The Strategy of Firms in Contextual Advertising Auctions and Incentives Facing Advertising Providers*

Charlie Gibbons†
University of California, Berkeley

October 14, 2011

Preliminary draft—comments are appreciated.

Abstract

This paper considers the strategy of online advertising providers, firms, and consumers in the context of ad listings produced by a generalized second price auction. The first part of the paper develops a model of consumer responses to ad listings and product offerings by 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. Next, we turn to incentives facing the ad server to alter the relevances and margins of the firms and the search costs and valuations of the consumer pool. We compare these incentives to the desires of firms and consumers. We also consider whether ad servers desire thick or thin product markets. These incentives have important implications for competition policy and online content provision.

1 Introduction

Advertising is essential in funding online content, from social networking sites to newspaper articles to streaming music, as well as 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. The content provider does not choose the contextual ads that appear on its site. Rather, a portion of a page delivering content is reserved for ads provided by an *ad server*. The ad server allocates slots within the space to firms using a generalized second-price auction and the revenues are shared with the content provider. In some cases, like Google, the

*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, and seminar participants at UC, Berkeley and the Institute for Humane Studies Summer Research Fellowship conference for their suggestions on this project.

†`cgibbons@econ.berkeley.edu`

ad server and content provider are one and the same. To understand how contextual advertising can supply revenues for online content, models of consumer behavior and auction theory must be combined.

We begin by formulating a model of consumer responses to contextual advertising. A stochastic element is introduced by assuming that some firms may not offer the precise variant of a product that a consumer desires; a consumer may be searching for a sweater, but may want a different style than the one offered by an advertiser. A fraction of consumers are satisfied by a firm's offering, while others continue searching down the list. This model is similar to the earliest model of consumer responses to ad listings proposed by Aggarwal et al. (2008). The probability that a consumer likes the product offered by a firm is called its relevance.

When all consumers have the same valuation, any subset of the population is equally valuable to a firm. But, when valuations vary, selection may cause some groups of the population to be more valuable than others. As consumers search down a list of contextual ads, they are not satisfied randomly. A high-value consumer purchases a product that is relevant for him, while a low-value consumer (one that has a valuation less than the price) keeps searching, despite finding a suitable product. The high-value consumers drop out and the low-value consumers fruitlessly remain, lowering the average value of the market confronting lower-listed sites. This is selection by attrition of high-value consumers. We explore how this selection effect impacts firm bidding strategies in the placement auction.

After considering these issues in consumer responses to ad listings, we turn to the auctions used to establish the orderings of the firms on the lists. 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.

We first consider market-wide competitive environments that pressure all firms to charge the same market price, though they may have different marginal costs. Under the special case the firms all have the same cost, firms sort by relevance, as in the results of Athey and Ellison (2009). Contrarily, 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 top down.

Having described equilibrium behavior of firms and consumers, we turn to the incentives facing the ad server. We consider the desire of the ad server to increase the relevance of its ads to consumers, to reduce search costs, to increase margins, and to supply a more valuable pool of consumers to advertisers. We also ask whether the ad server desires thick markets with many potential advertisers or thin markets with fewer firms. We determine if the incentives for the ad server align with those of the firms and the consumers in these regards.

This examination of consumer behavior and bidding strategies serves as a framework for future applications. Very little is known about competition in ad serving and the ability of contextual advertising to adequately fund online content. Though there are a number of papers that consider optimal auction design, few consider how consumer behavior motivates bidding. And no work relates these models to the scale and scope of online content provision and competition in ad serving. The models considered in this paper can be used to address issues in competition policy and business strategy.

2 The Goals of Contextual Advertising

The strategies of online advertisers have changed greatly over the past fifteen years. Initially, advertising online was guided by the same philosophy as that in newspapers and television. Flashy graphics aimed to grab a viewer's attention and to make him aware of a firm's products, known as brand promotion. Like newspaper ads, these ads were sold on a cost-per-impression (an "impression" is a viewing) basis, which aligned with the goal of advertisers to just be seen.

The internet provided a capability that newspapers did not: consumers could interact with an advertisement directly and could be directed to purchase a product immediately. This realization spawned the contextual advertising revolution. Rather than create awareness among a target audience, advertisers wanted consumers to find them. Advertisers sought venues where consumers were actively seeking their products. Users of search engines are actively looking for something—contextual ads could be used to help them find it.

In contextual advertising, the ads displayed are directly related to the content being viewed. Advertisers bid in generalized second price auctions for a place on a list of advertisers appearing for particular keywords. Additionally, advertisers may target consumers based upon their known demographics, location, or prior viewing habits. This matching serves to link advertisers with consumers that may actually be interested in their products.

Importantly, contextual ads provide information. In the models considered in this paper, firms are sorted in a list of several contextual ad slots in order of decreasing relevance to consumers. This sorting arises from the optimal bids of firms in the generalized second price auctions used to allocate the ads. Viewers are assumed to move down the list from top to bottom and, given this strategy, firms that are more relevant to consumers are willing to bid more to be at the top of the list.

Firms are looking for immediate, direct responses to their ads and the cost-per-impression pricing model does not reflect this goal. Firms may want a cost-per-action model, whereby a firm is only charged if someone views its ad, clicks on it, and actually makes a purchase. An example of this approach is the Amazon Associates program—content providers place links to Amazon's products on their pages and receive portion of the revenues generated via those links. This is not the most common model, however. Most contextual ads are priced on a cost-per-click basis. This is a middle ground between the model most in line with the advertiser's goals and the desire of an ad server to be paid every time that it displays an ad. Content providers share this revenue with the ad servers.

3 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.

