Ad Server and Firm Strategies in Contextual Advertising Auctions* (Job Market Paper)

Charlie Gibbons†
University of California, Berkeley

November 11, 2011

Preliminary draft—comments are appreciated.

Abstract

We consider the strategies of online advertising providers, firms, and consumers in the context of ad listings assigned by a generalized second price auction. The first part of the paper develops a model of consumer responses to ad listings and product offerings from the included firms and uses this behavioral model to derive optimal bidding functions for the firms. We show that the relationship between per-sale margins and product-consumer match probabilities ("relevances") must meet certain conditions to rationalize this equilibrium for consumers and firms; in particular, we give the conditions for consumers to rationally search from the top of the listing downward. Next, we turn to the incentives facing the ad server to alter the relevances and margins of the firms and the search costs and valuations of the consumer pool. While these incentives align with the desires of consumers, they may conflict with those for firms. We calculate the optimal number of slots for the ad server to offer, which is less than that desired by firms and consumers. We also show that the ad server has an incentive to subsidize its own competitor in the product market. These results have important implications for competition policy, innovation, and online content provision.

*I would like to thank Henry Brady, Jeff Deason, Aaron Edlin, Joe Farrell, Benjamin Hermalin, Ben Handel, Michael Katz, Rob MacCoun, Suzanne Scotchmer, Jas Sekhon, Jidong Zhou, and seminar participants at UC, Berkeley, the Searle Center for Law, Regulation, and Economic Growth's conference on Internet Search and Innovation, and the U.S. Federal Trade Commission's Microeconomics Conference for their suggestions on this project.

†cgibbons@econ.berkeley.edu

1 Introduction

Advertising is essential in funding online content, from social networking sites to newspaper articles to streaming music and search engines. On all these sites, there is a movement toward contextual ads that are related to keywords found on the page. These ads aim to generate immediate action by consumers, including clicking a link and performing an "action," such as purchasing a product from the advertiser's site. By leading to immediate sales, the listings of these ads can impact the market for the advertised product.

The provider of these ads, known as the *ad server*, may have an incentive to reduce the competitiveness of the product market in order to extract higher advertising revenues from firms that pay a premium to be listed in a less competitive space. Alternatively, the ad server could be encouraged to find ways innovate and increase sales in a competitive market to raise its ad revenue. A close examination of these issues is necessary as interest turns toward competition policy for the online advertising market.

In this paper, we focus on the incentives facing an ad server to manipulate the primitives of the market in order to generate higher advertising revenue. In particular, we consider the incentives to improve consumer-product match probabilities, reduce search costs, boost firm product margins, and target profitable consumer groups. We also examine how dispersion in the margins and match probabilities among firms impact ad revenues and how the ad server may limit the length of the ad listing to combat dispersion. Lastly, we consider an ad server that is also active in the market for the product being advertised (*e.g.*, Google providing ads for e-mail services, including its own Gmail offering) to determine whether the ad server can increase total profits from combined product and ad sales by artificially placing its own firm at the top of the listing. We determine whether all these incentives align with those of firms and consumers.

The incentives of all parties agree to reduce search costs. The ad server has an incentive to fulfill consumers' wishes for better product matches, while this innovation exceeds the level desired by firms. The incentives of the ad server may be contrary to those of consumers, however. We show that the ad server displays a limited number of ads, while more is always better for consumers and for firms in total. Additionally, the ad server has an incentive to subsidize its own competitor in the product market, which changes the sizes and distributions of producer and consumer surpluses.

This latter point has important implications for antitrust inquiries into Google in particular.

In answering these questions, we begin by formulating a model of consumer responses to contextual advertising. Consumers need to be matched to a product that they like and only a fraction of consumers like a product offered by a particular firm. For example, a consumer may be searching for a sweater, but may want a different style than the particular one offered by an advertiser. A fraction of consumers are satisfied by a firm's offering and have a positive valuation for that particular product, while the rest do not like it, giving them 0 valuation. The probability that a consumer likes the product offered by a firm (*i.e.*, has positive valuation for it) is called the firm's *relevance*.

For a consumer to make a purchase, the product must not only be relevant, but also its price must be below the consumer's valuation. We first consider a model in which a common market price prevails across all firms. Each firm knows its relevance as well as its cost. Based upon these factors, the firms bid for placement in the advertising list via a generalized second price auction. We incorporate our model of consumer behavior into the Varian (2007) framework to consider how heterogeneous margins and relevances of the firms impact bidding strategies.

Under the special case that all the firms have the same cost, firms sort by relevance, as in the model of Athey and Ellison (2009). Since better matches are placed at the top, consumers are most likely to find a relevant product quickly and consumer surplus is maximized, as is producer surplus. Alternatively, if firms all have the same relevance, they sort in decreasing order of cost, another ordering that maximizes total surplus. If both costs and relevance vary across firms, the ordering depends upon how these characteristics covary and we derive the maximum covariance between these factors that can exist in an equilibrium with consumers that search from the top down.

The ranking of firms, then, is endogenous to the model. We require this feature in order to understand the incentives that the ad server faces to improve or alter these rankings. Our model builds upon other models of ordered consumer search that assume exogenous rankings (see Arbatskaya, 2007; Armstrong, Vickers and Zhou, 2009; Xu, Chen and Whinston, 2011*a*). The work of Xu, Chen and Whinston (2011*b*) endogenizes firm rank, but the model of consumer behavior here is most similar to the work of Chen and He (2006) and Athey and Ellison (2009), models that also endogenize firm rank and incorporate the idea of relevance. All but the latter of these papers focus on firm pricing strategies that depend upon position; Athey and Ellison (2009) focus on optimal auction design in this context. Our paper instead focuses on the incentives for the ad server to innovate and to change the structure of the market, as discussed above.

Unlike these papers, we allow consumers to have different valuations for the product, conditional upon finding it relevant. This allows us to consider possible selection effects that impact the profitability of the pool of consumers that visit each site. For example, consumers with high valuations may quickly find products priced below their valuation, allowing them to make a purchase and to leave the market, while low valuation consumers continue onward. Alternatively, high valuation consumers have the most to gain from finding a relevant product and may continue searching longer than low-valuation consumers. These patterns may influence the profitability of being in a particular slot in the ad listing, an issue that has not been studied in the literature.

This examination of consumer behavior and bidding strategies serves as a foundation for future applications. Very little is known about competition in ad serving and the ability of contextual advertising revenue to adequately fund online content. Innovation is crucial in this relatively young industry and we must study the returns to innovation accruing to ad servers to determine the adequacy of technological advancement in this area. Additionally, if the ad server has the ability and incentive to influence competition in the product markets themselves, this can raise important antitrust concerns. The model and results of this paper can serve as a useful framework to examine issues in innovation, antitrust, and business strategy.

2 Model

The first step in analyzing the optimal bidding strategy of firms and the resulting incentives facing ad servers is to formulate a model of consumer responses to ad listings. We develop such a model in this section.

2.1 Framework

We begin by creating a model of consumer behavior in perusing the advertising listing and in purchasing a product offered by one of the firms. There is a unit mass of consumers, indexed by $i \in [0, 1]$. The firms are indexed $j \in \{1, \ldots, J\}$. Consumers view a listing of M advertisements and consider each product being advertised sequentially starting at the top of the list.[1]

Consumers have a sort of lexicographic preferences that are firm-specific. A product is either relevant for consumer i, yielding a positive valuation for that product v_i, distributed with cumulative distribution function F, or it fails to meet his needs and the consumer has no value for it at all. The needs of each consumer are met stochastically with probability q_j by firm j. That is, the valuation of the product offered by firm j to consumer i is

$$v_{ij} = \begin{cases} v_i & \text{with probability } q_j \text{ and} \\ 0 & \text{with probability } 1 - q_j. \end{cases}$$

We refer to the probability of firm j offering a relevant product to a consumer as the *relevance* of firm j q_j. While the relevance varies across firms, it is the same for all consumers facing a given firm. This model of preferences differentiates the firms and provides a rationale for multiple firms to each to receive a positive market share. Product differentiation can also be helpful for avoiding the Diamond (1971) and Bertrand paradoxes in which firms respond to consumer search with perfectly monopolistic or competitive pricing strategies, thereby eliminating the need for consumers to search at all (for a similar justification for product differentiation in search models, see Anderson and Renault, 1999).

A consumer searches by visiting site j, which is placed in slot m of the ad listing, then determines whether that product is relevant to him. If so, he compares the price of the product p_j to his valuation. If the price is below his valuation, he makes the purchase and the search ends. If the product's price exceeds his valuation or it does not meet his needs, he continues searching with probability s_m.[2] Note that s_m is a conditional probability: given that a consumer visited site m and did not make a purchase, s_m gives the probability that he continues searching. For completeness, define s_0 as the probability that a consumer visiting a site displaying the ads looks at the ad listing at all. These search probabilities are exogenous and independent of the valuation held by the consumer.

Let us summarize the quantities that define the structure of the model. First, we have the

[1] In a later section, we consider whether this search pattern is an equilibrium response to firm decisions.

[2] In this model, consumers do not search for the best deal; instead, they simply find a product that meets their needs at a sufficiently low price and make the purchase. If they do not find such a product, some fraction continue searching. Unlike Athey and Ellison (2009), the probability of searching forward is exogenous in this model.

consumer-specific valuations v_i for consumer i. Next, we have firm-specific costs c_j, prices p_j, and relevances to consumers q_j for firm j. Lastly, we have slot-specific search continuation probabilities s_m for slot m. All these quantities are exogenous. In the case of prices, we begin by considering the case in which all firms charge the same price p. Here, prices are set outside the model and there is a single price that prevails in the market.[3]

Distinguishing between firms and slots at first appears cumbersome. Suppose that the index of firm $j \in \{1, \ldots, J\}$ corresponds to the rank of $q_j(p - c_j)$ in decreasing order; that is, firms are indexed in decreasing order of their full expected margins.[4] We show in Section 3.1 that, in the equilibrium of interest in this paper, there is a direct mapping from this ordering to the slots in which the firms appear: the first M firms according to this ranking appear on the list in the slots that corresponds to their ranks, while the remaining firms do not appear on the list. Then, rather than cite "firm j in slot m," we can simply refer to firm j, as the slot m must equal rank j in equilibrium.

2.2 Product market outcomes

An important quantity of interest in online advertising is the *click-through rate* (CTR) for the ad in slot j, defined as the probability of a consumer clicking on slot j's ad. Recall that firm j appears in slot j in equilibrium. The CTR r_j for firm j is the probability that the consumer enters the list, does not purchase from firms 1 to $j - 1$, and continues on to site j.

In the case of the first site in the ad listing, the CTR is simply the proportion of the unit mass of consumers that search the list at all:

$$r_1 = s_0.$$

Next, consider the firm in the second slot. A consumer arrives at this site because either

[3]One way to understand this assumption is that firms charge the same price to consumers that visit their sites via the advertising list, to consumers that visit the sites directly without any search aids, and perhaps even to consumers that find the firms' products in brick-and-mortar retailers. Other authors have endogenized price into models of search, assuming that firms price discriminate based upon the particular ordering of a particular listing (see Chen and He, 2006; Arbatskaya, 2007; Armstrong, Vickers and Zhou, 2009; Xu, Chen and Whinston, 2011*a*,*b*). These works do not discuss whether such discrimination can be maintained; for example, consumers could find a relevant product via searching, but then visit the site directly by typing in its address manually to find a better deal.

[4]We assume that there are no ties.

