article is devoted to topics, for example, consumer welfare, reserve prices, and click weighting, that they do not address.

Several more recent papers have also examined issues that reflect that search engines are auctioning advertisements. White (2008) and Xu, Chen, and Whinston (2009) develop models that

7. 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. BCPP also contains an empirical analysis which includes methodological innovations and estimates of how value-per-click changes with position in Yahoo! data. Chen, Liu, and Whinston (2009) develop a model (without consumer search) which treats the fraction of clicks allocated to each advertiser as a design variable.

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

include both organic and paid search results. Xu, Chen, and Whinston (2008) develop a model in which advertisers are also competing in prices for the goods they are advertising and provide a number of interesting observations about how this may interact with the willingness to bid for a higher position.

II. 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 $[0, 1]$.

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 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.⁹ The k th highest bidder pays the $k+1$ st highest bid for each click it gets.¹⁰ To avoid some of the complications that arise in 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

9. EOS show that in their model the equilibrium of such an ascending auction is also the lowest-revenue envy-free Nash equilibrium of a complete-information, simultaneous-move bidding game. Because our model has endogenous click-through rates, the envy-free concept would require some modification to be applied to our model.

10. Note that this model differs from the real-world auctions by Google, Yahoo!, and MSN in that it does not weight bids by click-through weights. We discuss such weighted auctions in Section VI. 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, for example, where the bidders are retailers with similar business models.

will next drop out if no other firm has yet dropped out.¹¹ In the first stage, which we call stage $M + 1$, the firms simultaneously submit bids $b_{M+1}, \dots, b_{M+1N} \in [0, \infty)$ specifying a per-click price they are willing to pay to be listed on the screen. The $N - M$ lowest bidders are eliminated.¹² Write b^{M+1} for the highest bid among the firms that have been eliminated. In remaining stages k , which we will 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 baseline model. First, advertisers are symmetric except for their probability of meeting a need: profit-per-action is the same for all firms, and we do not consider the pricing problem for the advertiser. By focusing on the probability of meeting a need rather than pricing, we focus attention on the case—which we believe is most common on search engines (as opposed to price comparison sites)—where search phrases are sufficiently broad that many different user intentions correspond to the same search phrase, and so the first-order difference among sites is whether they are even plausible candidates for the consumer's needs. If we allowed for firms to set prices but required consumers to search to learn prices, firms would have an incentive to set monopoly prices, following the logic of [Diamond \(1971\)](#) and [Chen and He \(2006\)](#); thus, the main consequence of our simplification is the symmetry assumption. Generalizing the model to allow for heterogeneous values conditional on meeting the consumer's need 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; we leave that for future work.

11. Two 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 multiunit auctions.

12. If 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.

Our baseline model assumes that advertisers receive no benefit when consumers see their ad but do not click on it. Incorporating such impression values (as in BCPP) places a wedge between the externality created by a firm and its value to being in a given position. In addition, consumers get no information about whether listed firms are more or less likely to meet their needs from reading the text of their ads. In Section VI we consider extensions of our model where the ad text of firms is informative, leading to heterogeneous click-through rates, and its accuracy is endogenous; we also consider the effect of advertiser value for impressions in this context.

III. CONSUMER SEARCH AND THE ECONOMIC ROLE OF SPONSORED-LINK AUCTIONS

In this section we bring out the idea that search engines auctioning sponsored links are information intermediaries and one way in which they contribute to social welfare is by making consumer search more efficient. We do so by characterizing consumer welfare with sorted and unsorted lists. This section also contains 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.

III.A. Consumer Search and Bayesian Updating

Suppose that advertisers' bids in the position auction are strictly monotone in q .¹³ 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, $q^{1:N}$.¹⁴ The expected payoff from any additional click must be determined by Bayesian updating: the fact that the first website didn't meet a consumer's need leads them to 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

13. We will see in Section IV that our model does have an equilibrium in which this occurs.

14. 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.

to 1 with these probabilities. Define \bar{q}_k to 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:

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

PROPOSITION 1. 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 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 .) ■

The special case of the model when the quality distribution F is uniform is surprisingly tractable: there are simple closed-form expressions for the \bar{q}_k and this makes it possible to give more explicit formulas characterizing consumer strategies, welfare, and so on. Appendix A presents these results.

One way to motivate interest in equilibrium consumer search rules would be to assume that consumers are highly sophisticated and know the number N of bidders, the distribution F of firm qualities, and do all the Bayesian updating. Note, however, that such sophistication is not really necessary. In practice, consumers just need to have learned the probabilities $\bar{q}_1, \bar{q}_2, \dots, \bar{q}_M$, with which they meet their needs when clicking on each link. In practice, consumers will not get this fully right. For example, they will not know the exact number of potential advertisers

N and have learned different values of the \bar{q}_k applicable to each N . But we think consumers will have some ability to know whether a query is likely to have a small or large number of potential advertisers and to have roughly learned the probabilities conditional on some such distinctions. In our later auction design discussions we are motivated by similar reasoning: we think that real world consumers know very little about reserve price policies, but will eventually react to policy changes as the equilibrium theory predicts because they will eventually learn that links have become more or less likely to meet their needs.

III.B. Welfare Gains from Information Provision

Lists of sponsored links provide consumers with two types of information. They identify a set of links that may meet the consumer’s need, and they provide information on relative quality that helps consumers search through this set more efficiently. To bring out this latter source of welfare gains, it is instructive to consider how consumer search would differ if advertisements were instead 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 2. 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

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

If ads are sorted in order of decreasing quality then

$$E(CS(s)) = \begin{cases} 0 & \text{if } s \in [\bar{q}_1, 1] \\ \bar{q}_1 - s & \text{if } s \in [\bar{q}_2, \bar{q}_1] \\ \dots & \dots \\ (\bar{q}_1 - s) + (\bar{q}_2 - s)(1 - \bar{q}_1) + \dots & \dots \\ + (\bar{q}_k - s) \prod_{j=1}^{k-1} (1 - \bar{q}_j) & \text{if } s \in [\bar{q}_{k+1}, \bar{q}_k] \end{cases}$$

Proof. When links are not sorted equilibrium search is straightforward. Consumers with $s > \bar{q}$ never click. Consumers with lower s will click on links until their need is met or they exhaust the list. They 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 $(\bar{q} - s)(1 + (1 - \bar{q}) + (1 - \bar{q})^2 + \dots + (1 - \bar{q})^{M-1})$.

The payoffs in the sorted list are computed similarly, but reflect that the searches occur in a top-down manner and stop endogenously as described Proposition 1. ■

Comparing the two expressions brings out the welfare gains. Consumers with $s \in [\bar{q}, \bar{q}_1]$ get no utility at all from an unsorted list but positive utility from a sorted list because the higher quality of the top links makes clicking worthwhile. Consumers with low search costs also benefit from the sorted list because they find what they want more quickly. One case in which this is particularly clear is when N is very large. In that case, consumer surplus is approximately $1 - s$ with a sorted list (provided F has full support so that the highest order statistic is very close to 1) and approximately $1 - \frac{s}{\bar{q}}$ with an unsorted list because approximately $\frac{1}{\bar{q}}$ searches are needed to find an advertiser that can meet the consumer's need. Appendix A gives some more explicit expressions and a comparative graph for the case where the quality distribution, F , is uniform.

IV. EQUILIBRIUM OF THE SPONSORED SEARCH AUCTION

In this section we solve for the equilibrium of our base model taking both consumer and advertiser behavior into account. 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.¹⁵

IV.A. *Equilibrium in the Bidding Game*

Consider our formalization of an "ascending auction" in which N firms bid for the $M < N$ positions. When clicked on, firm i will be able to meet a consumer's need with probability q_i . We have exogenously fixed the per-consumer profit at 1, so q_i is like the value of a click in a standard position auction model.

