

Firm Strategy in Contextual Advertising Auctions*

Charlie Gibbons[†]
University of California, Berkeley

May 27, 2011

PRELIMINARY DRAFT—COMMENTS ARE APPRECIATED.

Abstract

We develop models of consumer responses to online contextual advertising lists in the presence of heterogeneity in the valuations of consumers and in the costs and prices of firms. Slot-specific click-through rates, firm demand, and expected value-per-click are characterized. We explore the role of selection by attrition of high- and low-value consumers in moving down the ad listings. We incorporate these models into a generalized second price auction in which firms bid for slots in the ad listings. This endogenizes the value that the firms receive from each position and enables the finding of equilibrium bidding strategies. These strategies are then used to consider how different competitive environments impact the pricing strategies of firms in the listing and the information conveyed to consumers by the ordering of the list. This work provides a foundation for further inquiries into issues in competition policy and online content business models.

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. To understand how contextual

*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

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 base model serves as a touchstone to compare models of increasing complexity. Here, consumers all have the same valuation for a product and firms all sell the product at the same price. A stochastic element is introduced by assuming that some firms may not offer the precise variant of a product that a consumer is looking for; 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).

We add three separate sources of heterogeneity in the model: different valuations for the product and different search costs among consumers and different prices offered by firms. We also consider a model of market segmentation where ads are targeted to different groups of users. Each source is considered separately using variants of the base model of consumer behavior. Heterogeneous prices in a world of homogeneous consumers is uninteresting and varying search costs, as proposed by Athey and Ellison (2009), do not change the spirit of the base model. Heterogeneity in valuations, however, do alter the model in important ways.

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.

The opposite situation is also possible. Suppose that, after visiting a few sites and gaining an expectation of the price of the good, the low-value consumers realize that they cannot afford the product that they are searching for. They drop out of the market knowing that they would not make a purchase. Here, the average value of the market increases as the low-value consumers drop out. This is selection by attrition of low-value consumers.

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. First, a linkage is provided between the early work of Varian (2007), who models the auction as a simultaneous game of complete information, and Athey and Ellison (2009), who instead see the game as one of incomplete information and sequential moves. Then, we incorporate the variants of our base model into the Varian (2007) framework to consider how heterogeneity impacts bidding strategies. Most sources of heterogeneity do not alter the strategies greatly and the differences are quite intuitive. However, attrition by low-value consumers can prevent an equilibrium from being reached in certain circumstances.

Chen and He (2006) propose a model that incorporates heterogeneity in consumer valuations and aim to endogenize firms' pricing and bidding decisions. In this ambitious effort, they do not account for issues of attrition, issues that have important implications for both bidding and pricing strategies. We provide intuitive implications of these phenomena to pricing strategies and provide firms' pricing strategies in a two-slot listing under attrition by high-value consumers. We also consider the outcomes of ad auctions under different market-wide competitive environments.

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 costs and relevance vary, the ordering depends upon how these characteristics covary; an illustrative example is provided.

Lastly, we fully endogenize pricing by assuming that firms discriminate among consumers that arrive from an ad listing and experience attrition by high-value consumers. Firms higher on the list meet a pool of consumers with higher average valuation, inducing them to charge a higher price if consumers do search from the top of the list to the bottom. As with the case above, if firms all have the same marginal cost, firms sort in decreasing order of relevance; if firms all have the same relevance, they sort in increasing order of costs. Consumers may or may not continue to search from top to bottom in this case, as prices may be decreasing down the list, offsetting the decreases in the relevance of the products.

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 con-

textual 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 considers the implications for online content provision and the scope for competition in advertising. 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 Homogeneous Consumers: A Base Model

We have consumer $i \in 1, \dots, I$ that views a list of ads alongside a piece of content. Suppose that the consumer clicks on each ad, starting from the top and working down the list.¹ There are J firms and M slots on the ad listing. Firms 1 through M are ordered by rank in the ad listing. If a consumer visits the site of firm j , he finds a product that satisfies his need with probability q_j ; this is the *relevance* of firm j . Let q_0 be the probability that the consumer is interested in the ads at all, the relevance of the list. Note that these probabilities are all the same across individuals.

In this framework, we can calculate the *click-through rate* (CTR) for ad j 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 $z_{i1}, \dots, z_{i,j-1}$ be Bernoulli random variables that indicate whether consumer i visiting these sites is satisfied; they have corresponding probabilities q_1, \dots, q_{j-1} of being equal to 1, which are the same for all consumers i . Let z_0 be a Bernoulli random variable

¹In this model and assuming this strategy by consumers, firms are listed in decreasing order of the probability of satisfying a consumer, *i.e.*, in decreasing order of q_j . Sequential searching down the list is an optimal response to this strategy of firms.

equaling 1 if a consumer enters the list at all; this occurs with probability q_0 . Then, the CTR is

$$\begin{aligned} r_j &= \Pr(C_{ij} = 1) = \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0) \\ &= q_0 \prod_{k=1}^{j-1} (1 - q_k). \end{aligned} \tag{1}$$

We can see that the CTR is decreasing in rank j ; firms lower on the list receive fewer clicks, a result that comports with reported properties of these ad listings.

