

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

Charlie Gibbons^y
University of California, Berkeley

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.

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.

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

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.

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. An advertiser 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, advertisers are assigned slots in decreasing order of their bids. An advertiser pays the bid of the next ranked advertiser each time that its own ad is clicked. Many prominent papers have focused on this framework (see, *e.g.*

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 $M + 1; \dots; J$ 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 firm 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 $b_j = b_j(j; b_{j+1}; q_1; \dots; q_j; s_0; \dots; s_{j-1})$, 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

⁵We 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.

⁶Bear 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

$$r_j(m_j a_j q_j - b_{j+1}) - r_k(m_j a_k q_j - b_{k+1}) \geq 0 \quad (3)$$

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 ,

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

Adding these inequalities together gives

$$(m_j q_j - m_k q_k)(a_j r_j - a_k r_k) \geq 0 \quad (4)$$

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

⁷The CTR for slot k depends upon the relevances of the firms $1; \dots; k - 1$. As a result, the CTR for slot k is is

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, 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 $c + q_j$. We could impart a causal story: it is more or less costly to produce a product that a high proportion of people like.

Or (b)27(y)-3g]TJ 0 -2110.9091 Tf 295.659 0 ae 295.6uc25(i)1(s)7(t)1(26(c0ss2 T3kn-2110.90n-2110.90n-2110.90n-

and sort in increasing order of relevance otherwise. Note that the lefthand side of this expression

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

changed.

To find $(a_{k-1}r_{k-1} - a_k r_k)$, 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

$$[1 - F(p)]s_0 - \sum_{p=1}^{k-2} s_p(1 - q_p)[1 - s_{k-1}(1 - q_{k-1})] : \quad (9)$$

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, $(m_k q_k)$ 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

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

Taking the derivative with respect to q yields

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

The quantity in question and thus ad revenue is increasing if⁹

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

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

⁹Recall that bids are $cakre\ c8332(do))69[(1)]TJ/F72\ 10.9091\ Tf\ 6ndc81$

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 $\frac{1 - 1.2 \cdot 0.2}{1 - 0.2}$ 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%.


¹⁰One 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.

¹¹This 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 publishers' net revenues fall. This is not necessarily the case. Figure 2 shows the total ad revenue, ad revenue elasticity, and total publisher gross and net revenues across changes in the base relevance of 0.2 to 0.5. "Total" refers to measures summed across all publishers. 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 publishers are increasing with the relevance. The ad revenue elasticity and total publisher 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 publisher net revenue is maximized, at a relevance of 0.19. At a relevance less than 0.19, the ad server and the publishers benefit from increases in q . Between 0.19 and 0.24, the ad server



(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



(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 frequen-

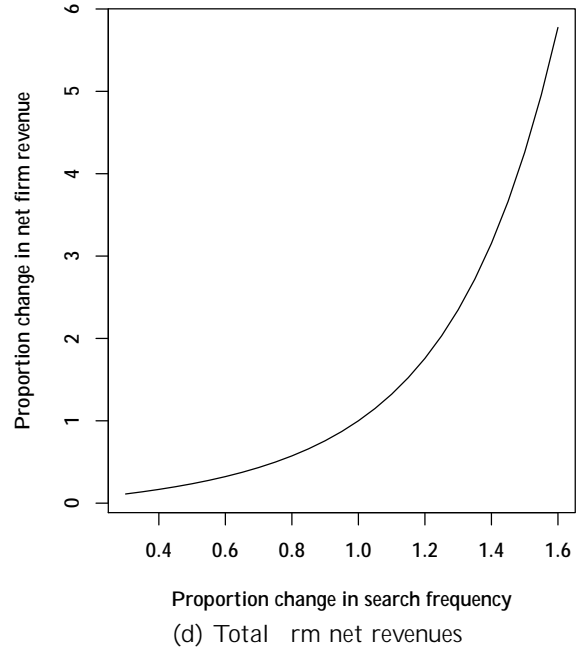





Figure 5: Impact of changes in search frequencies from $s = 0.6$ on aggregates

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

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

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.