- The product of firm 1 did not meet his needs or
- Though the product of firm 1 did meet his needs, it was too expensive (*i.e.*, the market price is above his valuation)

and he decided to continue searching. Since prices are the same across firms, the consumers in the group described by the second bullet above *never makes a purchase.*[5] The CTR is

$$r_2 = s_0 s_1 (1 - q_1) + s_0 s_1 q_1 F(p) = s_0 s_1 \left[(1 - q_1) + F(p) q_1\right].$$

We can generalize this expression to site $j \leq M$:

$$r_j = \prod_{k=0}^{j-1} s_k \left[\prod_{k=1}^{j-1} (1 - q_k) + F(p) \left[1 - \prod_{k=1}^{j-1} (1 - q_k) \right] \right]. \tag{1}$$

As there are only M sites listed, the CTR for firms $M+1, \ldots, J$ is 0. The two sets of terms inside the brackets above describe two groups of consumers: those that have yet to find a relevant product when reaching site j and those that had found a relevant product at some firm prior to firm j, but that have low valuations. Both are multiplied by a product of search frequencies to account for consumers that end their search without making a purchase.

The CTR is decreasing in list rank. This occurs for two reasons. One, some consumers find a suitable product, make a purchase, and quit searching ($q_k > 0$). Two, only a fraction of consumers continue searching down the list ($s_k \leq 1$) . A falling CTR is a well-known feature of ad listings and it is important that our model reflect this important empirical reality.

The CTR measures the size of the market that a firm faces. Now, consider the overall demand that each firm receives. In our model, consumers that purchase from firm j entered the list, were not satisfied by any of the previous $j-1$ firms, searched all the way to firm j, found a relevant product at firm j, and have a valuation above the price. Putting these pieces together, the

[5] It may seem odd that the model permits these consumers to fruitless search forward. This is to enable us to apply this model to a market with varying, perhaps endogenous, prices. We could consider the opposite scenario whereby low value consumers realize, after visiting the first site, that they can never find a product cheap enough for them to purchase and quit searching immediately. This is *attrition by low value consumers*, the opposite of the attrition by high value consumers described here. In Appendix A, we show that the framework developed in this section, notably the use of adjustment factors, is sufficiently broad that it can encompass attrition by low value consumers as well. In fact, with appropriate definition of the adjustment factor, the qualitative results in the paper do not change with this alternate conception of consumer behavior.

demand for firm 1 is

$$D_1(p) = s_0 q_1 [1 - F(p)]$$

and the demand for firm $1 < j \leq M$ is

$$D_j(p) = [1 - F(p)] s_0 q_j \prod_{k=1}^{j-1} s_k (1 - q_k).$$

A firm not on the list, $j > M$, does not face any demand. This is the gross sales that a firm makes and it varies across firms.

We see that demand is falling as we move down the list for the same reasons that the CTR was decreasing: some consumers are satisfied by previous firms and some consumers stop searching altogether. This naturally leads us to ask whether the ratio of overall demand to the CTR—the demand per click—is decreasing as well. As we see below, demand per click is also important in determining the value of a click to a firm, which is used in Section 3 to calculate the optimal bid in the placement auction. The demand per click is

$$\begin{aligned} \frac{D_j(p)}{r_j} &= \frac{[1 - F(p)] s_0 \prod_{k=1}^{j-1} s_k (1 - q_k)}{\prod_{k=0}^{j-1} s_k \left[\prod_{k=1}^{j-1} (1 - q_k) + F(p) \left[1 - \prod_{k=1}^{j-1} (1 - q_k) \right] \right]} q_j \\ &= a_j q_j, \end{aligned}$$

where

$$a_j = \frac{[1 - F(p)] \prod_{k=1}^{j-1} (1 - q_k)}{\prod_{k=1}^{j-1} (1 - q_k) + F(p) \left[1 - \prod_{k=1}^{j-1} (1 - q_k) \right]}. \tag{2}$$

The value $a_j \in (0, 1]$ is a slot-specific adjustment factor. Note that $a_j = 1$ for all slots if $F(p) = 0$; that is, if there are no low-value consumers. Otherwise, the adjustment factor is decreasing down the list.

The adjustment factor accounts for two features of the model. First, while firm j offers a product relevant to q_j consumers, a portion of these consumers have valuations lower than the price. Second, the distribution of valuations changes by the slot: consumers with valuations above the market price make purchases and quit searching, while low-valuation consumers continue searching. Thus, the profitability of the pool of consumers visiting each site decreases down the list. This implies that, at a fixed price, firms further down the list make fewer sales than they would expect based upon their relevances alone. From the perspective of the firm, too many consumers (*i.e.*, the low valuation ones) continue searching. This is called *attrition by high value consumers*.

The adjustment factor accounts for the changing profitability of the pool of consumers visiting the site in a particular slot. This framework is general enough such that adjustment factors can be calculated in other contexts when profitability is changing down the list. Rather than permit low value consumers to continue searching knowing that they can never find a product with a price low enough for them to make a purchase, we could have these consumers leave the market entirely after visiting the first listing. Appropriately defined adjustment factors can be used in this situation without changing the overall structure of the model or the auction in the following section (see Appendix A). As another example, it is plausible that high value consumers continue searching longer than low value consumers, increasing the profitability of slots lower on the list. Again, adjustment factors could be defined to capture this behavior.

As detailed in the next section, firms bid for slots in the ad listing, paying the ad server for each click received. This bid is going to depend upon the value of a click to the firm. We need to normalize the total profit received by the number of clicks received to calculate the expected value per click:

$$\frac{(p - c_j)D_j(p)}{r_j} = m_j a_j q_j, \tag{3}$$

where m_j is the margin for firm j.

Indexing the CTR and adjustment factor according to firm j obscures the calculation of these values. Looking to Equations 1 and 2, we see that r_j and a_j do not depend upon the firm j's own relevance q_j, but rather, they depend upon the relevances of the preceding firms. It is more accurate to say that the CTR and adjustment factor are specific to slot j, regardless of the firm that takes that position, given that firms $1, \ldots, j-1$ remain the same.

The prime reason for defining the adjustment factor is to break the expected value per click into a portion that is the same for the firm no matter where it is listed ($m_j q_j$) and a portion that depends upon its actual slot a_j. As mentioned earlier, the indexing of firms corresponds to their equilibrium slot placement, but proof of this fact in Section 3.1 requires a distinction between firm-specific and slot-specific quantities. If firm j is placed into slot k out of equilibrium, then the expected value per click would be $m_j a_k q_j$ and it would receive a CTR of r_k.[6]

2.3 Discussion

Previous work in the ad auction literature assumes that the value that a firm places on being at a particular ranking can be separated into a CTR effect and a firm-specific value effect. CTRs are assumed to decrease monotonically down a list, but a firm has the same value per click of being in any slot. Though a lower-ranked firm may receive fewer clicks, each click has the same value to that firm whether the firm was in the first slot or the last. In these models, consumers are identical and thus there can be no selection in the group that continues searching. If there is attrition by high-valued consumers, this structure is called into question.

One paper that does incorporate heterogeneous valuations is Chen and He (2006). Their framework combines consumers with differing valuations, but identical search costs that increase with the number of sites visited and endogenize pricing decisions by firms. They do not consider the potential for selection effects in the distribution of consumer valuations down the list. When Chen and He (2006) consider the firms' pricing decisions in their Equation 1, they assert that all firms face the same pricing decision, yielding no price dispersion, but they do not consider that firms may face different demand conditions depending upon their ranks and, as a result, the firms' maximization decisions will vary. In particular, firms further down the list face fewer high-value consumers and may be inclined to cut prices due to attrition of high-value consumers. Our base model does not endogenize pricing decisions, hence, we do not evaluate this strategy here.

3 Ad Auction Bidding Behavior

Contextual ads are sold using a Generalized Second Price (GSP) auction. A firm places a bid to be included in the ad listing based upon keywords that appear in the substantive content (search queries, articles, reviews, *etc*) on the page. In the simpler case developed by Overture for Yahoo, firms are assigned slots in decreasing order of their bids. A firm pays the bid of the next ranked firm each time that its own ad is clicked. Many prominent papers have focused on this framework (see, *e.g.*, Edelman, Ostrovsky and Schwartz, 2007; Varian, 2007; Athey and Ellison, 2009).

[6]Adding to the complexity, the out of equilibrium values for the adjustment factor a_k and CTR r_k would be different from their equilibrium values if one or more preceding firms, $1, \ldots, j-1$ also changed position.

In Google's auction, firms are ranked by the product of their bids and their "quality scores" and a firm pays the product of the bid and quality score of the next firm down the list on a per-click basis. Quality scores aim to estimate the expected CTR for a firm's ad.[7] For example, for the keyword "airplane," suppose that both Boeing and a toy airplane manufacturer would like to have their ads listed. Boeing may be willing to pay more for a listing because, if a click turns into a sale, the firm earns greater profit relative to the profit earned on the sale of a toy plane. Few viewers are interested in purchasing jumbo jets, however, so the Boeing ad receives few clicks, earning Google little revenue. Google could earn greater revenues by putting a firm with a lower bid but higher firm-specific CTR at the top of the list than a firm that bids high, but receives few clicks.

Advertisers can change their bids frequently; this might lead us to model the auction as an infinitely repeated game. The Folk Theorem, however, asserts that these games have many equilibria, rendering analysis extremely difficult. Resultingly, most work has focused on the single-shot version of the auction game to identify an equilibrium.

An early effort in the literature, Edelman, Ostrovsky and Schwartz (2007) places the GSP auction in the context of established auction designs, including the second price auction, Vickery-Clarke-Groves (VCG) mechanism, and the ascending English auction. They show that the GSP auction is not equivalent to the VCG mechanism. Unlike VCG, this auction does not have an equilibrium in dominant strategies and truth-telling is not an equilibrium. Under a set of restrictions, one of the equilbria that arises provides the same payoffs as under the dominant strategy VCG equilibrium. Edelman, Ostrovsky and Schwartz (2007) call these equilibria "locally envy-free equilibria." Varian (2007) independently identifies the same equilibria and calls them "symmetric Nash equilibria." The ad intermediary is better off at any other locally envy free equilibrium other than the one equivalent to the VCG equilibrium, while advertisers are worse off.

Most of the existing literature on advertising auctions has focused on the elements of optimal auction design. Alternative mechanisms have been offered that provide higher profits to ad intermediaries or more efficient assignments of ad slots. Other papers extend the standard GSP framework by incorporating the quality scores found in Google auctions or other weighting schemes and reserve prices. This paper focus on the properties of the standard auction mechanism, but

[7] The quality score of a firm is no longer this transparent. Many papers consider the optimal weights to maximize auction revenue, but we do not consider this literature here.

incorporates the consumer behavior underpinning click-through rates. While the structure of the auction is undoubtedly important for firms and the ad server, we ignore these complexities and use the simplified version of the auction developed by Yahoo/Overture in our analysis.

3.1 Ranking of firms in the ad listing

We begin by incorporating our model for CTR into the approach of Varian (2007), specifically, a one-shot, simultaneous move, complete information game. Of the J firms in the market, M appear on the ad list.[8] The CTR for firms $M+1,\ldots,J$ is 0, while a firm on the list in slot j experiences a CTR r_j following Equation 1. Varian (2007) assumes that the CTR is exogenous and decreasing down the list; in the preceding section, we provide a behavioral foundation for this assumption.

A firms is charged on a per-click basis at a price equal to the bid of the firm one slot down on the ad list.[9] Recall from Equation 3 that the expected value per click for firm j in slot j is $m_j a_j q_j$. And, based upon the discussion on page 8, the expected value per click of firm j in the off equilibrium slot k is $m_j a_k q_k$ with a CTR of r_k. In the symmetric Nash equilibria of Varian (2007), the expected profits in firm j's equilibrium slot must be weakly higher than those it receives in any other slot k:[10]

$$r_j(m_j a_j q_j - b_{j+1}) \geq r_k(m_j a_k q_j - b_{k+1}). \tag{4}$$

Note that the CTR and the slot-specific adjustment factor change with the slot for a given firm, but the relevance of the firm and its margin do not. The firm faces the following trade-off: Accepting a lower slot on the page requires a smaller payment for the slot. However, the firm receives fewer clicks in this space and faces a less profitable pool of consumers.