15. In 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.

Although one can think of our auction game as being analogous to the EOS model, but with endogenous click-through rates, the auction part of our 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.¹⁶ The equilibrium derivation, however, is similar to that of EOS.

Our first observation is that, as in the EOS 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 0 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 .¹⁷

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

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

17. In 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.

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

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

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

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

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 3. 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

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

Proof. 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 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

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

This argument is formalized in Appendix D. ■

REMARK 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 1, 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.

REMARK 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.

REMARK 3. Bidders shade their bids less when bidding for higher positions, that is, $b^*(k, b'; q)$ is decreasing in k with b' and q fixed, if and only if $\frac{G(\bar{q}_{k-1})}{G(\bar{q}_k)}$ is decreasing in k . This can be seen most easily by rewriting Equation 1 as

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

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 $\frac{G(\bar{q}_{k-1})}{(1 - q^{k-1:N})G(\bar{q}_k)}$, and industry sources report that it decreases moderately to rapidly in k , which is consistent with less bid shading at the top positions. However, because $q^{k:N}$ is also declining in k , the declining click-through rates could also be due to rapidly decreasing quality lower down the list.

REMARK 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.

V. RESERVE PRICES

We now turn to questions of auction design. Such questions are of practical interest for three reasons: they are of interest to firms designing auctions; auction design also affects the welfare of consumers and advertisers; and antitrust and regulatory authorities will need to understand both the incentives and welfare effects to interpret actions in what is a concentrated industry. Auction design questions in our model are also interesting theoretically because the standard principles of auction design can be substantially altered by the fact that changes to the auction design affect consumer beliefs about the quality of sponsored links and thereby affect the “values” of prizes that are being auctioned. In this section we discuss a common and important design decision: the setting of reserve prices.

In a standard auction model, reserve prices increase the auctioneer’s expected revenues. At the same time, however, they reduce social welfare.¹⁸ 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,

18. The reduction in the gains from trade could inhibit seller or buyer entry in a model in which these were endogenous, so reserve prices might not be optimal in such models. See [Ellison, Fudenberg, and Möbius \(2004\)](#) for a model of competing auction sites in which this effect would be important.

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.

Most of this section examines the special case of consumer search costs being uniformly distributed on $[0, 1]$ (in fact, our results generalize to the case where the search cost distribution takes the form $G(s) = s^d$, but we focus on the uniform case for simplicity of exposition).¹⁹ By specializing the search cost distribution, we are able to derive a neat theorem on the alignment of interests that is very general on other dimensions. This allows us to provide some complete characterizations, and it is also a nice way to highlight forces that remain present in more general specifications. Section V.F contains some results on general search cost distributions, illustrating some results that are robust and highlighting forces that can make others change.

V.A. *An Alignment Theorem*

In this section, we assume that the distribution G of search costs is uniform (without making assumptions about the advertiser quality distribution). 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. The intuition for the result can be given as follows. First, producer surplus and gross consumer surplus (ignoring search costs) are both determined by the probability that consumers have their needs satisfied. Second, consumers search optimally, so that their search intensity and thus average search costs increase with their (perceived) probability that their need will be met. Third, when the search cost distribution takes the form $G(s) = s^d$, average search costs increase with the probability that their need will be met in a constant proportion. Finally, the third fact implies that consumer surplus net of search costs is also proportional to the probability that needs are satisfied.

19. We also restrict our analyses to equilibria like those described in the previous section in which firms use strictly monotone bidding strategies and bid their true value when they will not be on the sponsored link list.

PROPOSITION 4. 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: $GCS = CS + \text{Search Costs}$.²⁰ Write GPS for the gross producer surplus: $GPS = \text{Advertiser Profit} + \text{Search Engine Profit}$. 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. ■

REMARK 1. Note that the alignment result does not require any assumption on the distribution F of firm qualities and is thus fully general in this dimension.

REMARK 2. The alignment result does not depend on the assumption that consumers and advertisers both receive exactly one unit of surplus from a met need. If advertisers receive benefit α from meeting a consumer's need, then $E(W) = (\frac{1}{2} + \alpha)E(GCS) = (1 + 2\alpha)E(CS)$.

REMARK 3. The alignment result with uniform search costs is a special case of a slightly more general result. If the search cost distribution is $G(s) = s^d$, then welfare and consumer surplus are proportional with $E(W) = (2 + \frac{1}{d})E(CS)$. The argument is

20. More precisely, GCS is the population average gross consumer surplus. It is a random variable, with the realized value being a function of the realized qualities. The other measures of welfare and consumer and producer surplus we discuss should be understood similarly.

similar and uses the fact that $E(s|s \leq q)$ remains proportional to q for this family of distributions. ($E(s|s \leq q) = \frac{q}{(d+1)}$.)

REMARK 4. The alignment result pertains to producer surplus, but it does not say anything about the distinct reserve price preferences of advertisers and search engines. As shown in more detail in the next section, advertisers and the search engine are typically in conflict with one another and with consumers about the level of reserve prices.

Our next result is a corollary that points out that commitment problems are absent in one particular (and infrequently studied) case: it shows that a search engine that has consumer surplus maximization as its objective function would choose the socially optimal reserve price even if the search engine lacked the ability to commit to a reserve price. One may initially wonder why we bother to point out this lack of a commitment problem. We have three motivations: it highlights a contrast between consumer surplus-maximizing search engines and search engines with other objective functions; the consumer surplus maximization objective function may have some practical relevance in the search engine industry; and the corollary also turns out to be a useful technical tool. We elaborate on these after presenting the result.

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.

Proof. Proposition 4 says directly that r^W will be chosen if the search engine has commitment power. The fact that a CS-maximizing search engine won't have a commitment problem is less obvious, but nearly as immediate once one sets up the argument.

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' . If consumers expect that the search engine is using reserve price r^W , then the CS-maximizing search engine gets payoff $CS(r^W, r^W)$

if it indeed uses r^W and $CS(r^W, q')$ if it deviates and uses reserve price q' instead. We can see that the deviation is not profitable by a simple two-step argument: $CS(r^W, q') \leq CS(q', q') \leq CS(r^W, r^W)$. The first inequality is a consequence of consumer rationality—holding the policy fixed consumers do best if they know the policy and therefore act in the way that maximizes their payoff. The second is the conclusion of Proposition 4. ■

We noted that one motivation for presenting Corollary 1 is that it highlights that profit-maximizing (or social welfare-maximizing) search engines would have a commitment problem. An intuition for this is that consumers do not internalize the potential profits that advertisers and the search engine will get when they consider whether to click on a website. Hence, *holding consumer expectations fixed*, a deviation to a reserve price slightly lower than what consumers are expecting would typically increase profits (and social welfare) by leading consumers to click on more links. In equilibrium, of course, the search engine cannot benefit from deviating from the policy that consumers believe it to be using, so we end up with an equilibrium in which the search engine is worse off. The contrast between a consumer surplus-maximizing and a social welfare-maximizing search engine is interesting: an incentive to try to increase the broader welfare measure ends up reducing social welfare relative to what happens if the search engine acts to maximize just consumer surplus.

A second motivation for presenting Corollary 1 is that the behavior of consumer surplus-maximizing firms could be directly relevant in the search engine application. Search engines are engaged in dynamic competition to attract consumers. If current market shares are sensitive to the consumer surplus a search engine provides, higher current market shares lead to higher future market shares, and future profits are important relative to forgone current profits, then designing a search engine to maximize consumer surplus may be a rough rule-of-thumb approximation to the optimal dynamic policy of a firm competing aggressively for consumers.

A final motivation for presenting Corollary 1 is that it turns out to be a useful computational tool. Corollary 1 implies that we can find the policy that maximizes the social welfare function by finding the equilibrium policy in the no-commitment model with a consumer surplus-maximizing search engine. The latter turns out to provide an easier path to some results.