3.1 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, that view an advertising listing containing M slots. The firms are indexed j and firms 1 through M are indexed to reflect their ranking in the advertising list; higher ranked firms appear lower on the list. Suppose that consumers consider each product being advertised sequentially starting at the top of the list.[1]

Consumers have a sort of lexicographic preferences. A product either meets the needs of consumer i, yielding a positive valuation for that product v_i, distributed with cumulative distribution function $F(v)$, 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. This differentiates a firms and provides a rationale for multiple firms to each to receive a positive market share. We refer to this probability as the *relevance* of firm j. While the relevance varies across firms, it is the same for all consumers facing a given firm.

A consumer searches by visiting site j and determining whether the product meets his needs. 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_j.[2] For completeness, define s_0 as the probability that a consumer visiting a site serving the ads looks at the ad listing at all. We begin by considering a case in which all firms charge the same price p.[3]

3.2 Product market outcomes

An important quantity of interest in online advertising is the *click-through rate* (CTR) for ad j, defined as the probability of a consumer clicking on firm j's ad. Let C_{ij} be 1 if consumer i clicks on ad j and 0 otherwise and let r_j be the CTR. The CTR for firm j is the probability that the consumer enters the list and is not satisfied by firms 1 to $j-1$. Let $z_{i1}, \ldots z_{i,j-1}$ be Bernoulli random variables that indicate whether consumer i visiting these sites finds a relevant product; they have corresponding probabilities $q_1, \ldots q_{j-1}$ of being equal to 1, which are the same for all consumers i. A second set of Bernoulli random variables, $x_{i0}, \ldots x_{i,j-1}$ indicate whether the consumer continued searching beyond a particular site (site 0 indicates browsing the ad listing at all); these are 1 with probabilities $s_1, \ldots, s_{j-1}$.

By assuming that these match and search probabilities are the same for all consumers, we

[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.

[3] To reference a competitive environment, perhaps there is Bertrand competition without fixed costs or Cournot competition among firms with the same cost function.

effectively assume that, for a consumer i that visits site j, $(x_{ij}, z_{i,j+1}, v_i)$ are mutually independent. We do not consider, for example, the case of only high-valuation consumers searching forward or consumers that are especially likely to have their needs fulfilled searching onward. Additionally, high-value consumers are no more likely to find a relevant product than a low-value consumer. Correlations among these variables may be quite interesting indeed, but are ignored in this beginning model.

In the case of the first site in the ad listing, the CTR is

$$r_1 = \Pr(C_{i1} = 1) = \Pr(x_{0i} = 1) = s_0.$$

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

- 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 make a purchase*. The CTR is

$$\begin{aligned} r_2 &= \Pr(z_{i1} = 0, x_{0i} = x_{1i} = 1) + \Pr(z_{i1} = 1, v_i < p, x_{0i} = x_{1i} = 1) \\ &= s_0 s_1 (1 - q_1) + s_0 s_1 q_1 F(p). \end{aligned}$$

We can generalize this expression to site j:

$$\begin{aligned} r_j &= \Pr(x_{01} = \ldots x_{i,j-1} = 1) \big[\Pr(z_{i1} = \ldots, z_{i,j-1} = 0) + F(p) \left[1 - \Pr(z_{i1} = \ldots, z_{i,j-1} = 0)\right] \big] \\ &= \prod_{k=0}^{j-1} s_k \left[\prod_{k=1}^{j-1} (1 - q_k) + F(p) \left[1 - \prod_{k=1}^{j-1} (1 - q_k) \right] \right]. \end{aligned} \tag{1}$$

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 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 demand for firm 1 is

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

and the demand for firm $j > 1$ is

$$\begin{aligned} D_j(p) &= \Pr(z_{i1} = \dots z_{i,j-1} = 0, z_{ij} = 1, x_{i0} = \dots x_{i,j-1} = 1, v_i \geq p) \\ &= [1 - F(p)] s_0 q_j \prod_{k=1}^{j-1} s_k (1 - q_k). \end{aligned}$$

It is also helpful to note that the elasticity of demand is

$$\frac{\partial D_j(p)}{\partial p} \frac{p}{D_j(p)} = -\frac{f(p)p}{1 - F(p)},$$

where $f(p)$ is the derivative of $F(p)$—the probability density function of the valuations at p. Notably, this elasticity does not depend upon the relevance of any firm; instead, it is a function only of the valuations of the product held by consumers and the market price.

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 demand to the CTR is decreasing as well. The demand-per-click is

$$\begin{aligned} \frac{D_j(p)}{r_j} &= \frac{\Pr(z_{i1} = \dots z_{i,j-1} = 0)[1 - F(p)]}{\Pr(z_{i1} = \dots z_{i,j-1} = 0) + F(p)\left[1 - \Pr(z_{i1} = \dots z_{i,j-1} = 0)\right]} q_j \\ &= a_j q_j. \end{aligned}$$

The expected value per click is

$$\frac{(p - c_j) D_j(p)}{r_j} = m_j a_j q_j, \tag{2}$$

where m_j is the margin for firm j and $a_j \in (0, 1]$ is a slot-specific adjustment factor. This is a ratio of the proportion of consumers that have relatively high valuations and have yet to find a relevant product to the proportion of consumers that have yet to find a relevant product and have any valuation or have found a relevant product, yet the market price is too high.

These values per click are decreasing in j (*i.e.*, the adjustment factor a_j is decreasing in j). This result arises from the fact that relatively high-value consumers make purchases and quit searching, while relatively low-value consumers continue to search down the list, never finding a product priced below their valuation. A disproportionate share of consumers that move down the list have low valuations. Thus, the fraction of clicks that turn into sales falls down the list; for a given price and cost, the expected margin from a click falls as a firm moves down the list. 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.*

3.3 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 whether the firm was in the first slot or the last. In these models, consumers are identical and 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 under attrition of high-valued consumers. Our base model does not endogenize pricing decisions, hence, we do not evaluate this strategy here.

4 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).

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.[4] 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

[4] 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.

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 incorporates the consumer behavior behind 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.