[8] We do not consider the case of "unsold pages," where there are fewer willing bidders than slots. Additionally, we assume that the highest $M+1$ firms all bid above the reserve price of the auction.

[9] Bear in mind that firms lower on the list have higher indices—firm j is one slot above firm $j+1$.

[10] The CTR for slot k depends upon the relevances of the firms $1,\ldots,k-1$. As a result, the CTR for slot k is different when different firms are in the preceding slots. If firm j moves up to slot k, then the ordering of firms $1,\ldots,k-1$, and thus r_k and a_k remain unchanged. If firm j moves down the list to slot k, this changes the firms that appear in slots $1,\ldots,k-1$. If firm j moves down the list, then r_k would be an out-of-equilibrium CTR. We will claim that, in equilibrium, higher-ranked firms must have weakly higher relevances. If a relatively high relevance firm j moves down the list to slot k, fewer consumers find a relevant product from firms $1,\ldots,k-1$. This implies that more consumers search forward and these out-of-equilibrium CTRs for slots $j,\ldots,k$ are higher than their equilibrium values. The following equation fails to distinguish between in- and out-of-equilibrium CTRs. Since the latter are weakly higher, this inequality remains valid.

Consider the equilibrium conditions for firm j moving to slot k and firm k moving to slot j:

$$m_j q_j (a_j r_j - a_k r_k) \geq r_j b_{j+1} - r_k b_{k+1} \text{ and}$$
$$-m_k q_k (a_j r_j - a_k r_k) \geq -r_j b_{j+1} + r_k b_{k+1}.$$

Adding these inequalities together gives

$$(m_j q_j - m_k q_k)(a_j r_j - a_k r_k) \geq 0. \tag{5}$$

Recall from Section 2.2 that the CTR r and the adjustment factor a are both decreasing down the list. This expression reveals that the relevance q times the margin m must move in the same direction, namely, decreasing down the list; the full expected margin is decreasing down the list. We assume that there are no ties in the full expected margin $m_j q_j$, ensuring that this ranking is unique.

3.1.1 Varying margins, constant relevance

An interesting special case is when $q_j = q$ for all j. Here, firms sort in decreasing order of margins. All firms charge the same price p and have the same relevance; consumers are indifferent to the order of firms that they search. In the case of indifference, assume that consumers still search from the top down. While the ordering of the firms has no impact on consumer surplus, producer surplus is largest when firms sort in increasing order of costs—that is, decreasing order of margin. This is precisely the result given by the auction, hence, total surplus is maximized.

3.1.2 Varying relevances, constant margins

At the other extreme, suppose that firms all have the same costs and thus the same margin, but have different relevances. The equilibrium condition reveals that the firms sort in decreasing order of relevance. Consumers prefer to visit the sites most likely to offer a relevant product. Given the bidding strategies of the firms, this would imply that consumers should search starting from the top of the list, confirming this outcome as an equilibrium. Since consumers visit a limited number of sites in order, the greatest number of sales occur when the most relevant firms are listed at the top; this ranking also maximizes both consumer and producer surpluses.

3.1.3 Varying margins and relevances

Of course, the intermediate cases are perhaps the most interesting and most difficult to characterize. Considering the expected ordering of firms, we ask how a firm's cost is correlated with its relevance. If these factors are negatively correlated correlated, then we expect the low cost, high relevance firms to be at the top and the high cost, low relevance firms to be at the bottom. This combines the results of the two previous subsections.

We can go further by considering the case that the cost of firm j is $c+\alpha q_j$. We could impart a causal story: it is more or less costly (depending upon the sign of α) to produce a product that a high proportion of people like. Or we could consider the model as one of association, used only to highlight existing correlations between relevance and cost. Our equilibrium condition becomes

$$\left[(p-c)(q_j-q_k)-\alpha\left(q_j^2-q_k^2\right)\right](a_j r_j - a_k r_k) \geq 0.$$

If the CTR is falling down the list, as it does when consumers search from top to bottom, then firms sort in decreasing order of relevance if

$$\frac{p-c}{q_j+q_k} \geq \alpha \tag{6}$$

and sort in increasing order of relevance otherwise. Note that the lefthand side of this expression is positive. The CTR is only decreasing down the list if consumers have an incentive to search downward; this is the case if relevance is weakly decreasing down the list. Hence, if α satisfies Equation 6, then this equilibrium exists. Intuitively, this condition states that the relevance and cost of a firm can covary positively, to a point, and still sort in decreasing order of relevance. Firms with smaller per-sale margins have higher expected margins due to their higher relevance.

3.2 Deriving equilibrium bids

To find the bids chosen by the firms, we return to Equation 4. Varian (2007) shows that, if this equation holds for a firm moving up one slot or down one slot (*i.e.*, from j to $j-1$ or to $j+1$), then it holds for a move to any slot or a move off the list entirely. Using the fact that firm j does not want to move to slot $j+1$ and that firm $j+1$ does not want to move to slot j, we find that[11]

$$m_j q_j \left(a_{j-1} - \frac{r_j}{r_{j-1}} a_j\right) + \frac{r_j}{r_{j-1}} b_{j+1} \leq b_j \leq m_{j-1} q_{j-1} \left(a_{j-1} - \frac{r_j}{r_{j-1}} a_j\right) + \frac{r_j}{r_{j-1}} b_{j+1}. \qquad (7)$$

These bounds can be solved recursively by recalling that $r_j = 0$ for the firms not listed, firms $j > M$, yielding

$$\frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} m_k q_k (a_{k-1} r_{k-1} - a_k r_k) \leq b_j \leq \frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} m_{k-1} q_{k-1} (a_{k-1} r_{k-1} - a_k r_k). \qquad (8)$$

Firm j can bid any value in this range without changing its slot or the slot of other firms.

We can learn many things from these recursive bounds. First, firms have positive profits. To see this, return to Equation 4 and set $k = M+1$, the first firm not listed. Here, $r_k = 0$, implying that $m_j a_j q_j \geq b_{j+1}$. Hence, the net profit from being in slot j, $r_j(m_j a_j q_j - b_{j+1}) \geq 0$. The lower bound of Equation 8 is less than or equal to the expected margin from slot j, $m_j q_j a_j$—firms may shade their bids. It is possible that the upper bound of Equation 8 is above the expected margin, implying that a firm may bid above its valuation. The logic here is that the firm in slot j must bid high enough so that the firm just above it in slot $j-1$ does not have an incentive to switch to slot j and sacrifice clicks to increase per-click profit. Nonetheless, the expected margin from the slot must be positive; remember, the firm does not pay its own bid, but rather the bid of the firm below it in the listing.

We can rearrange this equation to get the ad revenue raised from firm j:

$$\sum_{j+1 \leq k \leq M+1} m_k q_k (a_{k-1} r_{k-1} - a_k r_k) \leq r_j b_{j+1} \leq \sum_{j+1 \leq k \leq M+1} m_{k-1} q_{k-1} (a_{k-1} r_{k-1} - a_k r_k).$$

Total ad revenue has a lower bound of

$$\sum_{2 \leq k \leq M+1} (k-1) m_k q_k (a_{k-1} r_{k-1} - a_k r_k)$$

[11]This procedure actually gives the bounds for b_{j+1}; appropriate reindexing gives the result shown.

and an upper bound of

$$\sum_{2\leq k\leq M+1} (k-1)m_{k-1}q_{k-1}(a_{k-1}r_{k-1} - a_k r_k).$$

The pattern of bid shading is not obvious from Equation 8. We calculate the upper and lower bounds for the bid of the firm in each slot for a particular set of parameter values. For concreteness, let each firm have the same relevance $q = 0.2$, with the proportion of low-value consumers set to $F(p) = 0$ and all search frequencies set to $s = 1$. The margins of the firms vary from 0.9 in slot 2 to 0.1 in slot 10 in increments of 0.1.

Figure 1 plots the ratio of the bid of firm j to the expected margin $m_j a_j q_j = m_j q_j$ in this case (as there are no low-value consumers). The pattern in the shading of the bids is not monotone; shading is highest for firms in the middle of the listing, higher for top-ranked firms, and non-existent for the final firm (as was obvious from Equation 8). The range of possible bids gets larger moving down the list.

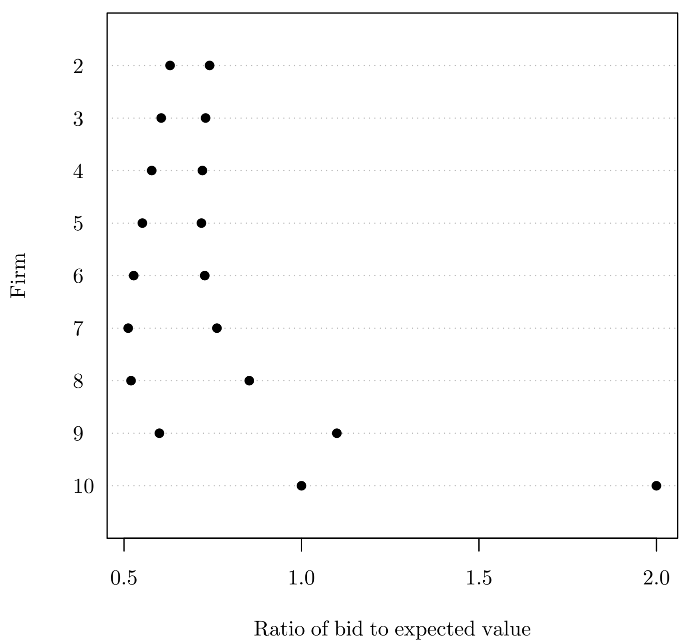

Figure 1: Ratio of bid to expected margin

Of course, firms do not pay their own bids; they pay the bids of the next firm down the list. We may wonder how the ratio of the cost of being in slot j, which is the bid of firm $j+1$, to the expected margin of firm j varies down the list. Using the same parameter values as Figure 1,

Figure 2 shows these results. Again, the pattern here is not monotone. Instead, this ratio falls through slot 7, then increases sharply. For both figures, the pattern depends upon the chosen parameter values, but these illustrations reveal that overall patterns are difficult to characterize.

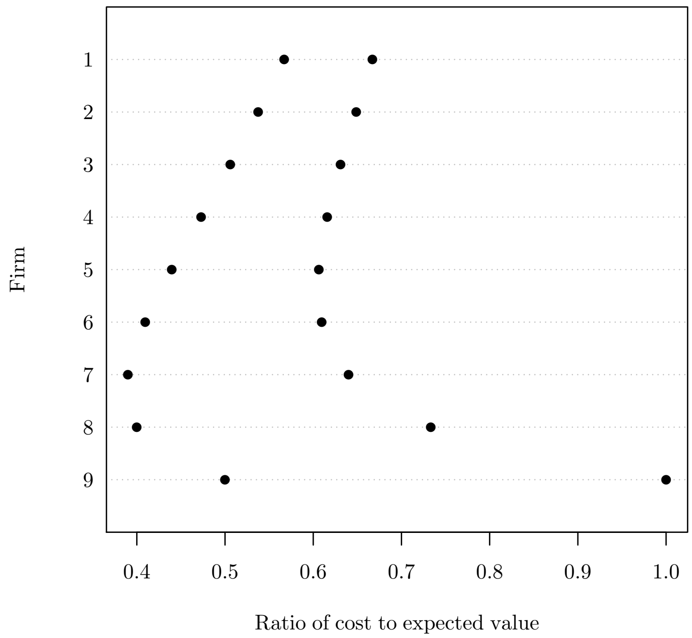

Figure 2: Ratio of cost per click to the expected value per click

3.3 Discussion

Varian (2007) arrives at these results by assuming complete information. He offers several justifications for this assumption. First, Google reports view and click rates on an hourly basis to bidders and, if bidders experiment with different bidding strategies, they can infer many of these quantities fairly quickly. Additionally, Google offers a "Traffic Estimator" that predicts the number of clicks and total costs for different bid-keyword combinations. Lastly, private, experienced search engine optimizers can offer clients assistance with bidding strategies.