V.B. *Socially Optimal Reserve Prices with One-Position Lists and Uniformly Distributed Search Costs*

To bring out the economics of setting reserve prices and the trade-offs for the welfare of participants, 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

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

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, a consumer surplus-maximizing search engine will only display a link if the net benefit to consumers 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 in the no-commitment model therefore requires that $r = \frac{1}{2}E(q^{1:N} | q^{1:N} > r)$. By Proposition 4 and Corollary 1 we then have:

PROPOSITION 5. Suppose that the list has one position and that the distribution of search costs is uniform. Then, consumer surplus and social welfare are maximized for the same reserve price. The optimal r satisfies

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

REMARK 1. Note that Equation 2 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 1 it is not necessary to consider how consumers Bayesian update when links do not meet their needs.

REMARK 2. For any quality distribution F with full support on $[0, 1]$ we will have $E(q^{1:N} | q^{1:N} > r) \approx 1$ for N large. Hence, for N large the optimal reserve price will be close to $\frac{1}{2}$.

REMARK 3. The probability that a link is displayed is $1 - F(r)^N$. The mass of consumers who will click on a link if one is

displayed is $E(q^{1:N}|q^{1:N} > r)$. Hence, a formula for expected consumer surplus given an arbitrary reserve price r is:

$$\begin{aligned} E(CS) &= (1 - F(r)^N) G(\bar{q}_1(r)) (E(q^{1:N}|q^{1:N} > r) \\ &\quad - E(s|s < E(q^{1:N}|q^{1:N} > r))) \\ &= \frac{1}{2} E(q^{1:N}|q^{1:N} > r)^2 (1 - F(r)^N). \end{aligned}$$

V.C. Welfare and the Distribution of Rents

We noted that expected consumer surplus and expected producer surplus are proportional, and hence maximized at the same reserve price. The “producer surplus” in that calculation is the sum of search engine revenue and advertiser surplus. In this section we note that these two components of producer surplus are less aligned: search engines may prefer a reserve price much greater than the social optimum and advertisers may prefer a reserve price much smaller than the social optimum.²¹

The producer surplus in our model, $GPS(r)$, is a sum of two terms: advertiser surplus $AS(r)$, and search-engine profit $SR(r)$. First, we note that a general result on preference conflicts is another simple corollary of our alignment theorem.

PROPOSITION 6. Consider the model with uniformly distributed search costs. Suppose that the reserve price r^π that maximizes the search engine’s expected profit does not coincide with the socially optimal reserve price r^W . Then, both consumers and advertisers are worse off under the profit-maximizing reserve price, that is, $E(CS(r^\pi)) < E(CS(r^W))$ and $E(AS(r^\pi)) < E(AS(r^W))$.

Proof. If r^π and r^W do not coincide then $E(W(r^\pi)) < E(W(r^W))$. The fact that expected consumer surplus is lower at r^π than at r^W is an immediate corollary of Proposition 4: $E(CS(r)) = \frac{1}{3}E(W(r))$. The advertiser surplus result is nearly as easy. The fact that $E(GPS(r)) = \frac{2}{3}E(W(r))$ for any r gives $E(GPS(r^\pi)) < E(GPS(r^W))$. Gross producer surplus is the sum of advertiser surplus and search engine profit. The latter is higher at r^π than r^W , so it must be that advertiser surplus is lower at r^π than at r^W . ■

21. Edelman and Schwarz (2010) illustrate this preference divergence in simulations for the case they study, exogenous click-through rates.

In the special case where the search engine displays only a single link we can get some additional insight by writing out an explicit division-of-surplus function:

$$\begin{aligned} E(SR(r)) &= (1 - F(r)^N) G(\bar{q}_1(r)) E(\max(q^{2:N}, r) | q^{1:N} \geq r) \\ &= \tau(r) E(GPS(r)) \end{aligned}$$

$$E(AS(r)) = (1 - \tau(r)) E(GPS(r)),$$

where the division-of-surplus function $\tau(r)$ is given by

$$\tau(r) \equiv \frac{E(\max(q^{2:N}, r) | q^{1:N} \geq r)}{E(q^{1:N} | q^{1:N} \geq r)},$$

The $\tau(r)$ function has $r \leq \tau(r) \leq 1$ and hence satisfies $\lim_{r \rightarrow 1} \tau(r) = 1$ for any quality distribution, that is, the search engine gets almost all of the (very small) surplus when the reserve price is very high. For many distributions, including the uniform, the $\tau(r)$ function is strictly increasing on $[0, 1]$, although this is not true for all distributions.²² When the division of surplus function $\tau(r)$ is increasing in r , the profit-maximizing reserve price for the search engine is greater than the social optimum.²³

In Section V.A we noted that consumer surplus maximization might be a reasonable approximation to the objective function of a search engine in a competitive dynamic environment. If a search engine instead maximizes a weighted average of consumer surplus and profit and the consumer surplus and profit functions are single-peaked, then a lesser weight on consumer surplus will result in a reserve price that is worse for social welfare, consumers, and advertisers, but better for the search engine. This comparison could be relevant for evaluating changes in industry structure that make search engines less willing to invest in attracting consumers instead of maximizing short-run profits.²⁴

22. The $\tau(r)$ function is also increasing when $F(r) = r^\alpha$ for any $\alpha > 0$. A simple example to show that it is not always increasing would be a two-point distribution with mass $1 - \epsilon$ on $q = \frac{1}{2}$ and mass ϵ on $q = 1$. For this distribution we have $\tau(\frac{1}{2}) \approx 1$ and $\tau(\frac{1}{2} + \epsilon) \approx \frac{1}{2}$.

23. More precisely, the profit-maximizing reserve price is always weakly greater than the social optimum if $\tau(r)$ is weakly increasing and strictly greater if one assumes other regularity conditions. For example, if $E(CS(r))$ is differentiable at the social optimum and $\tau(r)$ has a nonzero derivative at this point.

24. For practical application one would also want to augment the model to include fixed costs that advertisers must pay to participate. There would then be additional effects that stem from the search engine's willingness to invest in attracting advertisers.

Appendix B presents a figure illustrating the magnitudes of the welfare changes in a couple examples.

V.D. Optimal Reserve Prices 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 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 . The discussion shows:

PROPOSITION 7. 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

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

V.E. More General Policies

In the foregoing analysis 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 noted in the previous section) is that reserve prices should be set so that the reserve price for the m th position is approximately (but 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.²⁵

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, due in part to different reserve prices for the top positions and the side positions.

25. This contrasts with [Edelman and Schwarz \(2010\)](#), who show that a single reserve price is optimal for all positions when clicks are exogenous.

V.F. Reserve Prices under General Distributions

The analyses of reserve prices assume that consumer search costs are uniformly distributed. When search costs are not uniformly distributed, the consumer optimal and socially optimal reserve prices will no longer exactly coincide. Appendix B presents three results illustrating how results may change with general search costs.

First, we note that the conclusion that the consumer-optimal reserve price is positive is fully robust: it holds for any search cost distribution. Second, we note that the socially optimal reserve price is not necessarily positive. An intuition for why this conclusion may change builds on our earlier comment that consumers do not take into account firm profits when they choose whether to click on a link. In some situations increasing the reserve price can (counterintuitively) make consumers less likely to click on links, and this can reduce welfare. Third, we note that the search engine optimal reserve price can also be 0. The effect that makes this possible is that by making consumers more willing to click on lower-ranked links, the institution of a reserve prices can reduce the incentive that firms have to bid for higher positions.

The examples we use to show that reserve prices can reduce search engine profits and social welfare are highly special examples. Our motivation for presenting them is not to suggest that the outcomes that occur in them are likely to occur in practice, and we emphasize that consumer and social preferences will be roughly aligned whenever the search cost distribution is approximately uniform or approximated by any CDF of the form $G(z) = z^\alpha$. Our motivation for presenting the examples is instead to illustrate the mechanisms that drive them so that readers may be aware of them in case they are important in some situations.