4.1 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.[5] 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.[6] The firm has strategy $b_j^* = b_j(j, b_{j+1}; q_1, \ldots, q_j, s_0, \ldots, s_{j-1})$, its bid, which is a function of the slot, its relevance and the relevances of the preceding firms, the search frequencies of the consumers, and the price that it pays per click (*i.e.*, the bid of the firm appearing one slot lower on the list).

Recall from Equation 2 that the expected value per click for firm j in slot j is $m_j a_j q_j$. In the symmetric Nash equilibria of Varian (2007), the expected profits in firm j's equilibrium slot

[5] 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.

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

must be weakly higher than those it receives in any other slot k:[7]

$$r_j(m_j a_j q_j - b_{j+1}) \geq r_k(m_j a_k q_j - b_{k+1}). \tag{3}$$

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.

Consumers are assumed to search sequentially down the list, implying that the CTR is falling down the list. Using this fact along with the equilibrium conditions for firms j and k,

$$\begin{aligned} m_j q_j(a_j r_j - a_k r_k) &\geq r_j b_k - r_k b_{k+1} \\ -m_k q_k(a_j r_j - a_k r_k) &\geq -r_j b_k + r_k b_{k+1}. \end{aligned}$$

Adding these inequalities together gives

$$(m_j q_j - m_k q_k)(a_j r_j - a_k r_k) \geq 0. \tag{4}$$

Recall that we found 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.

4.1.1 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

[7] 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 $1, \ldots, k-1$ remains 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.

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.

4.1.2 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.

4.1.3 Varying margins and relevances

Of course, the intermediate cases are 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 and low relevance firms to be at the bottom.

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 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

$$\frac{p-c}{q_j+q_k} \geq \alpha \tag{5}$$

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 5, 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.

4.2 Deriving equilibrium bids

To find the bids chosen by the firms, we return to Equation 3. 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[8]

$$m_j q_j \left(a_{j-1} - \frac{r_j}{r_{j-1}} a_j \right) + \frac{r_j}{r_{j-1}} b_{j+1} \leq b_j \leq m_{j-1} q_{j-1} \left(a_{j-1} - \frac{r_j}{r_{j-1}} a_j \right) + \frac{r_j}{r_{j-1}} b_{j+1}. \tag{6}$$

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

$$\frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} m_k q_k (a_{k-1} r_{k-1} - a_k r_k) \leq b_j \leq \frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} m_{k-1} q_{k-1} (a_{k-1} r_{k-1} - a_k r_k). \tag{7}$$

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

4.3 Discussion

Firms have positive profits. To see this, return to Equation 3 and set $k = M + 1$, the first firm not listed. Here, $r_k = 0$, implying that $q_j \geq b_{j+1}$. 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 7 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 7 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

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

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.

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.

5 Incentives of ad servers to change search structure

There are three quantities that define the consumer side of the market: the tastes of consumers are given by the match probabilities q, search costs are implicit in the search frequencies s, and the valuations for the product v. A four quantity in the model describing 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 7 reveals that the lower bound for the advertising revenue generated by firm $j-1$ would be

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

The increase in revenue from firm $j-1$ is

$$\sum_{j\leq k\leq M+1}\Delta(m_kq_k)\big[(a_{k-1}r_{k-1}-a_kr_k)+\Delta(a_{k-1}r_{k-1}-a_kr_k)\big]+m_kq_k\Delta(a_{k-1}r_{k-1}-a_kr_k); \qquad (8)$$

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. 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.

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

$$[1 - F(p)]s_0 \prod_{p=1}^{k-2} s_p(1 - q_p)\,[1 - s_{k-1}(1 - q_{k-1})]\,. \tag{9}$$

5.1 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.

5.1.1 Intuition from the model

Consider this change in the context of Equation 8. Since q_k goes up, $\Delta(m_k q_k)$ is positive and the first component of the sum is positive. For expository purposes, let all firms have the same relevance q. Then, Equation 9 becomes

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

Taking the derivative with respect to q yields

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

The quantity in question and thus ad revenue is increasing if[9]

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

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

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

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 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.[10] Equation 9, thought of more simply, is the difference in CTRs between firms $j-1$ and j.[11] 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.

5.1.2 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 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 at 0.2. We consider an increase in this value by 20%.

Figure 1 gives the impact of this change on ad revenues, bids, and gross and net (of advertising costs) firm revenue. 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 increase 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%.

[10] 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.

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

Firm revenues increase for the first 4 firms, but fall for the remainder; 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 revenue after the increase in relevances. Total firm net revenue across all the firms actually fell by 2.2%.