Athey and Ellison (2009) do not assume perfect information in their model. Instead, they choose to study an incomplete information game.[12] Specifically, each firm j's relevance, q_j, is private information, while all margins $m_j = 1$. Additionally, rather than consider a simultaneous-move game, they consider a multistage game where each slot is sold separately from bottom to top.

[12]Varian (2007) offers a version of his model in an incomplete information setting as an appendix to his paper. In his case, however, the CTRs remain common information and exogenous, while Athey and Ellison (2009) turn to an incomplete information context precisely because the CTRs are functions of the unknown relevances.

This change lets the authors condition on bids for lower slots in solving for higher bids; we sense why this is important by recalling the equilibrium conditions of the complete information model given by Equation 8.

Most papers in the position auction literature take the CTRs to be separate, exogenous parameters. Here and in Athey and Ellison (2009), the CTRs are derived using behavioral assumptions about consumers and are functions of the firms placed above a given slot, *i.e.*, the CTR for firm j depends upon the relevances of firms 1 through $j-1$. This implies that the value that a firm obtains from being in a given slot depends upon those ranked above it. As Athey and Ellison (2009) note, this fact implies that the auction is no longer based upon private values.[13] The common values aspect is easily separated from the private value of appearing in the list and this wrinkle is not difficult to handle.

The authors use a perfect Bayesian equilibrium solution concept to solve the game. Using this procedure, the authors provide a bid function that is qualitatively the same as that of Varian (2007) and thus the behavioral model of our paper could easily be placed within their framework without fundamentally altering the results that follow.

4 Incentives of ad servers to change search structure

There are three quantities that define the consumer side of the market: the tastes of consumers given by the match probabilities q, search costs that are implicit in the search frequencies s, and the valuations for the product v. A fourth quantity in the model that describes the supply side is the per-sale margin m. These factors are not immutable, however; altering these quantities may change the profits and overall welfare generated in the product market. In this section, we explore the incentives than an ad server has to alter these quantities and how these changes impact firms and consumers.

The prime incentive for the ad server to produce such changes stems from changes in advertising revenue. Equation 8 reveals that the lower bound for the advertising revenue generated

[13] Additionally, the GSP is no longer the most efficient mechanism for ranking ads (Aggarwal et al., 2008). We do not look for the optimal auction design in this paper.

after a change in the market structure would be

$$\sum_{2\leq k\leq M+1}(k-1)\big[m_kq_k+\Delta(m_kq_k)\big]\big[(a_{k-1}r_{k-1}-a_kr_k)+\Delta(a_{k-1}r_{k-1}-a_kr_k)\big].$$

The change in revenue, then, is given by

$$\sum_{2\leq k\leq M+1}(k-1)\Big[\Delta(m_kq_k)\big[(a_{k-1}r_{k-1}-a_kr_k)+\Delta(a_{k-1}r_{k-1}-a_kr_k)\big]+m_kq_k\Delta(a_{k-1}r_{k-1}-a_kr_k)\Big]; \tag{9}$$

the increase in the margin times a number that is a function of the new CTRs and adjustment factors plus the old margin times the change in the CTRs and adjustment factors. The first piece captures the increased value-per-click to a firm after the change. The latter captures whether the number of clicks has changed.

We find $a_{k-1}r_{k-1}-a_kr_k$ by noting that, by definition, $a_kr_k=\frac{D_k(p)}{q_k}$; an analogous result is found for firm $k-1$. The difference between these two quantities is

$$[1-F(p)]s_0\prod_{p=1}^{k-2}s_p(1-q_p)\left[1-s_{k-1}(1-q_{k-1})\right]. \tag{10}$$

To find $\Delta(a_{k-1}r_{k-1}-a_kr_k)$, we can analyze changes in Equation 10.

4.1 Proportional changes in the relevances

Suppose that the ad server has the ability to boost all firms' relevance by a certain percentage. This could occur by achieving a better matching algorithm, by using information known about a particular user, or, rather than increasing the relevances of given firms, by having bigger pool of advertisers, thereby yielding more high quality matches.

4.1.1 Intuition from the model

Consider this change in the context of Equation 9. Since q_k goes up, $\Delta(m_kq_k)$ is positive and the first component of the sum is positive. For expository purposes, let all firms have the same relevance q. Then, Equation 10 becomes

$$[1 - F(p)](1 - q)^{k-2} \prod_{p=0}^{k-2} s_p \left[1 - s_{k-1}(1 - q)\right].$$

Taking the derivative with respect to q yields

$$[1 - F(p)](k - 2)(1 - q)^{k-3} \prod_{p=0}^{k-2} s_p \left[s_{k-1}(k - 1)(1 - q) - (k - 2)\right].$$

A necessary condition for the quantity in question and thus ad revenue generated by a particular slot to be increasing is[14]

$$s_{k-1}(1 - q) > \frac{k - 2}{k - 1}.$$

This inequality does not hold in general; it holds only for highly ranked firms.

Firms receive a higher margin per click because a consumer is more likely to find a relevant product on its site. This implies, however, that more consumers are satisfied high on the list and do not visit lower-ranked sites. While the expected margin may be higher, the pool of consumers is smaller. These conditions act as opposing forces in changing the ad revenue generated by a firm.We expect the bids of high-ranked firms to increase more after the change in the relevances compared to lower-ranked firms. However, as bids are solved recursively, drops in the bids of lower-ranked firms temper increases in higher-ranked firms. As firm 1 does not experience any drop in its CTR or adjustment factor, we expect it to exhibit the greatest change in advertising revenue generated.[15] Equation 10, thought of more simply, is the difference in CTRs between firms $j - 1$ and j.[16] This result states that this gap is bigger for firms high on the list and smaller for firms low on the list after the change compared to the previous, lower set of CTRs.

4.1.2 Simulation of the change

While these calculations give us some intuition for the impact of a change in relevances has on ad revenues, let us consider a numerical example. Largely irrelevant to these calculations are the

[14]Recall that bids are calculated starting at $k = 2$; division by 0 is not an issue here.

[15]One issue not yet discussed is that, if an ad server can increase the relevance of its ads, then it may attract a larger pool of consumers to its site, increasing the size of the market for all firms.

[16]This occurs when $F(p) = 0$; everyone is a high-value consumer.

search frequencies s_k and the proportion of low-value consumers $F(p)$; set the former all to 1 and the latter to 0 for simplicity. Assume that all firms have the same relevance of 0.2. We consider an increase in this value by 20%.

Figure 3 gives the impact of this change on ad revenues, bids, and gross and net (of advertising costs) firm profit. First, we note that, in this case, the CTR drops by a factor of $\left[\frac{1-1.2\times0.2}{1-0.2}\right]^k$ for site k. After the 20% increase in relevance, firms bid at least 20% more.

The highest increases in bids come from firms in the middle. High ranked firms do not experience a large change in their CTRs. Middle ranked firms have large drops in their CTRs and need to bid higher to avoid slipping down the list and experiencing even greater changes. Firms low on the list had low CTRs anyhow and, while the drop may be relatively larger than for other slots, the absolute drop is smaller and these firms do not have as strong an incentive to bid to avoid it.

Ad revenue is a product of the CTR and the bid. Higher bids more than offset the reduced CTR for firms 1 through 6, increasing the ad revenue generated by these firms. For the last 3, ad revenue decreases. Total ad revenue increased by 21%.

Firm profits increase for the first 4 firms, but fall for the remainder; for the latter group, the higher match probability (and thus expected margin) is offset by fewer clicks. These first few firms generate more ad revenue and increases in gross profits are eaten away by higher advertising costs. Indeed, only the first 2 firms have higher net profits after the increase in relevances. Total firm net profits across all the firms actually fell by 2.2%.

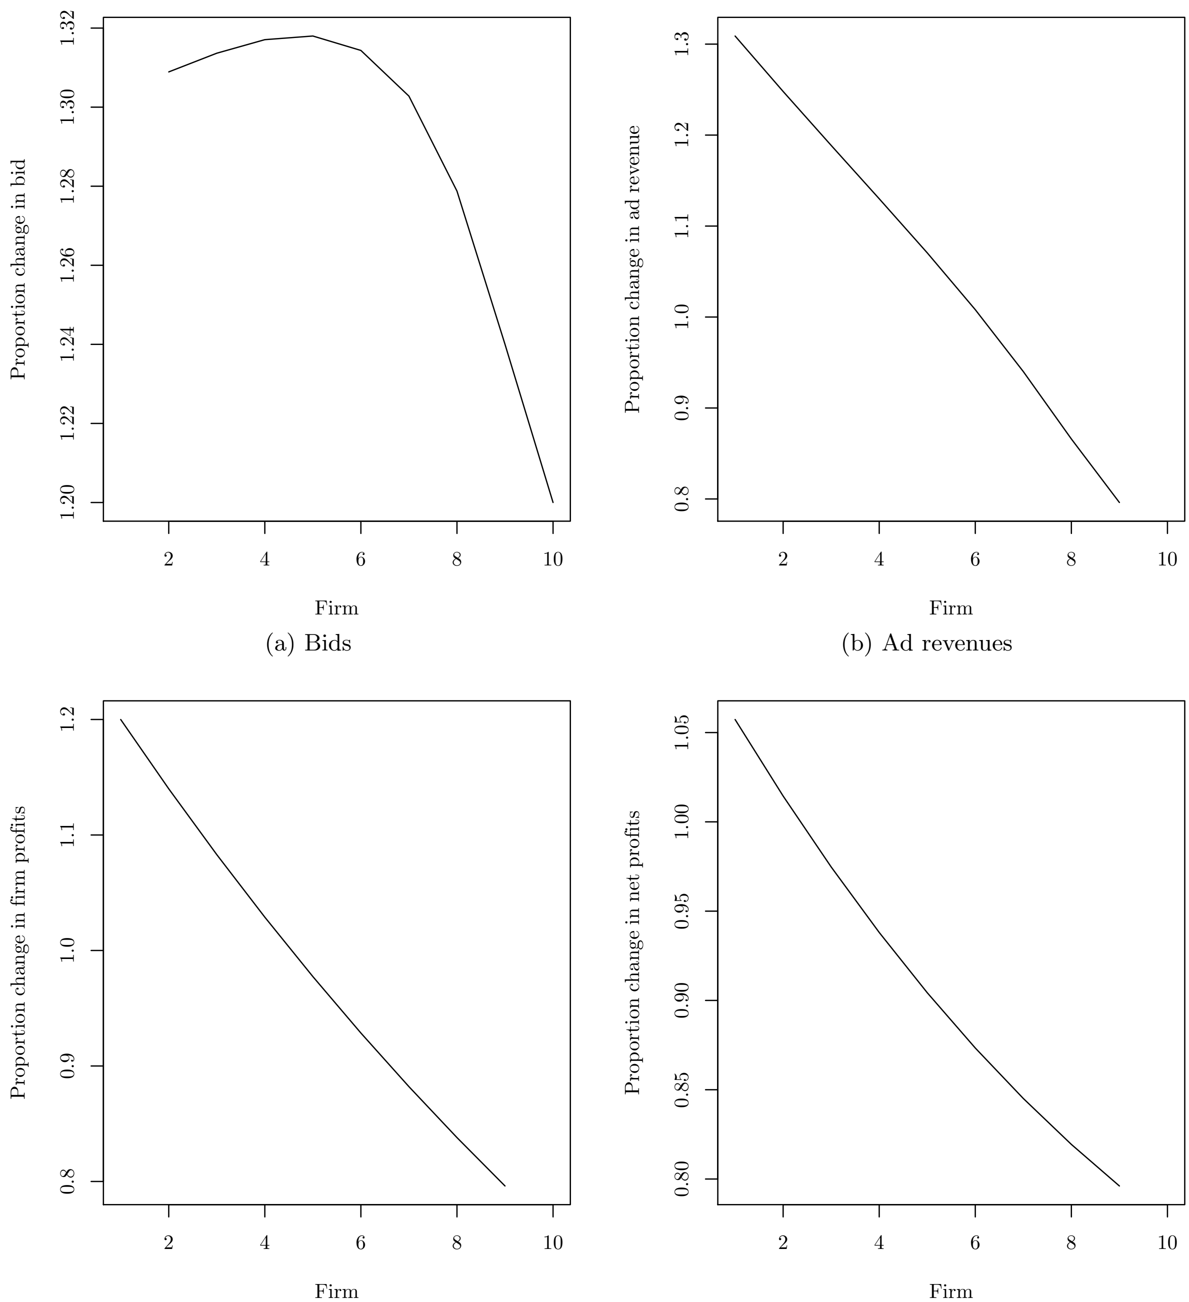

(a) Bids

(b) Ad revenues

(c) Firm profits

(d) Firm net profits

Figure 3: Impact of a 20% increase in relevance from $q = 0.2$

For this particular increase in the relevances, the ad server earns higher revenues, while firms' net revenues fall. This is not necessarily the case. Figure 4 show the total ad revenue, ad elasticity, and total firm gross and net profits across changes in the base relevance of 0.2 by factors of 0.5 to 1.5. "Total" refers to measures summed across all firms. By "elasticity," we mean the proportion change in ad revenues divided by the proportion change in relevance.[17]

Revenues for the ad server and profits for the firms are both increasing with the relevance. The ad revenue elasticity and total firm net revenues have maximum values, however. The ad revenue elasticity is maximized at a proportion increase of 1.2, an increase from 0.2 to 0.24. This is higher than the point where firm net profits are maximized, at a relevance of 0.19. These plots reinforce that firms in aggregate may be hurt by overall increases in match probabilities.