VI. CLICK-WEIGHTED AUCTIONS

Around 2003 Google was the first to implement a modified position auction: it assigned each advertisement i a quality score w_i ; ranked bids on the basis of the product $b_i w_i$ rather than on b_i alone; and modified per-click payments to reflect differences in the quality scores. This modification was tremendously important in practice: the textbook unweighted position auction lets obviously “wrong” bidders winning most auctions. For example, sites

selling ringtones and pornography sites might outbid camera stores for the right to be listed when consumers search for “digital cameras.” Although most consumers searching for digital cameras have no interest in buying ringtones or pornography, ringtone merchants and pornographers may get higher *per-click* profits than camera stores because profit margins are high in ringtones and pornography, and the few people who click on the ads are reasonably likely to purchase. Such an outcome would be highly inefficient: consumer surplus is very low, and neither the search engine nor the advertising merchants make much money because the number of clicks is so low.

The simplest rule of thumb for choosing weights that is often mentioned in the literature is that they can be set equal to the predicted click-through rate of the advertisement.²⁶ The rough motivation for this is straightforward: weighting bids by their click-through rates is akin to ranking them on their contributions to search engine revenues.²⁷ 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. We note that the argument for click-through weighting is not straightforward in our environment and identify several considerations that firms will want to take into account in designing a weighting scheme.

VI.A. A Model of Click-Weighting

To create a model in which click-weighting is natural, we modify our model to allow each firm to have a two-dimensional type (δ, q) . A firm of type (δ, q) is able to meet the needs of a fraction δq of consumers. Consumers get some information about whether a firm can meet their need at zero cost by looking at the ad. A fraction $1 - \delta$ immediately learn that the advertiser cannot meet their need. The remaining δ fraction learn that the firm might meet their need. If they click, the firm will meet their need with probability q . As before, consumers incur a cost of s if they click on an ad to learn whether it in fact meets their need. Whether each firm can potentially help a consumer is independent across

26. See BCPP (2007) and Jerzierski and Segal (2009) for empirical evidence on the magnitude of differences in click-through rates and many other insights.

27. In practice, weights are the output of highly complex algorithms that may incorporate many other factors in addition to predicted click through rates. Search engines invest a great deal in trying to improve these algorithms, which are seen as a source of competitive advantage.

firms. We assume that the δ of each ad is known to the search site and to consumers. As before each site's q is private information.

One example that readers could keep in mind is a consumer who has searched for "shoes." The text in some sponsored links will reveal that the store in question serves only women or sells only athletic shoes. This immediately tells some consumers that the link would not meet their need. A store that potentially serves all consumers would have a large δ . A store that serves a small niche, for example, ballroomdancingshoes.com, would have a small δ .

In this environment a click-weighted auction could use the δ 's as the weights. Firms submit per-click bids b_1, \dots, b_N .²⁸ 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, $\frac{\delta^{k+1} b^{k+1}}{\delta_i}$.²⁹

VI.B. Inefficiencies of the Click-Weighted Auction

In some models without explicit consumer search costs it is obvious that a click-weighted auction is efficient. Things are not so obvious in our model. We note in this section that the click-weighted auction is efficient in one limiting case, but otherwise there are at least two distinct sources of inefficiency.

Our efficiency result is that the outcome approximates the first best in the limit as search costs become negligible.

PROPOSITION 8. In the limit as $s \rightarrow 0$, social surplus of the click-weighted auction converges to the first-best.

The argument for this has two steps, but is straightforward. The first step is to note that the firms that win places on the sponsored-link list will be the M firms for which $\delta_i q_i$ is largest. To see this, note that firms must bid up to q_i if it is necessary to get on the sponsored-link list because the firms get a zero payoff if they do not make the list. The second step is to note what makes

28. 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.

29. Note that as in our earlier discussions of bids we use subscripts as indexes when the index is a firm identity and superscripts as indexes when the index is the rank of the firm in the bidding, for example, δ_1 is the click-through weight of firm 1 and δ^3 is the click-through weight of the firm that is the third-to-last to drop out (in the weighted bidding).

consumers best off in the limiting case is to have the highest probability of meeting their need. 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.

The efficiency result is, however, limited by two other observations:

- The set of firms on the sponsored-link list is not necessarily optimal away from the $s \rightarrow 0$ limiting case.
- Even in the $s \rightarrow 0$ limit the ordering of the sponsored-link list may not provide consumers with as much information about advertiser quality.

In more detail, the first observation is that social welfare will sometimes be improved if some firms are replaced by other firms with higher q 's. This can be an improvement if the reduction in the number of needs met is more than outweighed by a saving of search costs. The inefficiency is most stark when M is large and the set of advertisers consists of both high-quality specialist firms, for example, $(\delta, q) = (\frac{1}{4}, 1)$, and low-quality generalist firms, for example, $(\delta, q) = (1, \frac{1}{3})$ with a slightly higher probability of meeting any given consumer's need. The click-weighted auction will produce list of low-quality generalist firms. But it would have been more efficient to have given consumers a list of high-quality specialists. Either list lets most consumers find a firm that meets their need, and there are many fewer wasted clicks when consumers get the list of high-quality specialists. One practical application is that click-weighted auctions may allow firms like eBay and Nextag to win more sponsored-link slots than would be socially optimal.

The second observation relates to the fact that the unweighted position auction always has an equilibrium in which firms end up sorted in order of quality. This ordering conveys valuable information to consumers. In the click-weighted auction, such equilibria generally will not exist. Intuitively, consumers will infer that a low δ firm that has made it into the top M firms must have a fairly high q ; consumers will therefore sometimes click on such a firm first even if it is not at the top of the list; this makes it hard to get low δ firms to bid for position. Appendix C presents some formal results. One notes that equilibria in which no additional quality information is revealed always exist and are

more robust in this model. Another notes that an equilibrium with quality sorting is possible in one special case.

VI.C. A New Auction Design: Two-Stage Auctions and Efficient Sorting

To eliminate the information 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 with a second-stage auction allowing bidders to raise bids further, but using a different payment scheme so that the equilibrium will have the firm with the highest q winning. Appendix C presents a formal model illustrating how this could be done. The mechanism, however, is more complex than a standard position auction and informational requirements could make it difficult to implement in practice.

VI.D. Obfuscation

A related issue of substantial practical importance is whether the snippet of text that accompanies each sponsored link conveys useful information about the link. Consumers benefit when the ad text lets them avoid unproductive clicks. And the search engine and advertisers benefit when consumers are more willing to click on seemingly relevant ads. In this section we discuss advertisers' incentives to make ad text accurate and informative. We note that click-weighted auctions can create incentives for obfuscation.³⁰

Our formal analysis is presented in Appendix C. The model augments our base model in two ways. First, we assume firms receive some small benefit a from each click regardless of whether the firm meets the consumer's need. This could reflect advertising revenues, future sales, sales unrelated to the need, or other factors. Second, we model the endogenous choice of an obfuscation strategy as an ability to choose the fraction $\delta_i \in [q_i, 1]$ of consumers who will think after reading the ad text that firm i might meet their need. We assume that firm i cannot affect the fraction q_i of consumers whose needs it will meet, so choosing a larger δ_i is implicitly choosing a lower probability of meeting the consumer's need conditional on the ad text being consistent with the need.

30. See [Ellison and Ellison \(2004, 2009\)](#) for a discussion of obfuscation including a number of examples involving e-retailers, and [Wilson \(2008\)](#) and [Ellison and Wolitzky \(2009\)](#) for models in which firms intentionally make search more costly.

Consumers cannot observe obfuscation choices: they believe that each firm is using the equilibrium level of obfuscation.