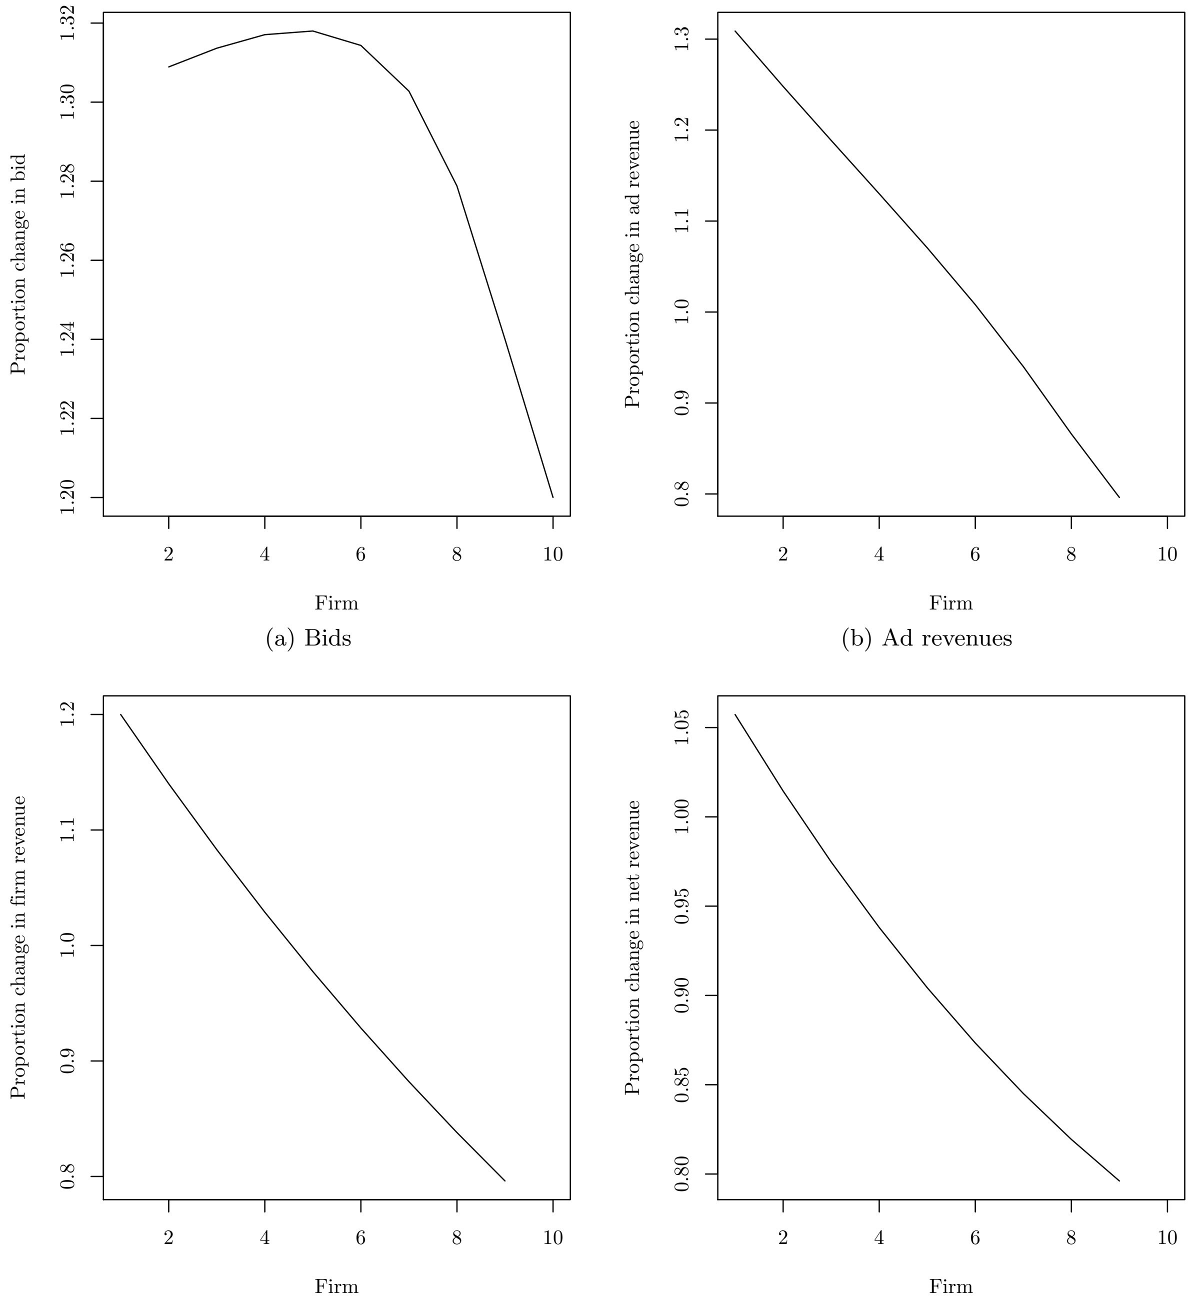

(a) Bids

(b) Ad revenues

(c) Firm revenues

(d) Firm net revenues

Figure 1: 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 2 show the total ad revenue, ad elasticity, and total firm gross and net revenues across changes in the base relevance of 0.2 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.[12]

Revenues for both the ad server and the firms are 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 revenue is maximized, at a relevance of 0.19. At a relevance less than 0.19, the ad server and the firms benefit from increases in q. Between 0.19 and 0.24, the ad server benefits from increases, while firms lose net revenue. Above 0.24, all parties are hurt by increases in relevance.[13] These plots reinforce that firms in aggregate may be hurt by overall increases in match probabilities.

Next, in Figure 3, 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 3a, we see how sales change by firm in the context of the analysis leading to Figure 1. 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 3b, an analogue to the panels of Figure 2. Consumers are unambiguously better off if the ad server increases the relevances.

[12] 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.

[13] Analysis of the benefits to the ad server of increases in q rely on assumptions about the cost of increasing q. The analysis here is predicated on the cost structure described in the preceding footnote.