Next, in Figure 5, we turn to the consumer side of the market. Since we assume that prices are constant across firms and unchanging with q, consumer welfare is higher if sales are higher. In Figure 5a, we see how sales change by firm in the context of the analysis leading to Figure 3. We see that the first 4 firms increase their sales, while the remaining firms have lower sales. Sales increase because the match probability increases, but fall for lower-ranked firms because fewer consumers reach these firms without already being satisfied (*i.e.*, the CTR is lower). Overall sales are increasing in q, as seen in Figure 5b, an analogue to the panels of Figure 4. Consumers are unambiguously better off if the ad server increases the relevances.

[17] If improving matches has linear cost with no fixed costs, then this would be the elasticity of revenue with respect to costs. This cost function for improving matches by the ad server is highly unlikely, but this calculation provides useful insights nonetheless.

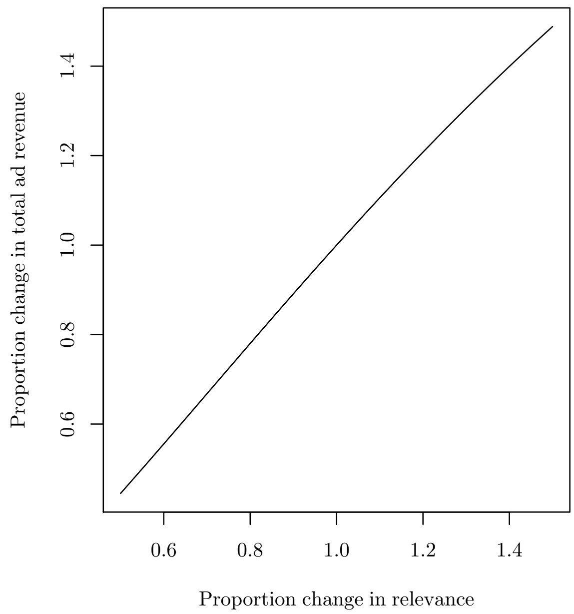

(a) Total ad revenue

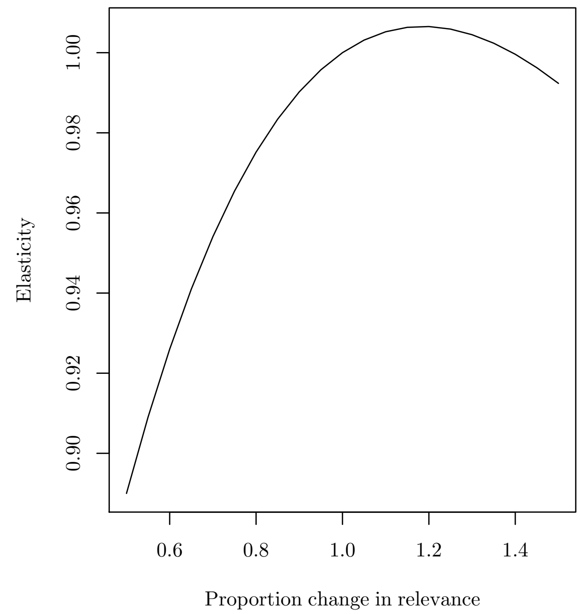

(b) Ad elasticity

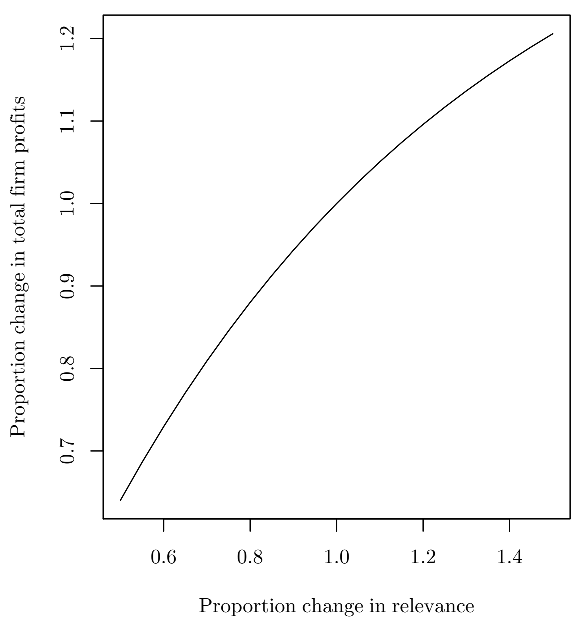

(c) Total firm profits

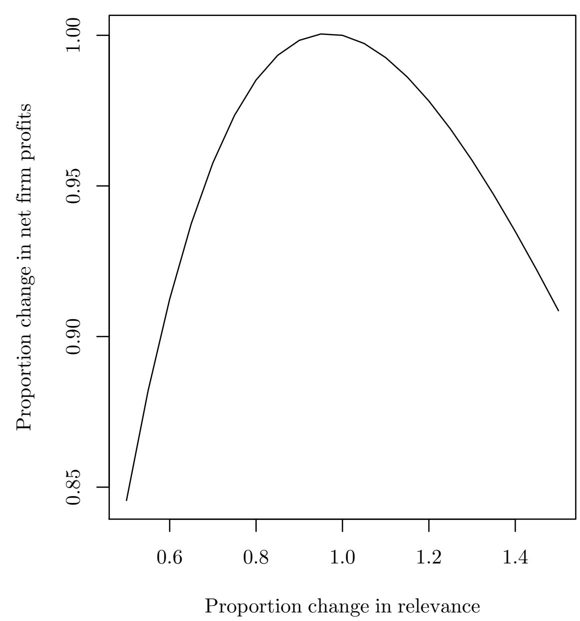

(d) Total firm net profits

Figure 4: Impact of changes in relevance from $q = 0.2$ on aggregates

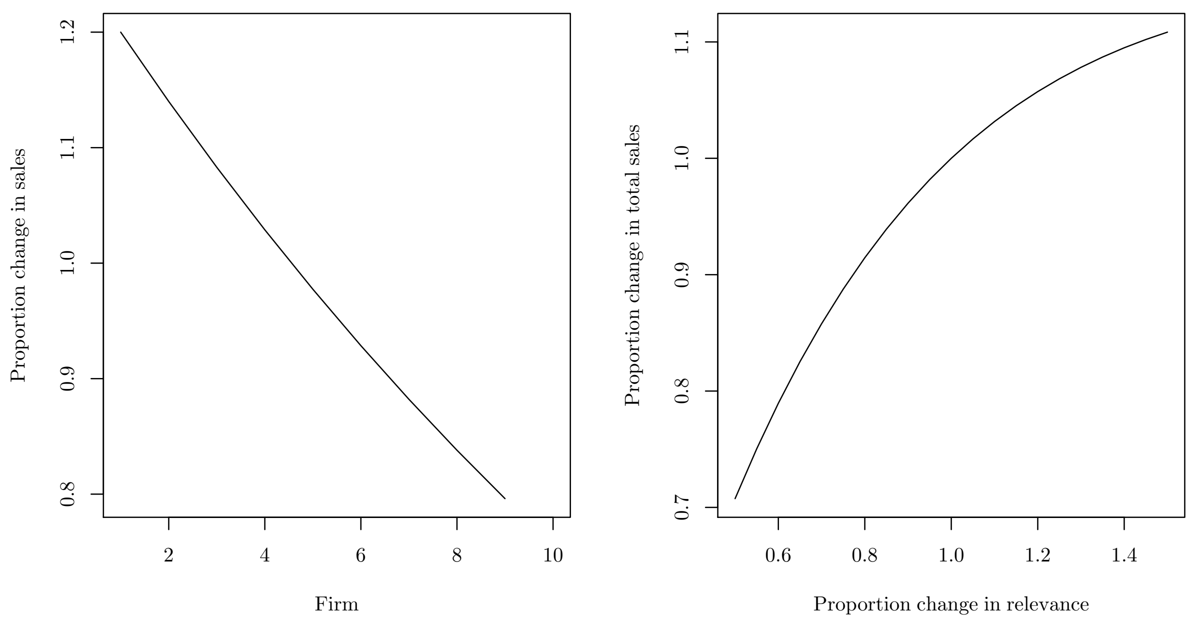

(a) Sales by firm, $q = 0.2$ increased by 20%

(b) Total sales for a range of proportional changes in $q = 0.2$

Figure 5: Impact of changes in relevance for consumers

4.2 Proportional changes in search costs

The ad server may also be able to reduce search costs. Practically, this may mean caching pages for faster loading, subsidizing high-speed internet access, or making consumers more proficient searchers. Unlike in the case of increasing relevance, this change does not alter firms' expected margins. Instead, it just increases the size of the customer base visiting each site. We imagine that such a change should leave both firms and the ad server better off.

4.2.1 Intuition from the model

Again, return to Equation 9. The full margin $m_k q_k$ does not change, leaving only Equation 10 to consider. For simplicity, let all the search frequencies be the same. This equation becomes

$$[1 - F(p)]s^{k-1} \prod_{p=1}^{k-2} (1 - q_p)[1 - s(1 - q_{k-1})].$$

Taking the derivative with respect to s gives

$$[1 - F(p)]s^{k-2} \prod_{p=1}^{k-2} (1 - q_p)[(k - 1) - ks(1 - q_{k-1})].$$

Hence, a sufficient condition for the quantity to be increasing in search frequencies is

$$s(1 - q_k) < \frac{k - 1}{k}.$$

This inequality only holds for firms low on the list. It is perhaps surprising that ad revenues do not unambiguously increase for all firms. This is because, while the firm receives a higher CTR, the cost in terms of lost sales to moving a slot down the list is less severe (as a higher proportion of consumers now visit this site) and bids are shaded more.

4.2.2 Simulation of the change

Following an analogous analysis to that for altering the relevance, we consider an example with $q = 0.2$ and a base search frequency of 0.6, both of which are the same across firms. We consider increasing the search frequency frequency by 20%. The results are given in Figure 6.