Our analysis of this model brings out two results:

- There would be no obfuscation in a position auction without click-weighting.
- Firms will engage in obfuscation in the click-weighted auction.

The first result should be intuitive—firms do not want unproductive clicks if they have to pay for them. The easiest way to see this mathematically is to condition on the firm's position k on the list and normalize the mass of consumers who will click on the firm's link if it is consistent with their need to 1. The firm's revenue from product sales is always q_i . There is an additional benefit $\delta_i a$ from the extra revenue source. And the firm pays $\delta_i b^{k+1}$ for the clicks it gets. Total profits are $q_i + \delta_i(a - b^{k+1})$. This is decreasing in δ_i whenever $b^{k+1} > a$, which will always be true when the firms who don't make the list would have gotten at least a from each consumer.

The second result is more striking. The key observation is that the formula for per-click payments in the click-weighted auction, $\frac{b^{k+1}\delta^{k+1}}{\delta_i}$, makes firm i 's total payment completely independent of its click-through rate—the total payment is simply equal to $b^{k+1}\delta^{k+1}$.³¹ The gross benefit firm i receives from the clicks it receives, $q_i + \delta_i a$, is increasing in the number of clicks received. So the insensitivity of payments makes firms want to obfuscate and make the click-through rate as large as possible. A feature of this argument that may make it practically relevant is that it applies even if the benefit, a , from unproductive clicks is very small.

Search engines may attempt to combat obfuscation in various ways. One is to try to enforce rules forbidding misleading ad text by refusing to display ads flagged by a manual or automated review. Another is to adjust the pricing formula either so that firm i 's per-click payment decreases less than one-for-one with increases in the clickthrough rate, or so that the per-click payment is affected by relevance measures (which can be based on textual

31. Immorlica et al. (2006) can be seen as showing that this invariance is advantageous in other dimensions—it makes it possible to generate click-through weights in a way that is resistant to click fraud.

analyses or on the number of consumers who immediately return to the search page). Pay-per-action auctions are not a solution: the Appendix shows that they have the same obfuscation problem.

VI.E. Product Variety

In practice, consumers who type in a given keyword will have heterogeneous needs. For example, a consumer who types the keyword “shorts” may be interested in upscale women’s clothing, athletic shorts, or perhaps even short films. Intuitively it would be desirable to serve such a population by presenting a diverse set of ads. The model we have presented so far cannot capture this intuition—our assumption that the probability that each website meets a consumers need is independent rules out the possibility that particular pairs of sites, for example, two sellers of athletic shorts, are likely to meet the same needs. Appendix C works through a simple extension of the model in which there are different categories of advertisers. We show that an optimal weighted auction will give a higher weight to sites that increase product variety. Note that in this context the weight given to site i cannot be defined just as a function of site i ’s characteristics: contributions to product variety depend on the characteristics of the other advertisers who appear on the sponsored-link list.

VI.F. 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 VI.A, 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 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 $q^{k:2N}$ 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.

VII. CONCLUSIONS

In this article, 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 obfuscation—since advertisers' steady-state payments do not vary with their click-through rate (they are determined by the revenue bid of the bidder in the next-lower position), advertisers have no incentive to design ad text to help consumers avoid unnecessary clicks. The working paper version of this article, [Athey and Ellison \(2009\)](#), discusses additional auction design questions as well, including policies toward search-diverting sites and the use of minimum relevance thresholds when consumers are uncertain about the distribution of advertiser quality.

A more basic theme of our article is that sponsored link auctions create surplus by providing consumers with information about the quality of sponsored links. Sorting links on the basis of weighted bids is an effective mechanism for providing

such information. But once one thinks about search engines as intermediaries whose role is to make consumer search more efficient, it immediately becomes salient that there could be effective ways to perform this role with quite different designs. For example, search engines could try to convey finer information about quality. They could try to convey raw data like bids, conversion rates, estimated textual relevance, or aggregates of these. Data could be conveyed numerically or via visual schemes varying the placement, size, or color of ads. Landing page previews could be added to help consumers assess the relevance of links, and so on. The scope for creative exploration is enormous and should make this an interesting area for pure and applied research for many years to come.³²

APPENDIX A: UNIFORMLY DISTRIBUTED QUALITY

This appendix presents some additional results on our model for the case when the the distribution F of the advertisers' qualities is uniform on $[0, 1]$. This special case is surprisingly tractable. This enables us to derive more explicit versions of several propositions.

A. Consumer Search

The tractability of the model with uniform F stems from the fact that there is a simple closed-form expression for consumers' expectations of the quality of the k th link conditional on not having found that the first $k - 1$ links do not meet their need.

PROPOSITION A1. For uniform F , if consumers search an ordered list from the top down, then

$$E(q^{k:N} | z^1 = \dots = z^{k-1} = 0) = \frac{N + 1 - k}{N + k}$$

$$\text{Prob}\{z^1 = \dots = z^{k-1} = 0\} = \prod_{j=1}^{k-1} \frac{2j - 1}{N + j}$$

32. See [Rayo and Segal \(2010\)](#) for one interesting approach to the more abstract question of how consumers might be provided with information on advertiser quality.

Proof. When F is uniform, $f^{k:N}(x) = \frac{N!}{(N-k)!(k-1)!} (1-x)^{k-1} x^{N-k}$. The general formula in Proposition 1 gives

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

The conditional probability that shows up in the numerator and denominator is:

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

Hence, both the numerator and the denominator are some 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 $\frac{a!b!}{(a+b+1)!}$. Evaluating the integrals gives the first formula in the statement of the proposition.

The second formula in the proposition follows from computing the integral that is the dominator of the foregoing formula or it can be proved more quickly by noting that

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

■

Having an explicit expression for the conditional expectations makes it easy to give an explicit description of consumer behavior.

PROPOSITION A2. If the firms are sorted by quality in equilibrium and the distribution F of firm qualities is uniform, then a consumer with search cost s stops clicking when she reaches position $k^{max}(s)$, where

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

The proof of this result is immediate: the consumer will want to search the k th website if $\frac{(N+1-k)}{(N+k)} > s$. This holds for $k < k^{max}(s)$.

B. Consumer Welfare

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 $E(q^{1:N}) - s = \frac{N}{(N+1)-s}$. If the first link is unsuccessful, which happens with probability $\frac{1}{(N+1)}$, then (using Proposition A1) the consumer gets utility $E(q^{2:N} | z^1 = 0) - s = \frac{(N-1)}{(N+2)} - s$ from clicking on the second. Adding up these payoffs over the number of searches that will be done gives the following result:

PROPOSITION A3. If the distribution of firm quality F is uniform, the expected utility of a consumer with search cost s is:

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

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 $E(CS(s)) > \frac{5}{6} - s$ for all s . Figure A.1 plots the relationship between $E(CS)$ and s for $N = 4$.

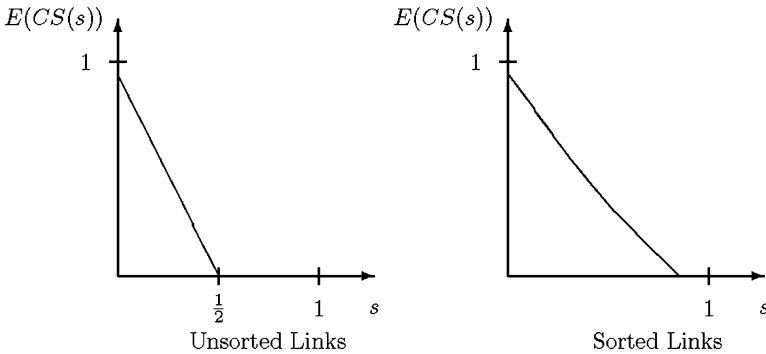


FIGURE A.1

Consumer Surplus with Sorted and Unsorted Links: $N = 4$

C. Click-Through Rates

If we add an assumption that consumer search costs are also uniformly distributed, then we also obtain simple explicit expressions for click-through rates in our model. Proposition A4 gives two explicit formulas: one for the expected click-through rate; and another for the conditional expected click-through rate when the k th link has quality q . The former might be compared with data on actual click-through rates. The is relevant to the firm’s bidding problem.