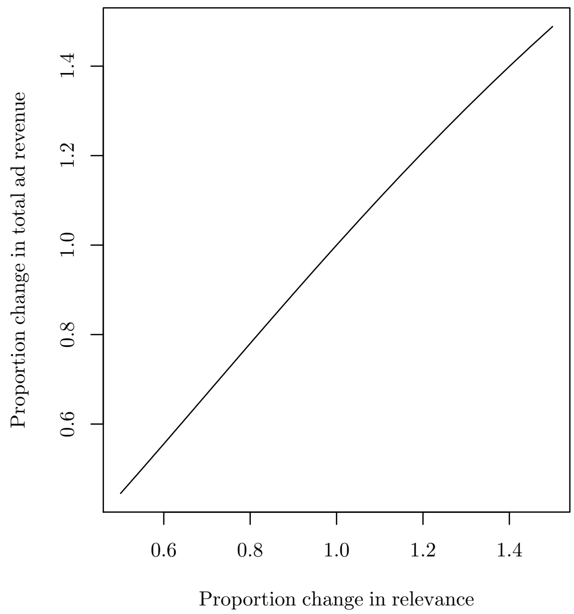

(a) Total ad revenue

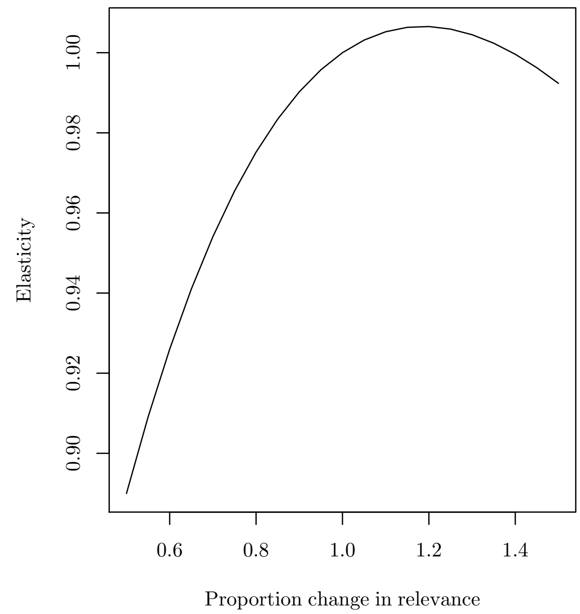

(b) Ad elasticity

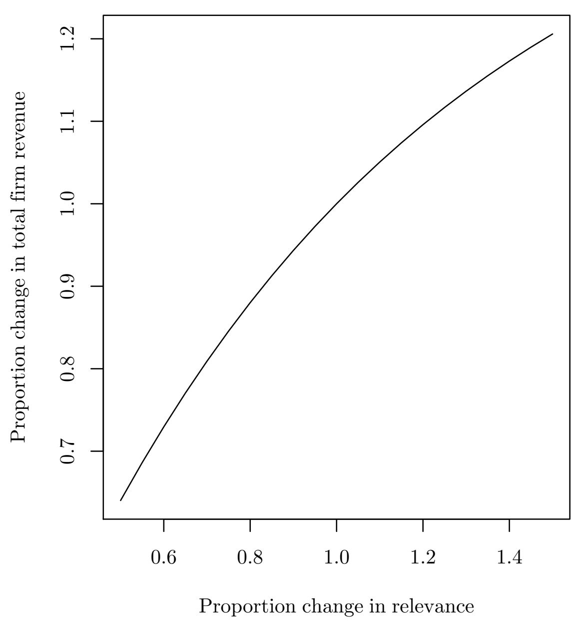

(c) Total firm revenues

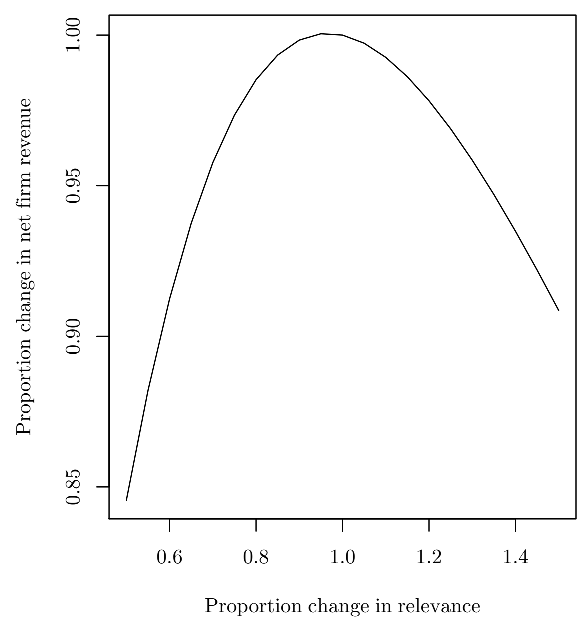

(d) Total firm net revenues

Figure 2: Impact of changes in relevance from $q = 0.2$ on aggregates

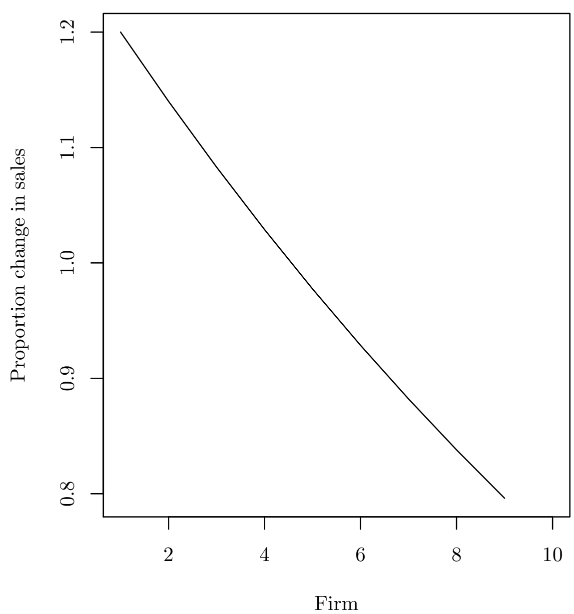

(a) Sales by firm, $q = 0.2$ increased by 20%

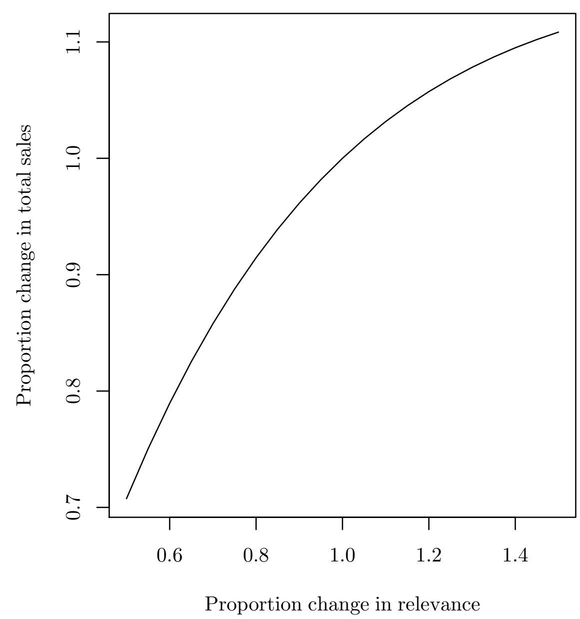

(b) Total sales for a range of proportional changes in $q = 0.2$

Figure 3: Impact of changes in relevance for consumers

5.2 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.