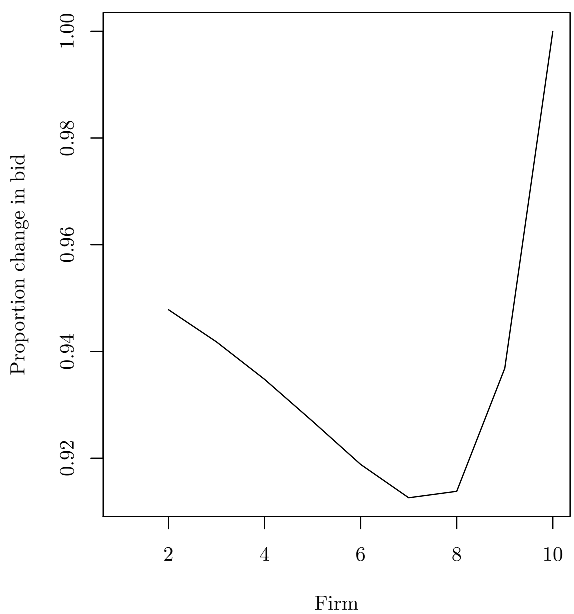

(a) Bids

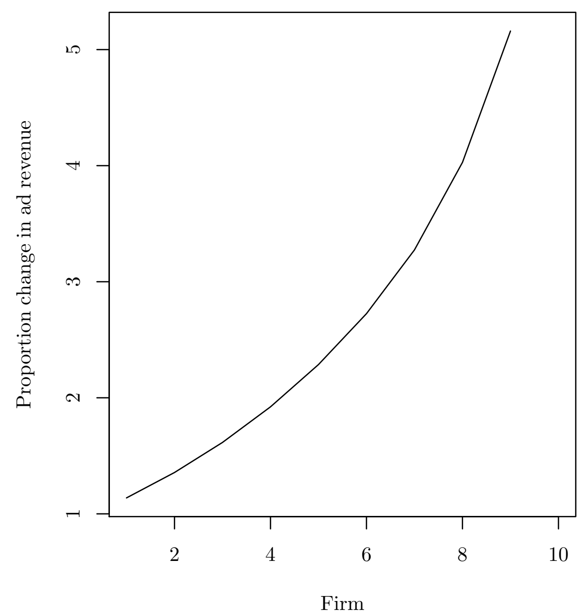

(b) Ad revenues

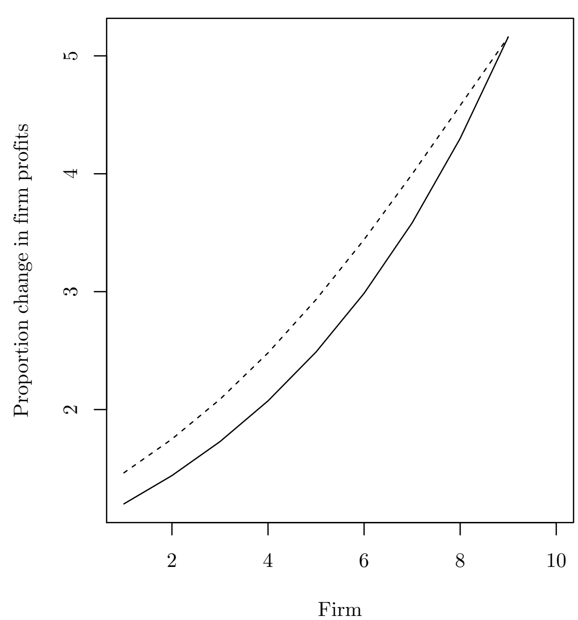

(c) Firm revenues, gross (solid), net (dashed)

Figure 6: Impact of a 20% increase in search frequencies from $s = 0.6$

Bids decrease for all firms except the first excluded firm. Recall that this firm bids its true valuation per click for being included in the list; since this has not changed, neither has its bid. Reduced bids are more than offset by higher CTRs, as evidenced by the fact that ad revenue from every site increases—by dramatic proportions for many slots. Site 1 has the smallest increase in ad revenue, a change of 14%, smaller than the change in visitors (20%). All other firms increase the ad revenues that they generate by a larger percentage than the change in search frequencies. This is sensible, as changes in search frequency compound and the proportion increase in the size of the consumer group after the change gets larger down the list. Firm net profits increase by a larger percentage than gross profits. Unlike the case of increasing relevance, firms keep a large share of the gains from increasing search frequencies.

We can explore these properties in aggregate across a variety of changes in search frequencies; we show the resulting patterns in Figure 7. These plots show that constant increases in search frequencies benefit both the ad server and the firms. Net revenue is the most responsive of all, suggesting that gains in search frequencies mostly benefit the firms. A larger proportion of consumers make purchases when a larger proportion search forward and more purchases leads to higher consumer surplus. We find that all parties benefit from reduced search costs or, more precisely here, higher search frequencies.

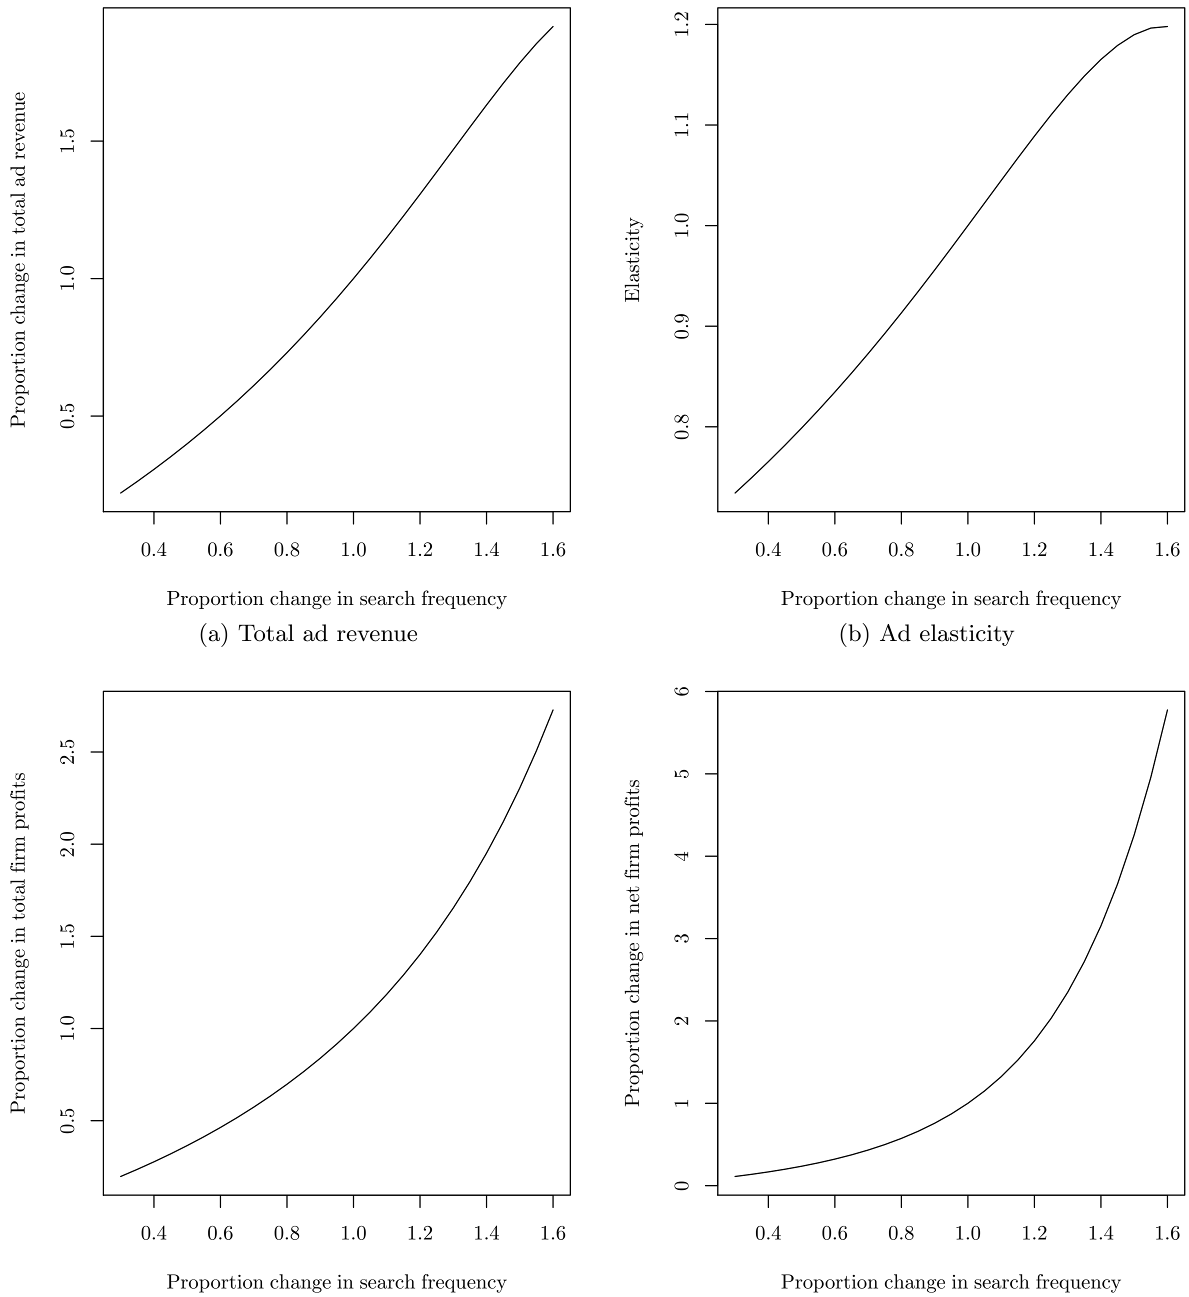

(a) Total ad revenue

(b) Ad elasticity

(c) Total firm profits

(d) Total firm net profits

Figure 7: Impact of changes in search frequencies from $s = 0.6$ on aggregates

4.3 Proportion increase in high value consumers

The ad server may be able to increase the profitability of the consumers that visit its site. It may be able to target high valuation demographics in a variety of ways, such as providing targeted content or advertising to this select group. We see how changes in $1 - F(p)$, the proportion of high value consumers, change the revenues accruing to the ad server and the firms.

Looking to Equation 8 and using the fact that $a_k r_k = \frac{D_k(p)}{q_k}$, we see that $1 - F(p)$ factors out of the sum. Hence, a proportional change in the probability of high value consumers leads to a change of the same proportion in ad revenue. Firm profits is given by $D_k(p)m_k$. Here, too, demand is directly proportional to $1 - F(p)$.

Gross firm profits and ad revenues increase by the same proportion as and thus net profits, too, grows by the same proportion. If a larger fraction of the consumers are high valuation types, a larger proportion make purchases. All parties are improved if the proportion of high value types increase.

Note that a higher fraction of consumers make purchases, reducing the CTRs for lower ranked firms. And yet these firms are better off because the clicks that they do receive are more valuable and attrition by high value consumers is attenuated.

4.4 Proportion increase in margins

A final variable to consider is the margins of the firms. The result is quite similar to that found in the preceding subsection. Equation 8 clearly shows that a proportional increase in margins (by lower costs; price remains fixed) leads to the same proportion increase in ad revenue. Gross firm revenue increases by the same proportion, implying that net revenue increases by this proportion as well. Firms and the ad server are better off. Consumers are not paying higher prices and the same fraction make purchases as before, so they are indifferent to the change.

5 Impact of dispersion of firm characteristics on bids

The lower bound of Equation 7 demonstrates that all firms may shade their bids, except for the first excluded firm. The lower bound reveals that the bid is nearly a weighted average of the value of being in slot j to firm j and the bid of firm $j + 1$. If all the other firms' expected margins are close to that of the first excluded firm and this firm bids its true value of being in the final slot on the list, the magnitude of the shading is likely reduced. We can consider how dispersion in margins and relevances across firms impacts the proportion of firm revenue that the ad server can extract through bid revenue.