PROPOSITION A4. Assume s and $q \sim U[0, 1]$. Consider an equilibrium of the bidding game in which firms are sorted in order of quality. Then, the expected number of clicks $D(k)$ that will be received by the k th website is

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

The expected number of clicks conditional on the k th highest quality website having quality q is

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

Proof. Expected click-through rates can be computed as a product of two terms: the probability that a consumer’s search costs is such that he/she would be willing to click on the k th link if the he/she does not meet his or her need at the first $k - 1$ websites; and the probability that the first $k - 1$ clicks will all be unsuccessful. With uniformly distributed search costs the former is simply the expected quality of the k th link conditional on $k - 1$ failures. Hence, the unconditional expected clickthrough rate is simply the product of two expressions from Proposition A1.

The second expression is derived by thinking of the first $k - 1$ websites as ordered at random rather than ordered by decreasing quality. With this reordering, whether each of the $k - 1$ websites meets the consumer’s need is independent. Hence, the probability that none of the $k - 1$ sites meets the need is $(1 - E(q_j|q_j > q))^{k-1} = \left(\frac{1+q}{2}\right)^{k-1}$. The probability that a consumer’s search cost is such that she would click on the website is as in the previous formula. ■

D. Equilibrium Bidding

Plugging the conditional expectation formula into the formula for equilibrium bidding also yields a formula for equilibrium bidding which simplifies.

PROPOSITION A5. When both qualities and search costs are uniform it is a PBE for firms to choose dropout points according to

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

Proof. The formula is obtained by substituting the expression for $G(\bar{q}_k)$ from Proposition A1 into the general formula for the equilibrium in Proposition 3 and simplifying:

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

■

E. Optimal Reserve Prices with One-Position Lists

The analysis of the optimal reserve price with a one-position list (and uniformly distributed search costs) also becomes very tractable when the quality distribution is uniform.

PROPOSITION A6. Suppose that the list has one position, the distribution of search costs is uniform, and the distribution F of firm qualities is also uniform. Then the welfare maximizing reserve price r is the positive solution to $r + r^2 + \dots + r^N = \frac{N}{(N+2)}$.

Proof. Using the decomposition

$$\begin{aligned} E(q^{1:N}) &= \text{Prob}\{q^{1:N} < r\}E(q^{1:N} | q^{1:N} < r) + \text{Prob}\{q^{1:N} > r\} \\ &\quad E(q^{1:N} | q^{1:N} > r) \end{aligned}$$

we find that the conditional expectation of the order statistic is:

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

The formula follows from the general formula of Proposition 5 of the main text after a bit of algebra. ■

The formula implies that the welfare-maximizing reserve price is one-third when $N = 1$. With $r = \frac{1}{3}$, consumer expectations will be that $q \sim U[\frac{1}{3}, 1]$, and consumers search if and only if $s \in [0, \frac{2}{3}]$. Hence the average search cost is indeed $\frac{1}{3}$ as the general formula of Proposition 5 requires.

The expected consumer surplus, search engine profits, and social welfare are all higher with a small positive reserve price than with no reserve price when the quality distribution is uniform. The general expression for consumer surplus with a reserve price of r becomes

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

Writing $SR(r)$ for the search engine’s revenue when it uses a reserve price of r we find

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

Each of these expressions is increasing in r for small r .

APPENDIX B: RESERVE PRICES WITH NONUNIFORM SEARCH COSTS

The analysis of reserve prices in the text focuses on the case when the distribution G of consumer search costs is uniform on $[0, 1]$. In this appendix we present some results involving general search costs distributions and a figure illustrating the welfare trade-off in a few examples.

A. Consumer Optimal Reserve Prices

Our first result is that consumer surplus maximization does require a positive reserve price for any G . We prove this by showing that consumer surplus is increased when small positive reserve prices are implemented.

PROPOSITION A7. Consumer surplus is maximized at a strictly positive reserve price.

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 $r = 0$.³³ 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 . ■

B. Social Welfare

Proposition 5 of the main text implies that the socially optimal reserve price is always positive when G is uniform. Here, we note that social welfare need not be maximized at a positive reserve price for arbitrary G . An intuition for why reserve prices can be harmful is that 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.

PROPOSITION A8. 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

33. Formally, 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.

have an effect that works through changes in 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 $O(r)$. Appendix D 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 proposition 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 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.³⁴ There is a substantial welfare loss in this subpopulation if the search engine uses a large reserve 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. ■

The example used in the proof is obviously quite special. We do not mean to suggest that the adverse effects of reserve prices are likely to dominate the beneficial effects in the real world. Rather, the motivation for presenting the proposition was simply to note that one property of socially optimal reserve prices in the model with uniform search costs cannot be generalized to apply for all G

34. The 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 $2\text{Min}\{r, \epsilon\}\frac{\epsilon}{2}$.

and to illustrate a channel through which reserve prices can have a negative impact on social welfare.

C. Reserve Prices

It is also theoretically possible that using a small positive reserve price 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 $M-1$ st position. Because bids depend recursively on lower bids, this can reduce bids on higher positions as well.

PROPOSITION A9. There exist distributions F and G for which search engine revenue is decreasing in the reserve price in a neighborhood of $r = 0$.

Again, the proof consists of a very special 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 $E(q^{2:N})$.

The example above 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 omit the details of these modifications.

D. Numerical Examples of Welfare Trade-Offs

Figure A.2 illustrates the conflicting preferences of advertisers, consumers, and the search engine in two specifications of the model. In each panel, we graph expected advertiser surplus, expected consumer surplus, and expected search engine profit as a function of the reserve price and draw vertical lines at the values of r that maximize each of these functions.

The right panel is for a model with three firms drawn from a uniform quality distribution. In this model, consumer surplus turns out to be fairly flat over a wide range of reserve prices. The advertiser-optimal and search engine profit-maximizing prices are quite far apart, but consumer surplus at both of these points is not very far from its optimum. An intuition for the flatness of the consumer surplus function is that reserve prices in this range are rarely binding, and hence there is little direct effect on the probability of a link being displayed and little indirect effect via changes in consumer beliefs. The main effect of a shift from consumer-optimal to profit-maximizing reserve

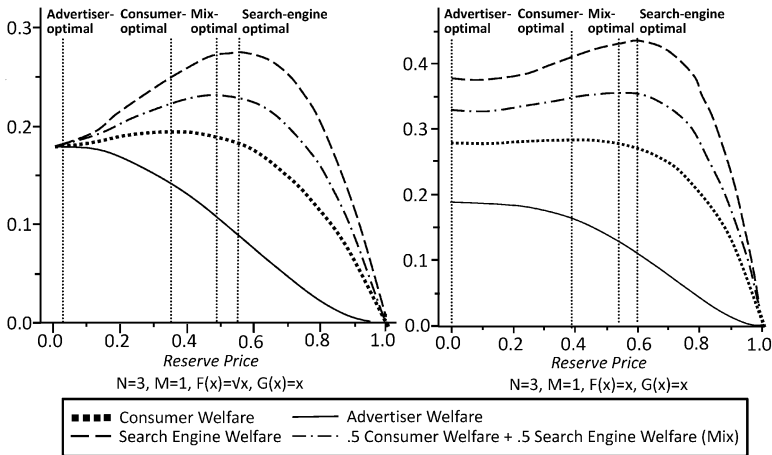


FIGURE A.2

Welfare and Distribution of Surplus for Two Specifications

prices is a redistribution in surplus from advertisers to the search engine.