5.2.1 Intuition from the model

Again, return to Equation 8. The full margin $m_k q_k$ does not change, leaving only Equation 9 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

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

Hence, the quantity is increasing in search frequencies if

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

Like in the previous example, this only holds for firms high on the list. It is perhaps surprising that ad revenues do not unambiguously increase for all firms. This is because the cost to a firm in terms of lost sales by falling a rank on the list is reduced when more customers visit its site. The incentive to be high on the list is reduced, lowering the bids.

5.2.2 Simulation of the change

Following an analogous analysis to that for altering the relevance, we consider an example with $q = 0.2$ and here a base search frequency of 0.6 that are the same across firms. We consider increasing the search frequency frequency by 20%. The results are given in Figure 4.

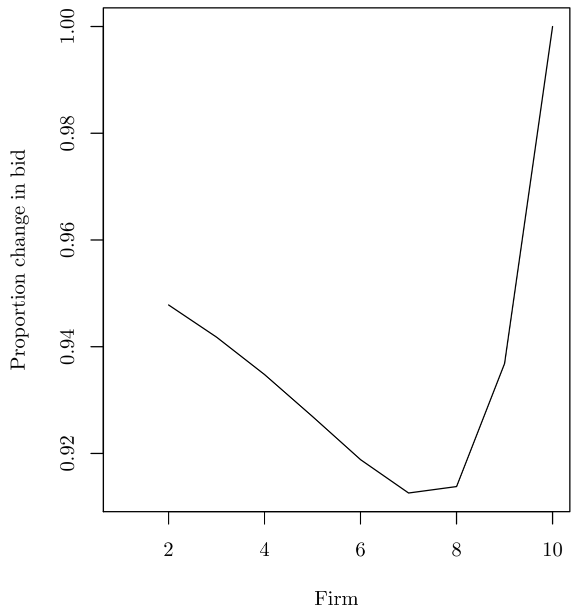

(a) Bids

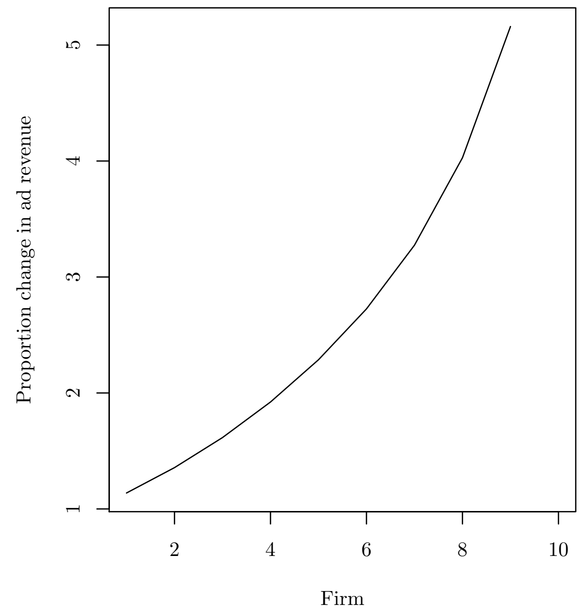

(b) Ad revenues

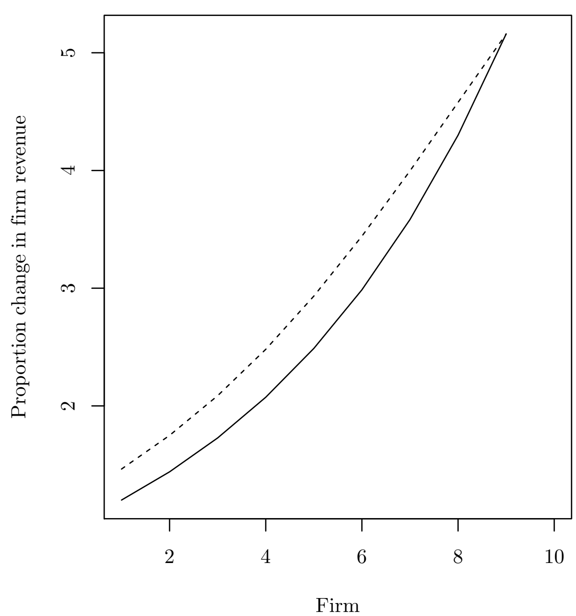

(c) Firm revenues, gross (solid), net (dashed)

Figure 4: 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 in many sites. 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 in the search frequency gets larger down the list. Firm net revenues increase by a larger percentage than gross revenue. Unlike in 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 5. 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, increasing consumer surplus. We find that all parties benefit from reduced search costs or, more precisely here, higher search frequencies.