This result is similar to that of Aggarwal et al. (2008). Instead of considering the relevance of a site that is revealed *ex post*, these authors instead interpret consumers as visiting sites probabilistically from the top to the bottom of the list. To apply this motivation here, consumers decide to visit site j with probability q_j and, having visited site j , to visit another site with probability $1 - q_j$. These probabilities sum to 1, but Aggarwal et al. (2008) do not impose this restriction, allowing shoppers to browse several sites.²

In this model, we assume that all consumers have the same valuation for a product that satisfies their needs equal to 1. Supposing that there is a unit mass of consumers, demand for the product from firm j is the proportion of consumers that visit firm j 's site times the probability that a consumer is satisfied:

$$D_j = r_j q_j = q_0 \left[\prod_{k=1}^{j-1} (1 - q_k) \right] q_j.$$

Assuming that firms all price at 1, the expected value per click is the demand divided by the CTR; here, it is q_j for firm j .

4 Incorporating Heterogeneity into the Base Model

We extend the model of Section 3 by introducing heterogeneous consumers and firms. The base model already includes one form of heterogeneity—the relevance of a particular product q varies across firms. Here, we explore additional sources of variation. In Section 4.1, we consider consumers that have the same valuations for a relevant product, but the consumers are segmented by tastes

²Though Aggarwal et al. (2008) claim that their model produces decreasing CTRs, additional assumptions are necessary to generate this result if these probabilities do not sum to 1.

and firms cater to specific target markets. In Section 4.2, we show that the ability for firms to choose different prices does not alter the results of Section 3. We add search costs and Bayesian updating following the model of Athey and Ellison (2009) in Section 4.3. Lastly, heterogeneity in valuations of consumers introduces issues of selection by attrition, as shown in Section 4.4.

4.1 Market Segmentation

The first extension to the simple model that we offer is differential responses by consumers based upon the content or wording of the ads. Using our notation, we permit q_j to vary across consumers. To consider this case, let consumer i belong to group $g_i \in 1, \dots, G$; similarly, firm j belongs to group g_j in that same set. Suppose that firms can and do truthfully signal their group membership using their ad and accurate matching occurs.³ Consumers only click on ads that are relevant to them; consumers in group g click on ad j if and only if firm j also belongs to group g . As an example, an ad list corresponding to the keyword “shoes” might have ads for running shoes, dress shoes, and sneakers. A consumer in search of dress shoes does not click on ads that offer sneakers or running wear. Upon visiting a corresponding site, a consumer is satisfied with probability q_j as above. Here, z_{ij} is a Bernoulli random variable equaling 1 with probability q_j if $g_i = g_j$ and with probability 0 if $g_i \neq g_j$.

The CTR is given by

$$\begin{aligned}
 r_j &= \Pr(C_{ij}) = \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0, g_i = g_j) \\
 &= \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0 | g_i = g_j) \Pr(g_i = g_j) \\
 &= t_{g_j} q_0 \prod_{k=1}^{j-1} (1 - q_k)^{\mathbb{I}\{g_k = g_j\}} 1^{\mathbb{I}\{g_k \neq g_j\}}, \tag{2}
 \end{aligned}$$

where t_{g_j} is the proportion of consumers belonging to group the same group as firm j , g_j . This is analogous to the simple model, but the CTR here is a function of number of preceding sites of the same type, rather than the full number of higher-ranked sites. Also, it is weighted by the size of the target market (*i.e.*, the proportion of corresponding types in the population). Market demand

³We do not consider whether firms have an *incentive* to provide a perfect signal to consumers. Relevant here is the work of Ellison and Ellison (2004) on ambiguity in search descriptions.

is given by

$$D_j = r_j q_j = t_{g_j} q_0 \left[\prod_{k=1}^{j-1} (1 - q_k)^{\mathbb{I}\{g_k = g_j\}} 1^{\mathbb{I}\{g_k \neq g_j\}} \right] q_j.$$

Market demand now captures heterogeneity in the types of consumers that would visit and purchase the product from a given site. This heterogeneity, of course, cannot be described generally and may add significant complexity in considering firm strategies. Note that the expected value of a click to firm j is q_j , as in the simple model.

One hypothesis to consider is that, as keywords become more specific, the number of groups in the search population falls. For example, the search term “shoes” could have different groups of individuals seeking sneakers, running shoes, dress shoes, work boots, etc. But, the search term “running shoes” engenders less heterogeneity in the market and the market segmentation issue may be less important in this situation.

The CTR in Equation (2) is clearly decreasing within a group, but not necessarily across groups. Integrating this model into one of producer behavior may show whether the CTR is generally decreasing. The ranking that a firm chooses, however, likely depends not only upon the size of their target group, but also that group’s utility. For example, a firm with a small, but high-value audience may offer higher bids for placement than a firm with a broad base of low-value customers. If this model fails to capture a decreasing CTR, a prominent feature of the search advertising market, it may not be useful in modeling firm behavior.

4.2 Heterogeneity in Prices

A simpler case is when price varies, but consumers are homogeneous. Throughout this paper, we assume that a consumer purchases a product as soon as he is satisfied and its price is below his valuation; consumers do not continue on looking for a better deal. If a firm prices above the (sole) valuation of consumers, then it sells no goods. If a firm prices below this valuation, it loses out on profit relative to pricing precisely at the consumers’ valuation. All firms have an incentive to price at the same level. There is no price dispersion when consumers are homogeneous and the simple model of Section 3 is unchanged. Variations of the model require that consumers are heterogeneous or that consumers search multiple sites looking for better deals.