The left panel is for a model with three firms with qualities drawn from the CDF $F(q) = \sqrt{q}$. This distribution is more concentrated on low-quality realizations, which makes consumer surplus and advertiser profits more sensitive to the reserve price in the relevant range. In each panel we also graph an equally weighted average of consumer surplus and search engine profit. The curvature of the functions involved is such that the maximizer of this average is closer to the profit-maximizing level than to the socially optimal level.

APPENDIX C: CLICK-WEIGHTED AUCTIONS

In this appendix we present some additional formal results on click-weighted auctions.

A. *Inefficiency in the Ordering of Listed Firms*

One result highlighting the potential loss of efficiency about advertiser qualities is very simple:

PROPOSITION A10. The click-weighted auction always has an equilibrium in which all remaining firms drop out immediately as soon as just M firms remain.

The equilibrium strategies are the obvious ones: firms remain in bidding until the bid reaches q if there are more than M firms remaining and then drop out immediately once M firms remain. When firms follow these strategies, consumers' beliefs about the quality of each remaining firm i conditioning on all available information X is $E(q_i|X) = E_{b^{M+1}}(q_i|\delta_i q_i > \delta^{M+1} b^{M+1})$. This is higher for firms with a lower δ , so it is an equilibrium for consumers to ignore the ordering of the firms on the list and search in increasing order of δ . With this consumer behavior, there is no benefit to bidding for a higher position and immediate dropout is optimal.

Immediate dropout equilibria also existed in the unweighted auction model—all consumers can search in a random order—but we mention them here because they seem more natural and robust when consumers have reason to believe that some advertisers are better than others and strictly prefer to search in the way that they do. One robustness criterion that would distinguish

the click-weighted model from the unweighted model is that behavior would not change substantially in the click-weighted model if consumers thought that there was an ϵ probability that the ordering was informative and a $1 - \epsilon$ probability that the ordering was not.³⁵

A second formal result illustrates that greater information revelation is also possible: the click-weighted auction model does have an equilibrium with full sorting in one special case. To define this, let \bar{s} be such that all consumers with search costs $s < \bar{s}$ will search all listed websites as long as their need has not been met.³⁶

PROPOSITION A11. Suppose that $N = M = 2$ and the support of the search cost distribution G is a subset of $[0, \bar{s}]$. 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 (conditional 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. ■

The full-sorting example uses several special assumptions. These are largely necessary to get full sorting. For example, one can show that there are no equilibria with full sorting when $\delta_1 \neq \delta_2$ if one assumes instead that $N > 2$ and/or that G has full support on $[0, 1]$.³⁷ A rough intuition for this is that the solution to the asymmetric first-order condition will not satisfy the

35. Chen and He (2006) note that the immediate-dropout equilibria in the unweighted model are nonrobust to assuming that an ϵ fraction of consumers always search in a top-down manner. This remains true in the click-weighted auction.

36. An $\bar{s} > 0$ with this property will exist if the δ 's are bounded away from 0. For example, it suffices to set $\bar{s} = E(q^M | \delta^M = 1, \delta^1 = \dots = \delta^{M-1} = \underline{\delta}, z^1 = z^2 = \dots = z^{M-1} = 0)$.

37. We thank Dmitry Taubinsky for these results.

symmetry condition necessary for full sorting, $\delta_i b_i^*(q) = \delta_j b_j^*(q)$, except in particular special cases. Accordingly, we feel that the more important lesson from this section is the first one: click-weighted auctions do not have the the nice information-revelation feature of the unweighted position auction.

B. A New Auction Design: Two-Stage Auctions and Efficient Sorting

Suppose $M = 2$ and $N > 2$ and firms 1 and 2 are the remaining firms. To design an auction in which there will be an equilibrium in which the firms are sorted by quality, it will suffice to choose asymmetric payment schedules $p_1(q)$ and $p_2(q)$ such that it will be an equilibrium for firms to announce their true qualities q_1 and q_2 if we ask firms to announce their qualities, put the firm with the highest announced quality in the first position, and assign a per-click payment of $p_i(q_{-i})$ to the winning firm i and $\frac{b^3 \delta^3}{\delta_j}$ to the losing firm j .

To see that such payment schedules exist, note that if a firm is the last to drop out, its expected profit is

$$\mu_{1i} \delta_i (q_i - p_i(q_{-i})),$$

where $\mu_{1i} = G(E(q_i | q_i > q_{-i}, X))$ with X the event that firms 1 and 2 are the two winning bidders. If the firm is the second to drop out, its expected profit is

$$(\mu_{2i} (1 - \delta_j) + \mu_{3i} \delta_j (1 - q_{-i})) \delta_i \left(q_i - \frac{b^3 \delta^3}{\delta_i} \right),$$

where $\mu_{2i} = G(E(q_i | q_i < q_{-i}, X))$ and $\mu_{3i} = G(E(q_i | q_i < q_{-i}, X, z_{-i} = 0))$. It is straightforward to choose $p_i(q)$ so that these two expressions are equal conditional on $q_{-i} = q_i$, in which case the necessary indifference condition for a truth-telling equilibrium is satisfied.³⁸ Note that to implement such rules the search engine needs to know the δ 's and also needs to know what click-through rates each firm will receive given each possible ordering. Knowing the δ 's is necessary for everything we've done in this section. The

38. One can show that the $p_i(q)$ function defined in this way is monotone. Indifference need not hold when the high δ firm drops out at a point when the other firm is known to have higher quality. To complete the specification without interfering with the selection of the final two firms we set $p_i(q) = \frac{b^3 \delta^3}{\delta_i}$ if $\delta_i < \delta_{-i}$ and $q < \frac{b^3 \delta^3}{\delta_i}$.

additional informational requirements will be more of an obstacle to implementing such schemes in practice.

When $s < \bar{s}$ for all consumers the payment functions take a particularly simple form:

$$p_i(q) = \frac{b^3 \delta^3}{\delta_i} + \delta_{-i} q \max \left\{ \left(q - \frac{b^3 \delta^3}{\delta_i} \right), 0 \right\}.$$

Using this formula we can see that firm 2 is favored at low quality levels when $\delta_1 < \delta_2$ in the sense that it makes a lower payment when the firms have equal qualities and these qualities are near the lowest possible. At high quality levels the bid preference may be reversed.

C. Obfuscation

We augment our base model in two ways. First, we assume that each firm i receives some benefit α from each click it receives independent of whether it meets the consumer's need. Second, we assume that each firm chooses 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 the 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 $\alpha = 0$ and no option other than full obfuscation, $\Lambda = \{1\}$.

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 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.³⁹ One thing that simplifies our analysis is that γ_k does not depend on the actual obfuscation level of the firm in position k .

39. Note 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.

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

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

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

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

In equilibrium b^{M+1} will be at least $a + q^{M+1:N}$, so this is negative and no obfuscation occurs in equilibrium.

If there was heterogeneity in the benefits a_i that firms receive from clicks that do not meet consumers' needs, then it is possible that firms with large a_i could engage in obfuscation. But note that it would still be necessary for a_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 .⁴⁰ Conditional on being in the k th position ($k \leq M$), firm i 's payoff is

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

This expression is monotone increasing in δ_i . Hence, in equilibrium we get full obfuscation: all firms choose $\lambda_i = 1$.

Search engines have been developing the capability to track sales made by their advertisers. This enables pay-per-action auctions: firms submit bids b_i which represent payments to be made to the search engine only if a consumer clicks on their link *and has their need met*. Suppose that a search engine records the fraction of clicks which result in needs being met, y_i , and uses this as an additional weighting factor just as click-through rates are used in the click-weighted auctions: the search engine ranks the

40. Note 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.

firms on the basis of $\delta_i y_i b_i$ 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

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

This expression is virtually identical to the expression for the standard click-weighted auction. The result on obfuscation carries over.