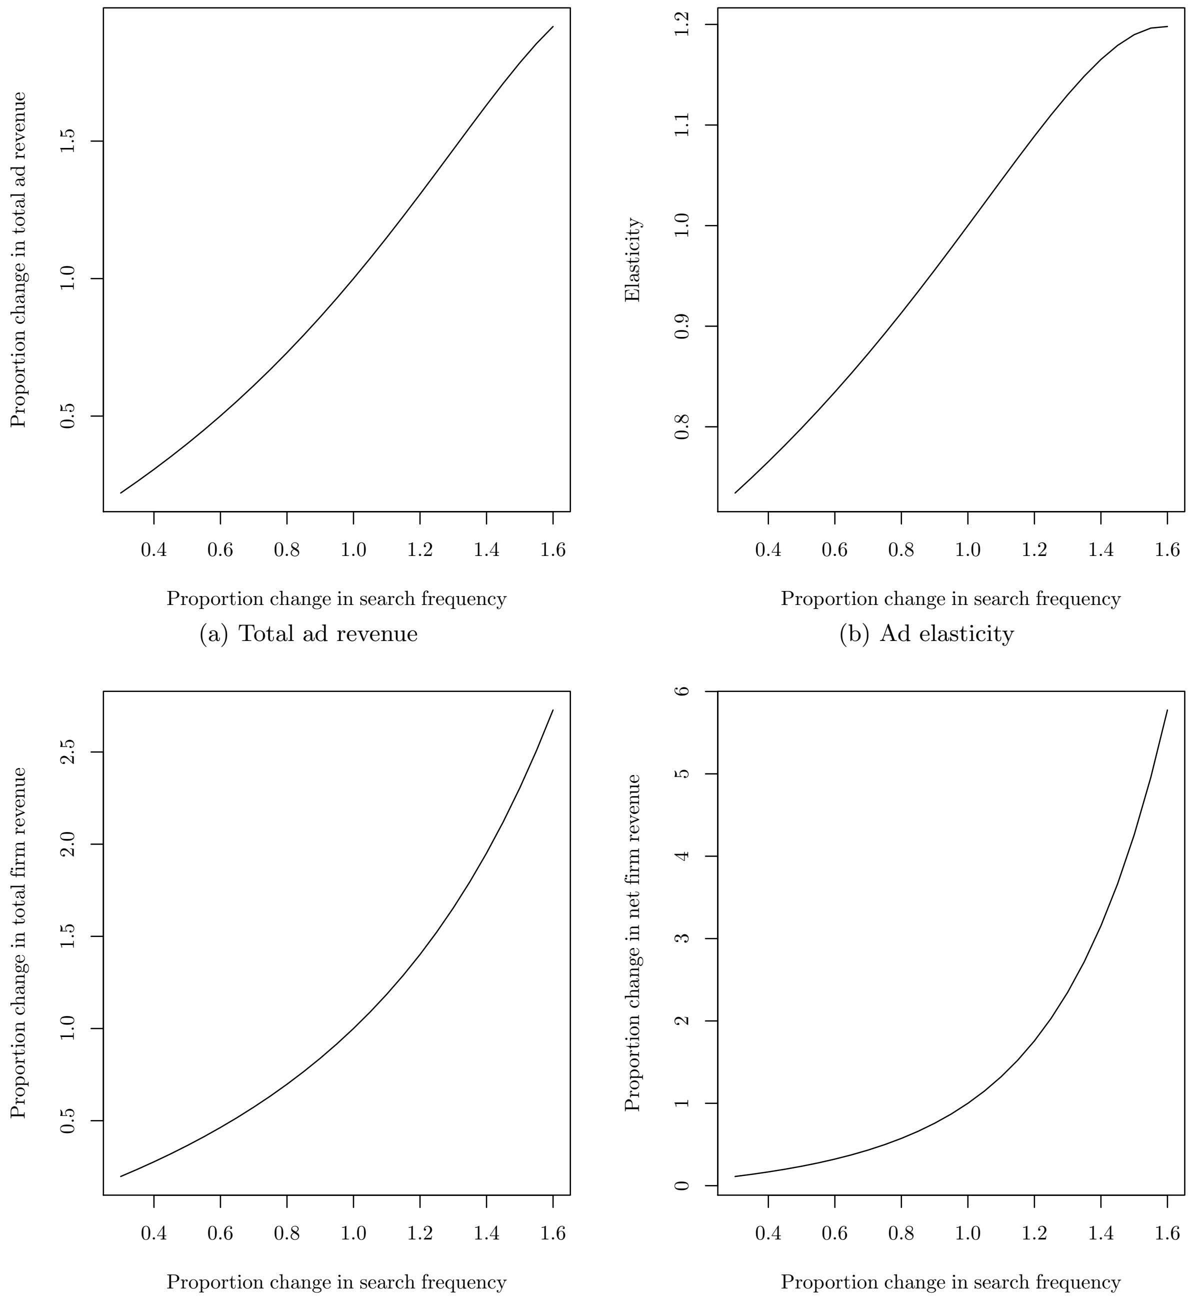

(a) Total ad revenue

(b) Ad elasticity

(c) Total firm revenues

(d) Total firm net revenues

Figure 5: Impact of changes in search frequencies from $s = 0.6$ on aggregates

5.3 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 itself. We see how changes in $1 - F(p)$ change revenues to the ad server and the firms.

Looking to Equation 7 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 revenue is given by $D_k(p)m_k$. Here, too, demand is directly proportional to $1 - F(p)$.

Gross firm revenues and ad revenues increase by the same proportion as the proportion of high value consumers and thus net revenue, 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.

5.4 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 7 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.

6 Impact of dispersion of firm characteristics on bids

The lower bound of Equation 6 demonstrates that all firms can 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 variation in margins and relevances across firms impacts the proportion of firm revenue that the ad server can extract through bid revenue.

6.1 Dispersion in margins

First, consider firms that all have the same relevance of 0.2, but have margins that vary. Figure 6 considers a range of variances for these margins. Margins are distributed uniformly with mean 0.5 with 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.