4.3 Search Costs with Bayesian Updating

The preceding models assume that consumers search costlessly and do not update their beliefs about the probability of being fulfilled by subsequent sites. Without search costs and a set of beliefs about being satisfied, these models permit infinite search by consumers. Athey and Ellison (2009) incorporate these properties into their model of consumer behavior. In this section, we outline the model of consumer behavior used in their paper.

After visiting $j - 1$ sites and remaining unfulfilled, consumer i weighs the expected probability of being fulfilled by site j , \bar{q}_j (which is also the marginal benefit for an assumed payoff of 1), against the (constant) marginal search costs of continuing to look for a satisfactory product, s_i ; consumer i visits site j if and only if $\bar{q}_j > s_i$. Search costs have a cumulative distribution H in the population. Of course, the base model is a special case of this model where H has a unit mass at 0.

Bayesian updating implies that

$$\begin{aligned}
 \bar{q}_j &= \mathbb{E}[q_j | z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0] \\
 &= \int_0^1 x \Pr(q_j = x | z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0) dx \\
 &= \int_0^1 x \frac{\Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0 | q_j = x) \Pr(q_j = x)}{\Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0)} dx \\
 &= \frac{\int_0^1 x \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0 | q_j = x) \Pr(q_j = x) dx}{\Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0)} \\
 &= \frac{\int_0^1 x \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0 | q_j = x) \Pr(q_j = x) dx}{\int_0^1 \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0 | q_j = x) \Pr(q_j = x) dx}.
 \end{aligned}$$

Now, the CTR for site $j + 1$ is

$$r_j = q_0 \left[\prod_{k=1}^{j-1} (1 - q_k) \right] H(\bar{q}_j), \quad (3)$$

giving a result analogous to Athey and Ellison (2009).⁴ The CTR is lower with search costs and Bayesian updating, but the expected value per click remains q_j for firm j .

⁴The only difference between this result and that of Athey and Ellison (2009) is that those authors assume that $q_0 = 1$; all consumers enter the ad list.

4.4 Heterogeneity in Valuations

In the three extensions considered thus far, the expected value of a click is the same as in the base model. This is because the valuations that consumers hold for a product do not vary. Instead, suppose that valuations do vary across consumers, while prices are homogeneous across firms and search costs are 0. Consumers purchase a product from firm j if they are satisfied by that good, which occurs with probability q_j as above, but also if their valuations are higher than the price offered by firm j , $p_j = p$, assumed to be constant across firms. Consumer i has valuation v_i for the product offered by a firm that satisfies his need. These valuations have a cumulative distribution F .

First, consider the demand for the firm at the top of the ad listing. The CTR is the same as the simple case: $r_1 = q_0$. The demand for product 1 by consumer i is the probability that he enters the search list and is satisfied by the product and the price of the product p is less than his valuation:

$$D_{1i}(p) = q_0 q_1 \mathbb{I}\{v_i \geq p\},$$

where $\mathbb{I}\{\cdot\}$ is an indicator function. Market demand is

$$D_1(p) = \int D_{1i} dF = q_0 q_1 [1 - F(p)].$$

The value per click is

$$\frac{D_1(p)}{r_1} = \frac{q_0 q_1 [1 - F(p)]}{q_0} = q_1 [1 - F(p)]$$

See also that the elasticity of demand is

$$\frac{\partial D(p)}{\partial p} = -q_1 f(p) \frac{p}{q_1 (1 - F(p))} = p \frac{-f(p)}{(1 - F(p))},$$

which does not depend upon q_1 . The value per click here is lower than in the simple model, which was q_1 , because there are no firms that do not purchase the product due to their low valuations of the good. The $[1 - F(p)]$ term emphasizes that not all consumers have valuations above the market

price. Now we consider whether lower-ranked firms also have lower values per click relative to the simple model.

4.4.1 No Use of Price Knowledge by Consumers

We begin by assuming that consumers do not realize that prices are the same across firms or they do not use this information in deciding to drop out of the market. Consider firm 2. 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).

Since prices are the same across firms, the consumers in the second group *never make a purchase*.

The CTR is

$$\begin{aligned} r_2 &= \Pr(z_{i0} = 1, z_{i1} = 0) + \Pr(z_{i0} = 1, z_{i1} = 1, v_i < p) \\ &= q_0(1 - q_1) + q_0q_1F(p). \end{aligned}$$

The CTR is higher than in the simple model; a site gets more clicks due to the inclusion of low-valued consumers that found a relevant, but too expensive product. But demand is lower than in the simple model:

$$\begin{aligned} D_2 &= \Pr(z_{i0} = 1, z_{i1} = 0, z_{i2} = 1, v_i \geq p) + \Pr(z_{i0} = 1, z_{i1} = 1, z_{i2} = 1, v_i < p, v_i \geq p) \\ &= q_0(1 - q_1)q_2(1 - F(p)). \end{aligned}$$

This implies that a given click is less likely to turn into a sale—clicks are less valuable. The expected value of a click is

$$\frac{D_2}{r_2} = \frac{q_0(1 - q_1)q_2(1 - F(p))}{q_0(1 - q_1) + q_0q_1F(p)} < q_2(1 - F(p)),$$

which is less than in the simple model even taking into account the fact that part of the market is never willing to pay for a good from any firm. It can be written more generally as

$$\frac{D_j}{r_j} = \frac{\Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0)(1 - F(p))}{\Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0) + F(p) [1 - \Pr(z_{i0} = 1, z_{i1} = 0, \dots, z_{i,j-1} = 0)]} q_j \quad (4)$$

Lower-ranked firms experience even lower expected values per click relative to the simple model than the top listed firm. A disproportionate share of individuals that move on to site 2 are low valuation consumers; high-value consumers leave the market once they find a relevant product, while low-value consumers do not. This is called *selection by attrition of high-valued consumers*. This effect is strengthened as rankings increase (*i.e.*, the fraction in Equation 4 increases in j). The value of a click is decreasing monotonically down the list.