5.1 Dispersion in margins

First, consider firms that all have the same relevance of 0.2, but have margins that vary. Figure 8 considers a range of variances for these margins. Margins are distributed uniformly with mean 0.5 and bounds determined by the standard deviation of the distribution of the margins. The first panel of the figure confirms our conjecture above: the less dispersion in the margins, the higher the share of firm revenue that is transferred to the ad server. As margins become more dispersed, bid shading becomes more extreme and ad revenues fall. The variation in bids relative to the variation in expected margins (here, relevance times the per-sale margin) exhibits no clear pattern and is varies little itself.

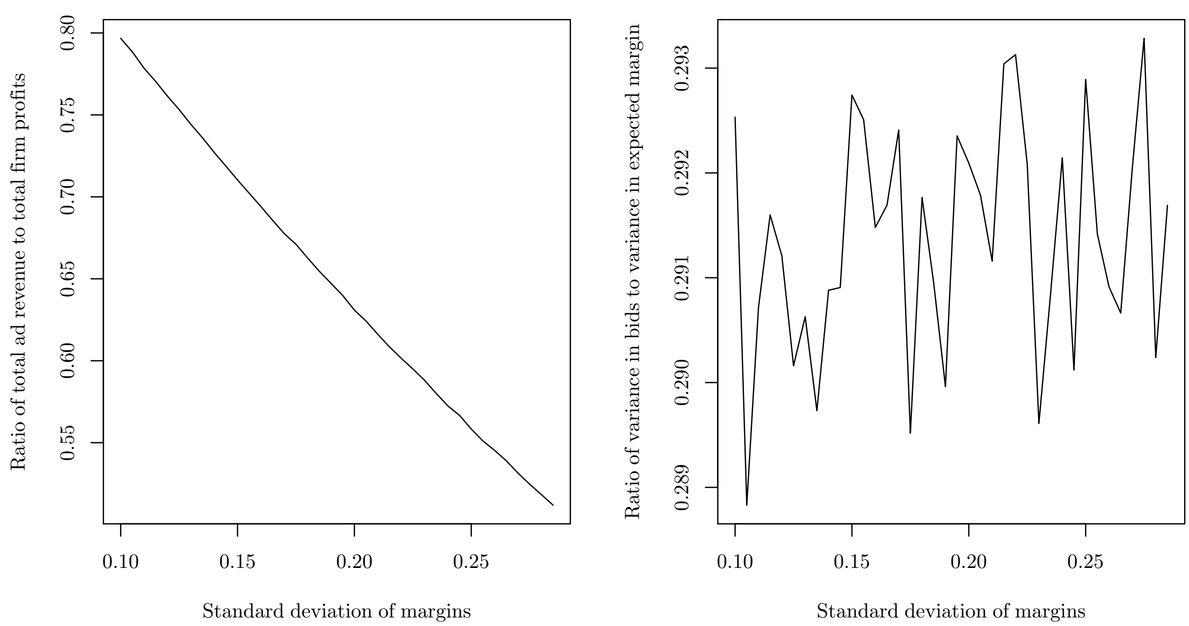

(a) Ratio of ad revenue to firm profits (b) Ratio of bid variance to variance in expected margin

Figure 8: Impact of variation in per-sale margin on auction bids and revenue

5.2 Dispersion in relevances

We might imagine that firms aim to produce products are broadly enjoyed by many consumers. The rise of the internet, however, has created new incentives and opportunities for firms to produce niche products that a relatively small segment of the population loves, while the rest does not (Bar-Isaac, Caruana and Cuñat, 2009). In our model, niche products would have lower relevances than broadly appreciated products and the advent of "long tail" niche products can lead to dispersion in the relevances in a market.

We can examine the impact of variation in the relevance of firms with constant per-sale margins of 0.5. Figure 9 shows the results of this analysis. Bid shading increases as relevances become more dispersed, just as in the case of dispersion in per-sale margins. The magnitude of this change is much smaller, however (compare the scale of the y axis in Figure 8a to that of Figure 9a). Also, the dispersion of bids relative to the dispersion of expected margins increases as relevance becomes more dispersed. This reflects the fact that shading responds by only a small amount to changes in dispersion of the relevances.

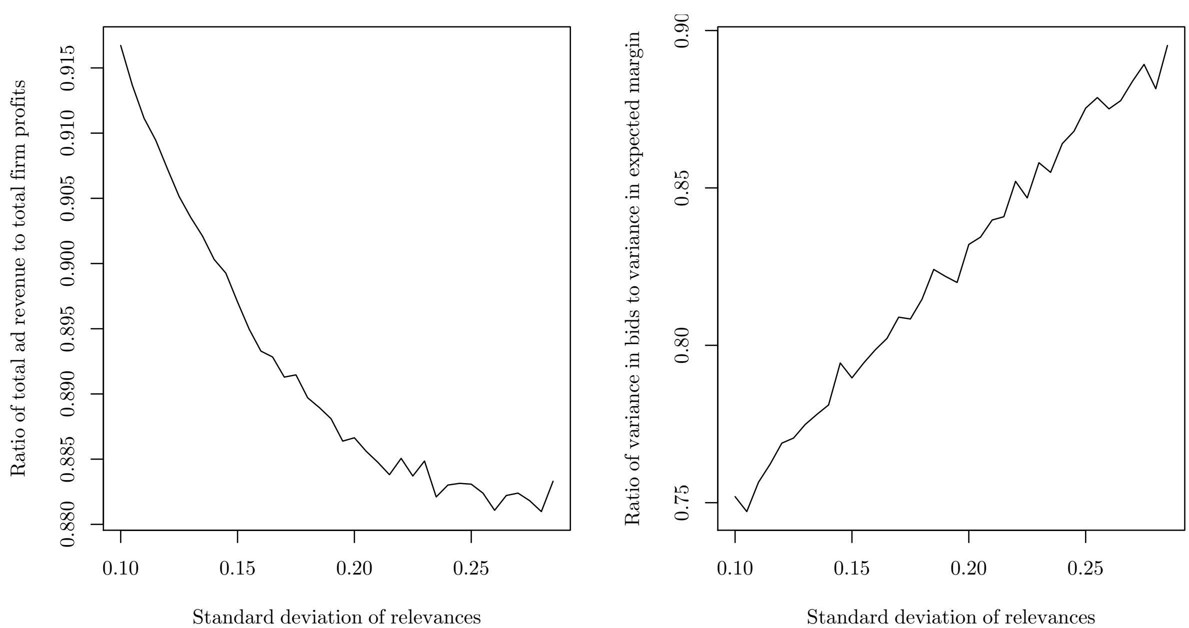

(a) Ratio of ad revenues to firm profits

(b) Ratio of bid variance to variance in expected margin

Figure 9: Impact of variation in per-sale margin on auction bids and revenue

6 Choosing the optimal number of ads to display

The ad server will choose an optimal number of ads to display. Each additional slot increases the number of clicks in the list, leading to more payments from firms, a benefit of increasing the length of the list. But the bids of all other firms change in response to the addition; indeed, they fall. The ad server chooses the optimal number of slots by balancing these factors.

Recall that the lower bound of the total ad revenue is

$$\sum_{2 \leq k \leq M+1} (k-1) m_k q_k (a_{k-1} r_{k-1} - a_k r_k).$$

When the list is M firms long, r_{M+1} is 0. Extending the list to $M+1$ firms, of course, results in a positive r_{M+1}. Plus, a new term is added to this sum equal to $(M+1) m_{M+2} q_{M+2} a_{M+1} r_{M+1}$, which uses the fact that the CTR for firm $M+2$ is 0. Hence, the change in total revenue with the addition of a new firm is

$$-M m_{M+1} q_{M+1} a_{M+1} r_{M+1} + (M+1) m_{M+2} q_{M+2} a_{M+1} r_{M+1},$$

which is negative M times the profit for firm $M+1$ when it is included on the list plus $M+1$ times the profit for firm $M+2$ if it was included in slot $M+1$ instead. This change is positive if

$$\frac{m_{M+2} q_{M+2}}{m_{M+1} q_{M+1}} \geq \frac{M}{M+1}.$$

One way to think about this condition is that the ad server chooses a set of firms so that the dispersion in total expected margins is relatively small. This accords with the findings of Section 5—dispersion is bad for ad server revenues.

Interestingly, none of the changes that we consider in Section 4 change the optimal length of the list. The length of the list does not change with proportional changes in relevances or margins and the length does not depend at all upon the proportion of low-value consumers or search costs. Hence, we might expect the length of ad listings to be relatively stable despite changes in the structure of the market. In Section 4, we considered 10 firms, each with a relevance of 0.2 and margins ranging from 0.1 to 1.0 in increments of 0.1. With these parameters, the optimal list length is not 9, as we considered, but instead would be 5.

7 Subsidizing the bid of a firm

To this point, we have considered the firms selling the advertised product to have separate interests from the ad server. Instead, suppose that the ad server also has a division that sells the product being advertised. For example, Google displays ads, but it also offers e-mail, mapping, and music services. One complaint lodged against Google is that it artificially boosts its own products to the top of advertising lists. In this section, we consider the strength of the incentives for the ad server to privilege its own products in the listing.

The obvious benefit to the ad server of its product being placed at the top of the list is that it gets additional sales. Let the ad server's product be firm j, following the labeling established throughout this paper. This firm would be in slot j earning a profit of

$$m_j D_j(p) = m_j[1 - F(p)]s_0 q_j \prod_{k=1}^{j-1} s_k(1 - q_k).$$

If, instead, firm j was placed into the first slot, it would receive a higher profit equal to

$$m_j D_1(p) = m_j[1 - F(p)]s_0 q_j.$$

The difference is

$$m_j[1 - F(p)]s_0 q_j \left[1 - \prod_{k=1}^{j-1} s_k(1 - q_k)\right].$$

This is increasing in the number of slots the firm is placed down the list j and, holding its slot fixed, the margin m_j and relevance q_j. It is also increasing in the relevances of the prededing firms. Boosting the firm to the top of the list generates more profit from sales of the product.

The cost of this maneuver is the reduction in bids by firms further down the list. The lower bound for ad revenue generated by the listing generated without privileging the ad server's firm is

$$\sum_{2 \le k \le j} (k-1) m_k q_k (a_{k-1} r_{k-1} - a_k r_k) + \sum_{j+1 \le k \le M+1} (k-2) m_k q_k (a_{k-1} r_{k-1} - a_k r_k).$$

This formulation does not include ad revenue stemming from the ad server's own listing, as we did not count this as a cost above (this explains the change to the multiple $k-2$ when summing the ad revenue generated by the bottom j firms). If the ad server's firm moved to the top of the list, the bids of the firms $j+1$ to $M+1$ would be unchanged as the firms preceding each of these would remain the same, though in a different order. Only the bids and thus the revenue generated by firms 1 to $j-1$ would change. Total revenue after firm j's move to the top is

$$\sum_{2\le k\le j}(k-2)m_k q_k(a_k^* r_k^* - a_{k+1}^* r_k^*) + \sum_{j+1\le k\le M+1}(k-2)m_k q_k(a_{k-1}r_{k-1} - a_k r_k),$$

where the asterisks indicate the out-of-equilibrium (relative to the auction that does not privilege the ad server's own firm) values of a and r under the new ranking. Changes to the first sum reveal that firm 2's bid had depended upon $a_1 r_1$ and $a_2 r_2$, but now firm 2 goes to slot 3 and thus its bid depends upon $a_2 r_2$ and $a_3 r_3$, as an example. Also, as the new top firm no longer generates revenue, $k-1$ becomes $k-2$.