D. Product Variety

We consider here the simplest extension of our model with different categories of advertisers. There are three sites: site 1A, site 1B, and site 2. Suppose that a fraction $\delta_1 > \frac{1}{2}$ 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 $[0, 1]$. To simplify the analysis we suppose that all 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.

Consider a weighted $k+1$ st price ascending bid auction in which winning bidders are chosen by comparing b_{1A} , b_{1B} , and $w b_2$. As before, assume that the per-click payment of firm k is the $k+1$ st highest bid adjusted for the weight difference (if a difference exists). 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, there will again be an equilibrium with full sorting if firms 1A and 1B are the two remaining firms. 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 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, that is, $q_{1B} = wq_2$. 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 $\delta_1(1 - q_{1A})E(q_{1B}|q_{1B} < q_{1A}, q_{1B} = wq_2) = \frac{\delta_1(1-q_{1A})q_{1A}}{2}$.⁴¹ Including firm 2 can provide incremental utility to any type 2 buyer: the incremental benefit is $(1 - \delta_1)E(q_2|q_{1B} < q_{1A}, q_{1B} = wq_2) = \frac{(1-\delta_1)q_{1A}}{2w}$. 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 $\frac{2}{3}$ and the conditional expectation of $q_{1A}^2 = \frac{1}{2}$. Hence, there is a gain in consumer surplus from choosing $w > 1$ if $\delta_1(\frac{1}{3} - \frac{1}{4}) < (1 - \delta_1)\frac{1}{3}$. We have

PROPOSITION A12. The consumer surplus-maximizing weighted auction is one that favors diversity of the listings ($w > 1$) if $\delta_1 < \frac{4}{5}$.

Proof. 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:

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

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

The expected qualities are

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

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

The first two are again identical to the formulas for the unweighted case because this L only arises when $q_2 \in [0, \frac{1}{w}]$ 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 $q_2 \in [0, \frac{1}{w}]$ and values with $q_2 \in [\frac{1}{w}, 1]$. For example,

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

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

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

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

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

Writing $f_1(w)$, $f_2(w)$, $g_3(w)$, and $h_3(w)$ for the four lines of this expression note that the first three terms are equal to 0 at $w = 1$. $f_2(w)$ is identically 0. The derivative of the third evaluated at $w = 1$ is just $\frac{dg_3}{dw}|_{w=1}h_3(1)$. 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 $\alpha < \frac{4}{5}$. 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. ■

REMARK 1. The proof contains an explicit formula for consumer surplus that could be maximized over w to find the optimal weight for particular values of δ_1 .

REMARK 2. 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 is positively or negatively correlated with whether that consumer clicked on another 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, 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 $\delta_1 > \frac{1}{2}$. CTRs could also be estimated using an average of observed CTRs 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 CTRs when firms are in the second position.

APPENDIX D: PROOFS

Proof of Proposition 3. 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 ($k = 2$). Then we 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 $b^3 = b^*(3, b^4, q)$ so that firm i 's belief is that the other active firm has $q_j \sim F|_{q > q}$. Firm i 's expected payoff as a function of its dropout point \hat{q} can be written as $\frac{1}{1-F(\hat{q})} \pi(q_i, \hat{q})$ where

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

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

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

To show that this is non-negative, 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 $b^{k+1} = b^*(k + 1, b^{k+2}; \underline{q})$. 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 -1st highest quality, player i will be the first to drop out in stage $k - 1$ if he then follows the equilibrium strategy. Hence, the change in payoff is proportional to

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

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 before, this will be non-negative 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 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 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

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

is non-negative 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 $\hat{q} \leq q_i$ case. The expression is non-negative 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. ■

Additional Details on the Proof of Appendix Proposition A8. 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

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

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

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

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

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

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, 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. ■

REFERENCES

- Aggarwal, Gagan, John Feldman, S. Muthukrishnan, and Martin Pal, "Sponsored Search Auctions with Markovian Users," Workshop on Internet and Network Economics, 2008.
- Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani, "Truthful Auctions for Pricing Search Keywords," ACM Conference on Electronic Commerce, 2006.
- Athey, Susan, and Glenn Ellison, "Position Auctions with Consumer Search," National Bureau of Economic Research Working Paper 15253, 2009.
- Ausubel, Lawrence M. and Peter Cramton, "Auctioning Many Divisible Goods," *Journal of the European Economic Association*, 2 (2004), 480–493.
- Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek, "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence," mimeo, 2006.
- Chen, Jianqing, De Liu, and Andrew B. Whinston, "Auctioning Keywords in Online Search," *Journal of Marketing*, 73 (2009), 125–141.
- Chen, Yongmin, and Chuan He, "Paid Placement: Advertising and Search on the Internet," mimeo, 2006.
- Cramton, Peter, Yoav Shoham, and Richard Steinberg, "An Overview of Combinatorial Auctions," *ACM SIGecom Exchanges*, 71 (2007), 1–12.
- Demange, Gabrielle, David Gale, and Marilda Sotomayor, "Multi-Item Auctions," *Journal of Political Economy*, 94 (1986), 863–872.
- Diamond, Peter, "A Model of Price Adjustment," *Journal of Economic Theory*, 3 (1971), 156–168.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97 (2007), 242–259.
- Edelman, Benjamin, and Michael Schwarz, "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100 (2010), 597–602.
- Eisenmann, Thomas R. and Kerry Hermann, "Google Inc.," Harvard Business School Case 806–105, 2006.
- Ellison, Glenn, and Sara Fisher Ellison, "Search, Obfuscation, and Price Elasticities on the Internet," National Bureau of Economic Research Working Paper 10570, 2004.
- Ellison, Glenn, and Sara Fisher Ellison, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77 (2009), 427–452.
- Ellison, Glenn, Drew Fudenberg, and Markus Möbius, "Competing Auctions," *Journal of the European Economic Association*, 2 (2004), 30–66.
- Ellison, Glenn, and Alexander Wolitzky, "A Search Cost Model of Obfuscation," mimeo, Massachusetts Institute of Technology, 2009.
- Immorlica, Nicole, Kamal Jain, Mohammad Mahdian, and Kunal Talwar, "Click Fraud Resistant Methods for Learning Click-Through Rates," mimeo, Microsoft Research, 2006.
- Jerzierski, Przemyslaw, and Ilya Segal, "What Makes them Click: Empirical Analysis of Consumer Demand for Search Advertising," mimeo, Stanford University, 2009.
- Kameinica, Emir, "Contextual Inference in Markets: On the Informational Content of Product Lines," *American Economic Review*, 98 (2008), 2127–2149.
- Lehaie, Sebastien, "An Analysis of Alternative Slot Auction Designs for Sponsored Search," EC '06, 2006.
- Liu, De, and Jianqing Chen, "Designing Online Auctions with Past Performance Information," *Decision Support Systems*, 42 (2006), 1307–1320.
- Liu, De, Jianqing Chen, and Andrew B. Whinston, "Ex-Ante Information and the Design of Keyword Auctions," *Information Systems Research*, 21 (2010), 133–153.
- Rayo, Luis and Ilya Segal, "Optimal Information Disclosure," *Journal of Political Economy*, 118 (2010), 949–987.
- Varian, Hal R., "Position Auctions," *International Journal of Industrial Organization*, 25 (2007), 1163–1178.
- White, Alexander, "Search Engines: Left Side Quality versus Right Side Profits," mimeo, Toulouse School of Economics, 2008.

- Wilson, Chris M., "Ordered Search and Equilibrium Obfuscation," mimeo, University of Oxford, 2008.
- Xu, Lizhen, Jianqing Chen, and Andrew B. Whinston, "To Place Better or Price Cheaper? Bidding and Pricing under Keyword Advertising," mimeo, University of Texas, 2008.
- , "Too Organic for Organic Listing? Interplay between Organic and Sponsored Listing in Search Advertising," mimeo, University of Texas, 2009.