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. If there is attrition by high-valued consumers, this assumption is called into question.

4.4.2 Full Use of Price Knowledge by Consumers

Here we change an assumption from the previous section: now, assume that consumers do know that prices are the same across firms and use this information in deciding whether to continue searching for a relevant product. This assumption does not change the results for the first firm in the rankings, but, after visiting this firm, all consumers know whether they can ever be satisfied—they know whether their valuations are above the market price. Even if a consumer was not satisfied, he knows whether he is willing to pay for the good conditional on being satisfied. Hence, a consumer only moves to firm 2 if he wasn't satisfied but knows that he could be.

The CTR for firm 2 is

$$r_2 = \Pr(z_{i0} = 1, z_{i1} = 0, v_i \geq p) = q_0(1 - q_1)(1 - F(p)),$$

and demand is

$$D_2(p) = \Pr(z_{i0} = 1, z_{i1} = 0, z_{i2} = 1, v_i \geq p) = q_0(1 - q_1)q_2(1 - F(p)).$$

See that no consumer that was satisfied by the first product moves on; if his valuation is above the price, he makes the purchase. If the price is above his valuation, he knows that he can never be satisfied. Additionally, all consumers that were unsatisfied, but realize that their valuations are too low to ever be satisfied also drop out of the market. Hence, all consumers with valuations less than the market price quit searching after the first firm.

The expected value of a click for site 1 is $q_1(1 - F(p))$, while it is q_2 for site 2; in fact, the expected value of a click is the same as in the simple model for all but the first firm. Though q_1 may be greater than q_2 , the fraction that buy from site 2 may actually be higher; unlike the previous case, the value of a click may not be falling monotonically down the list. The customer base for firm 2 is more valuable than that for firm 1. The top firm provides a positive informational externality to the other firms. This is *selection by attrition of low-valued consumers*.

4.5 Heterogeneity in Search Models: A Summary and Discussion

We expand on our simple model of Section 3 by adding market segmentation and heterogeneity in the prices offered by firms, valuations held by consumers, and consumer search costs. Though other papers have considered some of these issues, our analysis reveals important selection mechanisms that have not been discussed in the literature. One result is that, if prices are permitted to vary across firms, but consumers remain homogeneous, firms all choose the same price and this case is identical to our base case. Other patterns of heterogeneity produce more interesting results, however.

Market segmentation, shows that the decreasing pattern of the CTRs may not arise if different groups are targeted by different ads.⁵ Unlike the other models that we or other authors consider, the CTR is a function both of the slot in the ranking and the firm that appears there. For example, the CTR for a Gucci ad appearing in the second slot of a listing is different than the CTR that an Adidas ad would experience if it appeared in that slot instead. We imagine that the

⁵Alternatively, the CTR may be decreasing, but firms are no longer ordered in decreasing values of q .

importance of market segmentation is reduced as keywords become more refined; *e.g.*, moving from “shoes” to “running shoes.”

In this model, we maintain the assumption of homogeneous valuations; though the market may be segmented, all groups have the same distribution of valuations among consumers. There are no selection issues in this model and the expected value of a click is the same as in the base case, equal to the relevance of firm j to consumers, q_j , times the assumed price of 1. Selection issues would still be avoided if we weakened our assumption and instead maintained that all consumers *within a group* have the same valuations, as only one group of consumers interacts with a given site.

Attrition does not occur in the model of heterogeneous search costs with Bayesian updating that is suggested by Athey and Ellison (2009). Here, the CTR is lower than in the base case, as consumers with high search costs leave the market. The expected value of a click remains the same as the base case.

Attrition arises when consumers have different valuations for the goods being offered. We see that the CTR for the top-ranked firm is the same as the base case, but demand is lower because only a fraction of consumers have valuations above the market price. This implies that the expected value of a click for this firm is lower than in the base case, but this is due to the fact that the market is essentially smaller than in the base case, rather than due to attrition. We see the effects of attrition for the second-ranked firm onward. The consumers that visit these sites are not random—they self-select based upon their valuations relative to the market price. In considering these firms, we distinguish between two possibilities.

In the first case, consumers do not realize that there is a single market price or they do not use this fact in deciding whether to continue their browsing. Even after finding a satisfactory product, consumers with valuations lower than the market price futilely continue searching for a good that meets their needs and is priced low enough. The presence of these consumers yields a CTR that is higher than that of the base model, but, since they never make a purchase, the expected value of a click is lower than in the base case. Too many consumers remain in the market (*i.e.*, those that have found relevant products, but they were too expensive to purchase) and they all have low valuations. This is selection by attrition of high-valued consumers.

In the second case, consumers know that there is a single market price. After visiting the

first site, consumers quit searching if

- The product is relevant and they make a purchase,
- The product is relevant, but the market price is above their valuations, implying that they will never make a purchase, or
- The product is not relevant, but the market price is above their valuations, implying that they will never make a purchase.