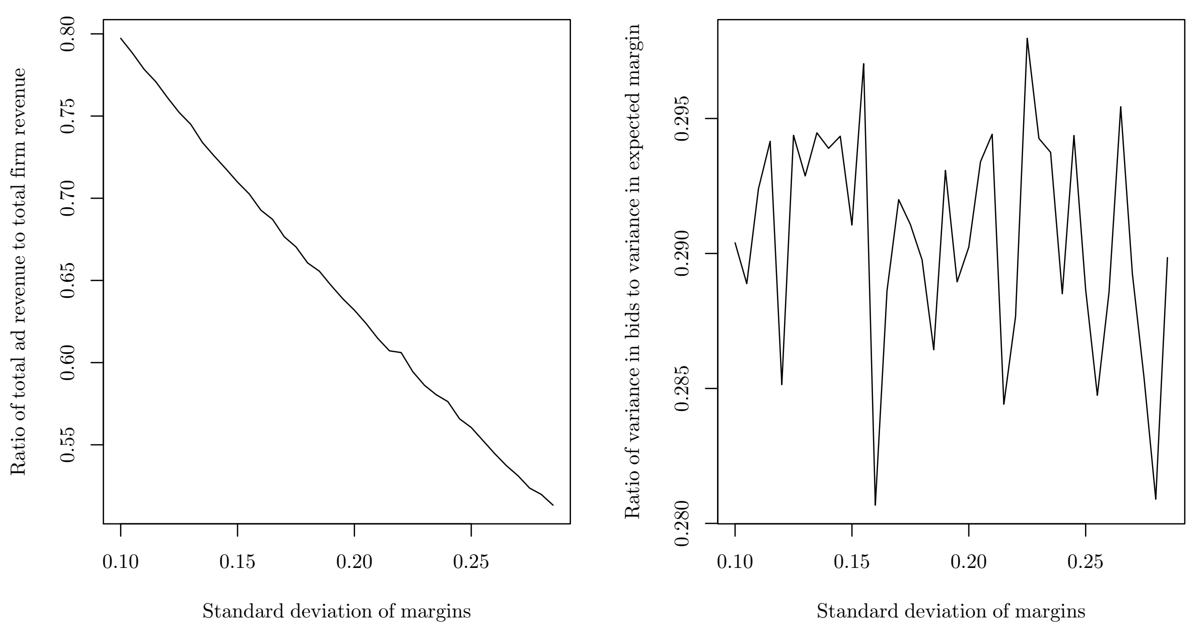

(a) Ratio of ad to firm revenues (b) Ratio of bid variance to variance in expected margin

Figure 6: Impact of variation in per-sale margin on auction bids and revenue

6.2 Dispersion in relevances

We can also examine the impact of variation in the relevance of firms with constant per-sale margins of 0.5. Figure 7 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 6a to that of Figure 7a). 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 a small amount to changes in dispersion of the relevances.

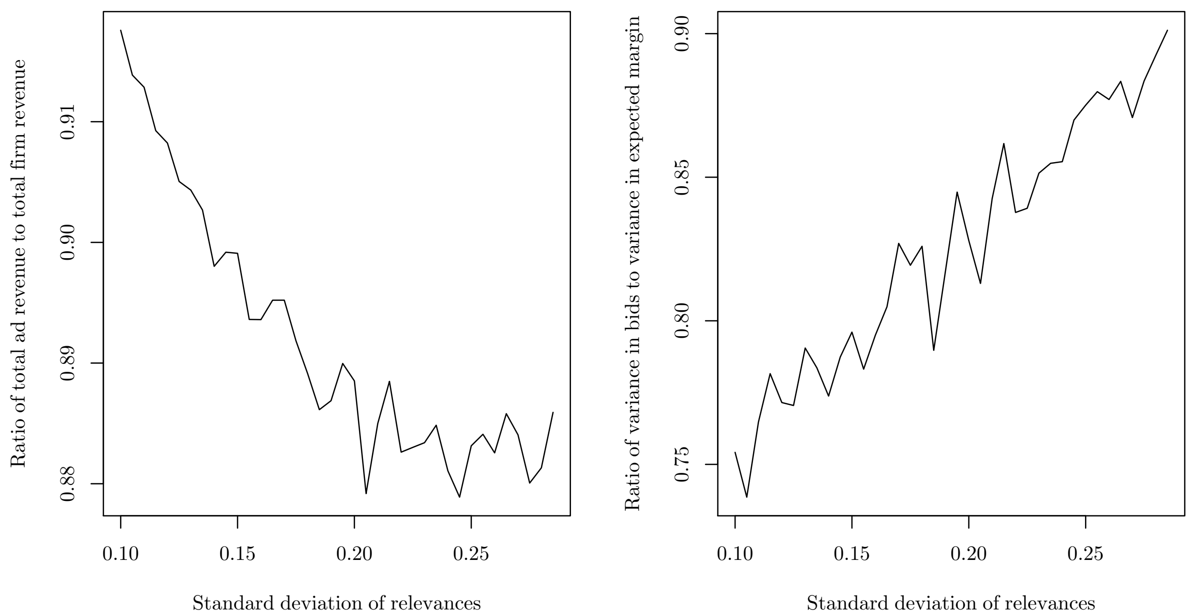

(a) Ratio of ad to firm revenues

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

Figure 7: Impact of variation in per-sale margin on auction bids and revenue

7 Conclusion

The revenue raised in contextual advertising auctions has become essential to funding online content, from blogs to news to search engines. In order to understand this business model, we need to use consumer behavior to derive the bidding strategies of firms and to examine the relationship between ad servers and content providers. It is essential that we analyze how this advertising generates its revenue to understand the provision of content. Additionally, we need to consider how this advertising impacts the markets for the products being advertised and how the markets being advertised impact advertising revenues. This paper addresses these questions.

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 maintaining the equilibrium behavior of consumers rationally searching from the top of the advertising list downward.

Given these firms' 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, while the ad server desires such improvements only to a point, and firms want even less improvement in developing 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 and thus thinner markets.

In this paper, markets have assumed to have been competitive. Ad servers may have an incentive to help themselves by subsidizing their own products in the advertising list. Even if the ad server does not have its own product in the listing, it may have an incentive to "pick winners," giving these firms monopoly profits, while the ad server extracts a large share. We might also consider whether the ad server has an incentive to exclude particular firms.

This paper can be extended to examine how revenue sharing between the ad server and the content provider fosters online content provision. How is revenue shared between these entities? How does competition in ad serving and content provision impact this share and the quality and quantity of online content generally? This paper offers a starting point for considering these issues in competition policy.

References

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

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

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

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.