The change in ad revenue, using the definitions of the parameters, is

$$\begin{aligned}
&\sum_{2\le k\le j}(k-2)m_k q_k(a_k^* r_k^* - a_{k+1}^* r_k^* - a_{k-1}r_{k-1} - a_k r_k) - \sum_{2\le k\le j} m_k q_k(a_{k-1}r_{k-1} - a_k r_k)\\
&= \sum_{2\le k\le j}(k-2)m_k q_k(1-F(p))s_0\left[\prod_{l=1}^{k-2} s_l(1-q_l)\right]\Big(\big(s_{k-1}(1-q_j)-1\big) + s_{k-1}(1-q_{k-1})\big(s_k(1-q_j)-1\big)\Big)\\
&\quad - \sum_{2\le k\le j} m_k q_k(a_{k-1}r_{k-1} - a_k r_k).
\end{aligned}$$

This change is clearly negative. Holding the other values fixed, this change is decreasing in the initial slot assignment to the ad server's firm, namely, j. Holding that slot assignment fixed, it is decreasing in q_j; *i.e.*, it is becoming costlier. This is because the ad server's firm in the top slot lets fewer consumers through, reducing downstream profits. It is decreasing in the full margin of the other firms $m_k q_k$; the more profitable the bumped firms, the more costly it is to move them to a lower slot. The question, then, is whether these costs are offset by higher profits from sales of the product being advertised.

Again, we can turn to an explicit example to calculate the change in profit from an ad server privileging its own firm. As in the other examples, consider search frequencies all equal to 1 and no low-value consumers. Also, fix every firm's margin to be 1 and set a range of relevances from 0.05 to 0.50 in increments of 0.05. Figure 10 shows the proportion change in total profits for the ad server from selling both ad space and the product being advertised by moving its firm to the top of the list relative to following the results of the auction without assistance given to the ad server's own firm. The x axis gives the unassisted ranking of the ad server's firm.

Increasing the rank of an already highly-ranked firm has small cost in terms of lost bids, but gains from increasing the size of the pool of consumers are also small. For low-ranked firms, the cost of lost bids becomes more substantial and begins to overwhelm gains from higher sales. Only by moving to the top a firm that was previously excluded from the list (here, firm 10) does the proportion change in profit start to increase again. This increase, too, should wain, however. An ad server has greatest incentive to subsidize its own firm by moving it to the top of the list when that firm would otherwise be neither high- or low-ranked or is just excluded from the listing.

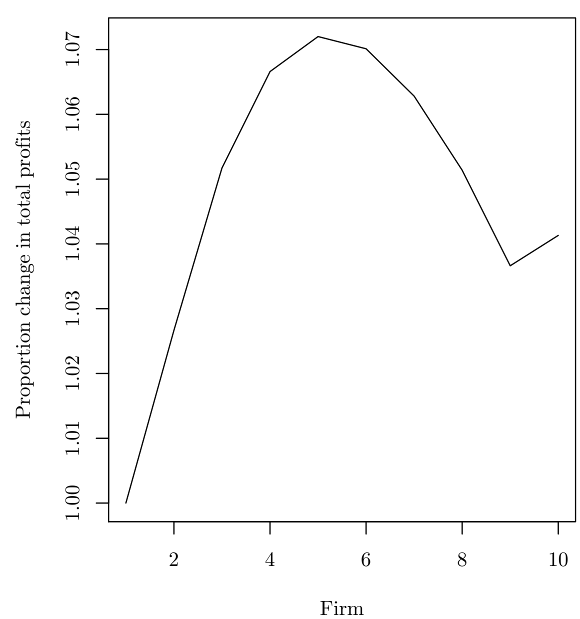

Figure 10: Change in ad server profits from both ad and product sales after privileging its own firm

Since this example uses search frequencies that are 1 all the way down the list, consumers continue to search until they find a suitable product. The margin of all firms is the same. Additionally, the relevances on the list do not change, only their order does (except when the excluded firm is added to the list and a formerly listed firm is removed). These facts together imply that total sales do not change and neither does consumer surplus; moving the favored firm to the top only changes the distribution of the revenues among the firms. This would not be the case if the search frequencies were not 1. If consumers did not search the full list, moving a less relevant firm to the top would reduce the total number of sales and consumer surplus, changing producer surplus and its distribution. If firms had different margins, then changing the order would also change the efficiency in the allocation of production. Internal subsidization of a firm by the ad server can have important implications for the sizes and distributions of consumer and producer surpluses.

Note that this analysis assumes that consumers continue searching from the first listing downward even after the ad server's firm is artificially raised to the top of the listing. We assume that consumers are not aware that they are viewing a less relevant ad and would perhaps instead be better off starting with the listing in the second position. This latter behavior would mitigate the incentive of the ad server to raise its firm to the top of the ad listing.

8 Extensions and Conclusions

The revenue raised in contextual advertising auctions has become essential to funding online content, from blogs to news to search engines. Innovation continues in this area to improve the relevance of the ads shown to consumers and to reduce search costs, each generating more product sales and higher ad revenue. While these innovations benefit consumers, ad servers may be able to manipulate the market for the products being advertised, potentially harming producers and consumers. We must understand the incentives for an ad server to effect the product market as we consider the role of competition policy in online advertising markets. To assess these issues, we incorporate consumer behavior with the bidding strategies of firms to calculate the revenue generated by contextual advertising auctions.

We begin by developing a model of consumer responses to ad listings and products being offered at the listed sites. Based upon these responses, we find the optimal firm bidding strategies. We show how these strategies depend upon per-sale margins and the probability of a consumer liking the product in question, known as the relevance of the firm. We characterize how the margin and relevance can covary while consumers still find it rational to search from the top of the advertising list downward.

Given these strategies, we consider the incentives facing the ad server. We find that it has an incentive to decrease search costs, increase firm margins (holding prices fixed), and cultivate a more valuable pool of consumers, actions that benefit itself, firms, and consumers alike. Consumers also desire improvements in match probabilities. The ad server has an incentive to develop such improvements only to a point, while firms want even less improvement in matching algorithms. The ad server seeks thick markets that generate top firms with little dispersion in margins and relevances, as this reduces the ability of firms to shade their bids. Firms, of course, desire more shading. Ad servers limit the length of the listing to mitigate shading, while firms in total, along with consumers, prefer to have more listings. Ad servers have an incentive to subsidize internal divisions that provide the product being advertised, changing the sizes and distributions of producer and consumer surpluses. We see that the preferences of consumers align with the incentives of the ad server in some cases, while they align with the incentives of firms against the ad server in others.

There are several important extensions that can be made using the model in this paper. The first would be to endogenize the search process of consumers by deriving an optimal stopping rule. Endogenized search decisions can impact the profitability of each slot and alter the incentives facing an ad server in other ways as well. Next, we might consider firms that offer the product at different prices and derive the conditions for it to remain optimal for consumers to continue searching from the top of the list to the bottom given the strategies of the firms.

Other papers have attempted to endogenize pricing decisions in the ordered search model (see, *e.g.*, Chen and He, 2006; Arbatskaya, 2007; Armstrong, Vickers and Zhou, 2009; Xu, Chen and Whinston, 2011*a*,*b*). This assumes that a firm price discriminates based upon how consumers found its product; it charges a different price to consumers that clicked on its link listed first in one ad listing than to those that clicked on its link listed third in another listing and a different price still to those consumers that visited the site directly without searching. While price dispersion has been characterized across online firms for homogeneous goods (see, *e.g.*, Baye, Morgan and Scholten, 2006), it has not been demonstrated for a single firm across means of finding the product (*e.g.*, different search engines, different ad listings, direct visits, *etc*). An interesting empirical exercise would be to look for evidence of this sort of price discrimination.

This model can also be extended to a situation where the ad server shares its profits with a content provider—Google sharing its profits from placing ads alongside an independent blog, for example. We can ask how these revenues are shared and how sharing alters the incentives for the ad server to pursue the actions discussed in this paper. Going further, we can ask how competition among ad servers to secure the space alongside independent content changes the incentives facing the ad servers and how it impacts the quality and quantity of online content provision, all important questions as the internet continues to grow.

References

Aggarwal, Gagan, Jon Feldman, S Muthukrishnan and Martin Pal. 2008. "Sponsored Search Auctions with Markovian Users." Working paper.

Anderson, Simon P. and Régis Renault. 1999. "Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlain-Diamond Model." *RAND Journal of Economics* 30(4):719–735.

Arbatskaya, Maria. 2007. "Ordered Search." *RAND Journal of Economics* 38(1):119–126.

Armstrong, Mark, John Vickers and Jidong Zhou. 2009. "Prominence and Consumer Search." *RAND Journal of Economics* 40(2):209–233.

Athey, Susan and Glenn Ellison. 2009. "Position Auctions with Consumer Search." Working paper.

Bar-Isaac, Heski, Guillermo Caruana and Vicente Cuñat. 2009. "Search, Design, and Market Structure." NET Institute Working Paper 09-17.

Baye, Michael R., John Morgan and Patrick Scholten. 2006. Information, Search, and Price Dispersion. In *Handbooks in Information Systems*, ed. Terrence Hendershott.

Chen, Yongmin and Chuan He. 2006. "Paid Placement: Advertising and Search on the Internet." NET Institute Working Paper 06-02.

Diamond, Peter A. 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3:156–168.

Edelman, Benjamin, Michael Ostrovsky and Michael Schwartz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars worth of Keywords." *American Economic Review* 97(1):242–259.

Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25(6):1163–1178.

Xu, Lizhen, Jianqing Chen and Andrew Whinston. 2011*a*. "Oligopolistic Pricing with Online Search." *Journal of Management Information Systems* 27(3):111–141.

Xu, Lizhen, Jianqing Chen and Andrew Whinston. 2011*b*. "Price Competition and Endogenous Valuation in Search Advertising." *Journal of Marketing Research* 48:566–586.

A Attrition by Low Value Consumers

In Section 2.2, we assume that consumers with valuations less than the market price p continue searching forward despite the fact that they can never find a product that is relevant and priced below their valuation. This assumption may seem odd. It does not impact the qualitative results of the paper, however.

Suppose instead that low value consumers realize that they are unable to ever find a suitable product after visiting the first site; that is, low value consumers stop searching after visiting the first site, knowing that, no matter how long they search, they will never make a purchase. This is the opposite case of Section 2.2; rather than experience attrition by high value consumers, here we find *attrition by low value consumers.*

This does not change the CTR, demand, or adjustment factor facing the first firm. For the second firm, the CTR is

$$r_2 = s_0 s_1 (1 - F(p))(1 - q_1);$$

the only consumers that continue onward are those that enter the list, have high valuations, did not find a relevant product at the first site, and continued on to the second site. Demand is

$$D_2(p) = s_0 s_1 (1 - F(p))(1 - q_1) q_2,$$

the proportion of consumers that both visit site 2 and find a relevant product. The expected margin per click is

$$\frac{m_2 D_2(p)}{r_2} = m_2 q_2.$$

Defining the adjustment factor analogously here as in Section 2.2 as the expected margin per click divided by the full margin of $m_2 q_2$ gives $a_2 = 1$. Indeed, $a_j = 1$ for all $1 < j \leq M$. Recall that the adjustment factor for firm 1 is $a_1 = 1 - F(p)$.

We can reapply the exact same results of the equilibrium analysis of Section 3.1 to this case. The main condition given by Equation 5 is that

$$(m_j q_j - m_k q_k)(a_j r_j - a_k r_k) \geq 0.$$

In that section, we concluded that the full expected margins mq were decreasing down the list because both a and r are decreasing down the list. In the where low value consumers drop out of the market after visiting the first site, the adjustment factor is increasing from site 1 to site 2: $1 - F(p)$ to 1.

Nevertheless, the product ar is decreasing. See that $a_1 r_1 = s_0(1 - F(p))$, while $a_2 r_2 = s_0(1 - F(p)) s_1 (1 - q_1)$; $a_1 r_1 > a_2 r_2$. Since $a_j = 1$ for $1 < j \leq M$, we have $a_j r_j = r_j$. The CTR is $r_j = [1 - F(p)] s_0 \prod_{k=1}^{j-1} s_k (1 - q_k)$, which is decreasing down the list, as consumers make purchases or quit searching.

The notion of an adjustment factor that adjusts the full expected margin by changes in the profitability of the consumer group visiting a particular site is flexible enough to incorporate many different behaviors of consumers. Hence, even in the case of attrition by low value consumers, the results of the auction given in Section 3 still apply and the qualitative results of the remainder of the paper remain.