This is an example of selection by attrition of low-valued consumers. The first firm provides information that consumers use to determine whether they will ever make a purchase. Since the low-valued consumers drop out, the CTR is lower for sites 2 onward relative to the base case and demand is lower by the same proportion. As a result, the expected value of a click is the same as in the base case. Recall, however, that the expected value of a click for the top firm is lower. Here, the top-ranked firm provides an information externality to lower-ranked firms that weeds low-valued consumers from the market, thereby providing a high-valued pool of customers to subsequent firms.

These cases are the two extremes: the price of one firm provides no information on the prices of other firms or it provides perfect information on the prices of others. In an updating framework with heterogeneous valuations and prices, the strength of the signal of a price falls between these two extremes. Hence, in this more complex scenario, we expect a combination of attrition by both high- and low-valued consumers. Since prices are no longer perfect signals as in the case of Section 4.4.2, the top-ranked firm no longer bears the full externality of providing the signal and the incentive to land the top spot are not deterred as strongly. This complicated model may be realistic, but it is analytically intractable.

Again, search costs with homogeneous valuations do not engender selection. But, if both search costs and valuations vary across consumers, the correlation of these factors can enhance attrition. If search costs and valuations are positively correlated, then high-value consumers are likely to leave the market for both reasons—attrition by high-valued consumers is strengthened. If they are negatively correlated, low-value consumers leave, strengthening the attrition by low-valued consumers.

The (small) literature on consumer behavior in contextual ad listings has not considered

these issues. In the model of Aggarwal et al. (2008), consumers are homogeneous and no selection issues arise. In Athey and Ellison (2009), consumers have heterogeneous search costs, but homogeneous valuations and, as we see in Section 4.3, this does not lead to selection issues.

The only other paper to consider heterogeneous valuations is Chen and He (2006). Their framework combines heterogeneous valuations with homogeneous search costs that increase with the number of sites visited and endogenous pricing decisions by firms. As we show, heterogeneous valuations generate attrition, implying that the distributions of valuations change with ranking in the ad listing. When Chen and He (2006) consider the firms' pricing decisions in their Equation 1, they do not incorporate this fact. They assume that all firms face the same pricing decision, yielding no price dispersion, but they do not consider that firms face different demand conditions depending upon their ranks and, as a result, their maximization decisions will vary. In particular, firms further down the list face fewer high-valued consumers and may be inclined to cut prices under attrition of high-valued consumers. Additionally, if their analysis was correct, if consumers know that prices are the same across firms, attrition by low-valued consumers might also apply. Clearly, these issues make the analysis much more difficult.

5 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. 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 ad is clicked. Many prominent papers have focused on this framework (see, *e.g.*, Athey and Ellison, 2009; Edelman, Ostrovsky and Schwartz, 2007; Varian, 2007).

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. Quality scores aim to estimate the expected CTR for a firm's ad.⁶ For example, for the keyword "airplane," suppose that both Boeing and a toy airplane company 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 than when a toy plane is sold. Few viewers are interested in purchasing jumbo jets, so the

⁶The quality score of a firm is no longer this transparent. Many papers consider the optimal weights to maximize auction revenue.

Boeing ad receives few clicks. 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.

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 equilibria 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 incorporates the consumer behavior behind click-through rates.

Handling the case where values per click, rather than just click-through rates, vary with the slot is more complicated. Since each slot has a different value to a firm, the bidder may want to place a different bid for each slot. This is not how ad auctions work in practice. As a result, we first consider a model of consumer behavior that does not exhibit attrition: the model of search costs and Bayesian updating found in Athey and Ellison (2009). The patterns of value-per-click given in Section 4.4 make incorporating this feature into auction games relatively straightforward, however, and we turn to these examples in Sections 5.3 and 5.4.

5.1 Position Auctions with Consumer Behavior

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, M appear on the ad list.⁷ The CTR for firms $J + 1, \dots, M$ is 0, while a firm on the list in slot j experiences a CTR r_j following Equation 3. Varian (2007) assumes that the CTR is exogenous and decreasing down the list; we provide a behavioral foundation for this assumption.

Prices are 1 and costs are 0, generating an expected value of a click equal to q_j . 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.⁸ The firm has strategy $b_j^* = b_j(j, b_{j+1}; q_j)$, its bid, which is a function of the slot, its relevance, and the price that it pays per click (*i.e.*, the bid of the firm appearing one slot lower on the list).

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 :

$$r_j(q_j - b_{j+1}) \geq r_k(q_j - b_{k+1}).$$

Analyzed another way, the benefits of a firm in slot j moving to another slot k must be weakly less than the costs. If a firm moves to a higher slot, it experiences a higher CTR, leading to higher expected sales, but pays a higher price. In moving down the list, the benefit to the firm is lower per-click costs at a cost of a lower CTR. Rewriting the equation above, we find that

$$q_j(r_k - r_j) \leq -(r_j b_{j+1} - r_k b_{k+1}). \tag{5}$$

If a firm moves up in the list relative to its equilibrium, it receives more clicks and the left-hand side is positive; counted on the right-hand side, the costs associated with this move must be larger than the benefit. If a firm moves down the list, it experiences a cost of a lower CTR (*i.e.*, the left-hand side is negative) with a benefit of lower payments to the ad intermediary—multiplying Equation 5 by -1 again shows the (positive) benefits of a move are less than the costs.

⁷We 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.

⁸First, note that firms lower on the list have higher indices—firm j is one slot above firm $j + 1$. Second, see that we use the Overture auction model here, rather than the Google model—*i.e.*, firms are ranked by bids, rather than quality-weighted bids. This comports with the expositions of Varian (2007) and Athey and Ellison (2009).

Consumers are assumed to search sequentially down the list. This strategy is validating by realizing that bidders sort themselves in decreasing order of relevances. Using the equilibrium condition,

$$\begin{aligned} q_j(r_j - r_k) &\geq r_j b_{j+1} - r_k b_{k+1} \\ -q_k(r_j - r_k) &\geq -r_j b_{j+1} + r_k b_{k+1}. \end{aligned}$$

Adding these inequalities together gives

$$(q_j - q_k)(r_j - r_k) \geq 0. \quad (6)$$

The CTR r_j is decreasing in j and this expression reveals that the relevance q_j must move in the same direction, namely, decreasing down the list, The auction is efficient in that the firms that value the slots the most get the highest slots.

Varian (2007) shows that if Equation 5 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 that⁹

$$(1 - \alpha_j)q_j + \alpha b_{j+1} \leq b_j \leq (1 - \alpha_j)q_{j-1} + \alpha b_{j+1}, \quad (7)$$

where

$$\alpha_j = \frac{r_j}{r_{j-1}} = (1 - q_{j-1}) \frac{H(\bar{q}_j)}{H(\bar{q}_{j-1})},$$

with the second equality found by using the CTR of Equation 3. These bounds can be solved recursively by recalling that $r_j = 0$ for the firms not listed, all $j \in M + 1, \dots, J$, yielding

$$\frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} q_k(r_{k-1} - r_k) \leq b_j \leq \frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} q_{k-1}(r_{k-1} - r_k)$$

⁹This procedure actually gives the bounds for b_{j+1} and appropriate reindexing gives the result shown.

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

We first note that firms have positive profits. To see this, return to Equation 5 and set $k = M + 1$, the first firm not listed. Here, $r_k = 0$, implying that $q_j \geq b_{j+1}$. Hence, the lower bound of Equation 7 is less than or equal to q_j —firms may shade their bids. It is possible that the upper bounds of Equation 7 is above q_j , 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.

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 managers can offer clients assistance with bidding strategies.

5.2 Incomplete Information Auctions

Athey and Ellison (2009) do not want to assume perfect information in their model. Instead, they choose to study an incomplete information game.¹⁰ Specifically, each firm j 's relevance, q_j , is private information. Additionally, rather than consider a simultaneous-move game, they consider a multistage game where each slot is sold separately from bottom to top. 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 7.

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.¹¹ The

¹⁰Varian (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.

¹¹Additionally, the GSP is no longer the most efficient mechanism for ranking ads Aggarwal et al. (2008). We continue to use the GSP mechanism here to reflect current practices in keyword auctions.

common values aspect is easily separated from the private value of appearing in the list and this wrinkle is not difficult to handle.

To reach a perfect Bayesian equilibrium, they find a condition such that a firm j is just indifferent to bidding the equilibrium bid b_j^* and $b_j^* + db$. The risk for firm j is that firm $j - 1$ tenders a bid within the db interval. If no other firm drops out, firm j is in slot j , experiences a CTR of r_j (which is a random variable that is dependent upon q_0, \dots, q_{j-1} , which themselves are unknown to firm j) and pays bid b_{j+1} . If another firm does drop out in the interval, firm j moves up to slot $j - 1$, pays b_j^* , firm j 's equilibrium bid (which is now given by the firm otherwise in slot $j - 1$) and has a CTR of r_{j-1} . Since firm $j - 1$ bid firm j 's equilibrium bid, it must also have a relevance of q_j . Indifference between these two outcomes for firm j for each $q_j = q$ implies that

$$\begin{aligned} \mathbb{E} [r_j(q_j - b_{j+1}) | q_j = q_{j-1} = q] &= \mathbb{E} [r_{j-1}(q_j - b_j^*) | q_j = q_{j-1} = q] \\ \mathbb{E} \left[q_0 \left[\prod_{k=1}^{j-1} (1 - q_k) \right] H(\bar{q}_j)(q_j - b_{j+1}) \middle| q_j = q_{j-1} = q \right] &= \\ \mathbb{E} \left[q_0 \left[\prod_{k=1}^{j-2} (1 - q_k) \right] H(\bar{q}_{j-1})(q_j - b_j^*) \middle| q_j = q_{j-1} = q \right] & \end{aligned}$$

The conditioning allows this condition to simplify to

$$(1 - q)H(\bar{q}_j)(q - b_{j+1}) = H(\bar{q}_{j-1})(q - b_j^*)$$

or

$$b_j^* = \left[1 - (1 - q) \frac{H(\bar{q}_j)}{H(\bar{q}_{j-1})} \right] q + (1 - q) \frac{H(\bar{q}_j)}{H(\bar{q}_{j-1})} b_{j+1}.$$

If we use the assumption that $q_{j-1} = q_j$ in the bid in the complete information game given in Equation 7, we see that the bounds shrink and reveal a unique value for the bid of firm j , the same value found here. Under a different conception of the auction game, Athey and Ellison (2009) arrive at a similar result as Varian (2007).

As above, the first excluded firm has an incentive to bid its true value. Unlike the model of

Varian (2007), for any firm included on the list, its bid is less than its true valuation—all included firms shade their bids. Solving recursively, we find that

$$b_j^* = \frac{1}{r_{j-1}} \sum_{j \leq k \leq M+1} q_k (r_{k-1} - \tilde{r}_k),$$

where

$$\tilde{r}_k = \frac{1 - q_k}{1 - q_{k-1}} r_k,$$

which reflects the fact that firm j uses $q_{j-1} = q_j$ in decision making, rather than the true value of q_{j-1} .

5.3 Position Auctions with Attrition by High-Valued Consumers

In Section 4.4.1, we see that, if consumers do not use the price of a firm to update their beliefs about the prices of other firms (which, in this example, are all the same), the high-valued consumers make purchases and leave the market, while low-value consumers fruitlessly continue searching, even after finding products that meet their liking. This is attrition by high-valued consumers. Here, we consider how this attrition changes bidding behavior in the auction.

For simplicity, we use the simultaneous-move, complete information game of Varian (2007). The value that a firm gets can be separated into the expected value per click given by the base model times a position-specific factor as shown in Equation 4. By this separability, let $m_s q_j$ be the expected value per click of firm j in slot s . Using the symmetric Nash equilibrium concept, firm j must prefer its slot to any other:

$$r_j (m_j q_j - b_{j+1}) \geq r_k (m_k q_j - b_{k+1}).$$

This equation incorporates the fact that the per-click value of being in slot j differs from that of slot k . Using the same logic as in the base case, we see that the relevances of firms are decreasing

down the list:

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

The CTR r_j is decreasing in j as is the high-value attrition factor m_j implying that q_j must be decreasing down the list as well. Hence, the auction remains efficient.

If the equilibrium inequality holds for j moving “one step” to slot $j + 1$ or to $j - 1$, then it holds for any move. For firms j and $j + 1$, equilibrium implies that

$$\begin{aligned} q_j(r_j m_j - r_{j+1} m_{j+1}) &\geq r_j b_{j+1} - r_{j+1} b_{j+2} \\ q_{j+1}(r_{j+1} m_{j+1} - r_{j+2} m_{j+2}) &\geq r_{j+1} b_{j+2} - r_{j+2} b_{j+3}. \end{aligned}$$

Adding the inequalities together and using the fact that $q_j \geq q_{j+1}$ shows that

$$q_j(r_j m_j - r_{j+2} m_{j+2}) \geq r_j b_j - r_{j+2} b_{j+3}.$$

a firm j does not want to move to slot $j + 2$ and repeated iteration implies that it would not move to any lower-ranked firm. A similar argument can be applied to show that firm j does not want to move to any higher slot on the list if it does not want to move to slot $j - 1$.

Using the “one step” solution as in Section 5.1, we find that

$$\left(\frac{m_{j-1}}{m_j} - \alpha_j\right) m_j q_j + \alpha_j b_{j+1} \leq b_j \leq \left(\frac{m_{j-1}}{m_j} - \alpha_j\right) m_j q_{j-1} + \alpha_j b_{j+1}, \quad (8)$$

where $\alpha = \frac{r_j}{r_{j-1}}$ as in the standard case. Since $\frac{m_{j-1}}{m_j} \geq 1$, more weight is placed on the expected value of being in slot j than in the base case. Taking into account that each slot is less valuable, firms are willing to bid more in the face of high-value consumer attrition. A click in slot $j - 1$ is worth $\frac{m_{j-1}}{m_j}$ times more than a click in slot j , so firms are more inclined to move up than in the base case. A higher slot means more clicks and more valuable clicks and both of these benefits offset the cost of higher bids.

5.4 Position Auctions with Attrition by Low-Value Consumers

Section 5.3 shows that attrition by high-value consumers does not change the structure of position auction equilibria a great deal. Now, we consider markets with attrition by low-value consumers using the model of Section 4.4.2.

The same framework used in Section 5.3 applies here, but we consider only two slots as a simple example. In the case of low-value attrition, $m_j = 1 - F(p)$ for $j = 1$ and $m_j = 1$ otherwise. As shown above, under the symmetric Nash equilibrium for firms 1 and 2,

$$(q_1 - q_2)(r_1 m_1 - r_2 m_2) \geq 0.$$

In that case, r , q , and m all decreased down the list. With attrition by low-valued consumers, r is decreasing, but m is weakly increasing. Using the CTR that we found in Section 4.4.2, the CTR of slot 1 is q_0 and the CTR of slot 2 is $q_0(1 - q_1)(1 - F(p))$. The inequality becomes

$$\begin{aligned} (q_1 - q_2)(q_0(1 - F(p)) - q_0(1 - q_1)(1 - F(p))) &\geq 0 \\ (q_1 - q_2)(q_0(1 - F(p))q_1) &\geq 0, \end{aligned}$$

implying that $q_1 \geq q_2$. It is straightforward to show that the rest of the list is ordered by decreasing relevance as well.

The result of Section 5.3 and the definition of m under low-value attrition reveal that the bid of the second-ranked firm is

$$((1 - F(p)) - \alpha_2) q_2 + \alpha_2 b_3 \leq b_2 \leq ((1 - F(p)) - \alpha_2) q_1 + \alpha_2 b_3. \quad (9)$$

Unlike the previous case, the weight on the expected value per click is less than the base case. The difference between the bid under low-value attrition and the base case is $F(p)q$, which is the difference in the expected value per click of being in slot 1 versus slot 2. Firm 1's lower expected value per click is entirely offset by the amount that it pays per click, the bid of firm 2.

Profits are $r_j(m_j q_j - b_{j+1})$. For firm 1, the expected value of a click $m_j q_j$ is lower than the base case by $F(p)$, but the bid of firm 2 is also lower than in the base case by $F(p)$ —profits

are the same for firm 1 in the case of low-value attrition and the base case. Even though firm 1 is providing an informational externality, it is being compensated for doing so by paying the reduced bid of firm 2. For firm 2, the CTR is $F(p)$ times that of the base case, but expected value per click and the bid of firm 3 (its true value) are the same. For firm 2, profits are lower under low-value attrition than the base case by a factor of $F(p)$ as a result of the lower CTR.

There is a complication, however. Bidders are placed into slots in decreasing order of their bids. We require the bid of firm 2 to be higher than the bid of firm 3. As in the cases considered above, firm 3, being the first unlisted firm, has an incentive to bid its true value, q_3 . For the entire range of bids to lie above that of firm 3, the following must hold:

$$\begin{aligned} ((1 - F(p)) - \alpha_2) q_2 + \alpha_2 b_3 &\geq b_3 \\ q_2 &\geq \frac{1 - \alpha}{(1 - F(p)) - \alpha} q_3. \end{aligned}$$

For any bid to be possible,

$$q_1 \geq \frac{1 - \alpha}{(1 - F(p)) - \alpha} q_3.$$

If this inequality does not hold, then the per-click profit for firm 1 would be too low and it would want to move to slot 2—there are just too many low-value consumers to make paying for those clicks worthwhile. The per-click profit for firm 2 is lower than for firm 1, so it would not want to move up to replace firm 1. If $F(p)$ is 0 (our base case), then these inequalities hold. The difference in relevances between the top two firms and the third (the first excluded firm) must get larger as a higher share of the population has valuations below the market price.

6 Competition Based upon a Wider Market

Using the tools that we develop in Sections 4 and 5, we are now prepared to consider a very simple competitive framework. We begin by considering a model with homogeneous consumer valuations. Suppose that the demand arising from any one ad listing has an insignificant effect on a firm's overall pricing strategy. For example, Nike is unlikely to alter the price that it charges for a pair of sneakers depending upon whether Reebok and New Balance appear alongside the firm in an

ad listing for the term “running shoes”; instead, the firm’s strategy is based upon a much wider market and it is unable to price discriminate among viewers of a particular ad listing.

Further, assume that the competition in the wider market results in firms charging the same price for their products.¹² Perhaps there is Bertrand competition without fixed costs or Cournot competition among firms with the same cost function. In the latter case, then the analysis of Section 5.1 applies directly and we find that firms are sorted in decreasing order of relevance, an efficient outcome.

Under Bertrand competition, firms charge the same price, but may have different costs and thus different margins. Note that these margins are firm-, rather than slot-specific. To incorporate this fact, we alter Equation 6 to be

$$m_j q_j (r_k - r_j) \leq - (r_j b_{j+1} - r_k b_{k+1}).$$

This gives an equilibrium result that

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

Now we see that expected margin is decreasing down the list, which is a combination of the relevance and the margin on a unit sold. By assumption, consumers face the same price from all firms and would prefer that firms sort based upon relevance alone, as in the previous case. To the extent that expected margin ordering differs from one based upon relevance, consumer welfare is reduced (relative to ordering by relevance).

An interesting special case is when $q_j = q$ for all j . Here, firms sort in decreasing order of margins. As the relevance of every firm is the same, a consumer is indifferent in the order of his search and we assume that he does search from top to bottom. Consumer surplus is unchanged, since all firms having the same relevance and price implies that consumers are indifferent over all orderings, but producer surplus is maximized. The opposite case, with consumers with vary relevances, but the same margin, constitutes our original model.

Of course, the intermediate cases are most interesting and most difficult to characterize.

¹²We assume that the homogeneous valuation that all consumers hold is above this market price.

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 $c + \alpha q_j$. 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

$$((p - c)(q_j - q_k) - \alpha(q_j^2 - q_k^2)) (r_j - r_k) \geq 0.$$

If the CTR is falling down the list, then firms sort in decreasing order of relevance if

$$\frac{p - c}{q_j + q_k} \geq \alpha \tag{10}$$

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 decreasing down the list. Hence, if α satisfies Equation 10, 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 margins per-sale margins have higher expected margins due to their higher relevance.

7 Competition Based upon List Placement

The scenario discussed in Section 6 assumes that firms' placements in the ad listings are irrelevant to their pricing strategies. Here, we assume that firms consider only sales made to consumers arriving via this listing in their pricing decisions. We may believe that a firm can perfectly discriminate among consumers based upon how they arrived to its site. To study this scenario, we begin by assuming that consumers do search the listing from top to bottom and solve for the firms' optimal strategies under this assumption. For the sake of simplicity, we consider the case of two firms.

7.1 Firm 1's maximization problem

The maximization problem for firm 1 is

$$\begin{aligned} \max_{p_1} (p_1 - c_1) q_1 \Pr(v_i \geq p_1) \\ \max_{p_1} (p_1 - c_1) q_1 (1 - F(p_1)). \end{aligned}$$

The first-order condition becomes

$$\begin{aligned} q_1 (1 - F(p_1)) - (p_1 - c_1) q_1 f(p_1) &= 0 \\ (p_1 - c_1) f(p_1) &= 1 - F(p_1) \end{aligned} \tag{11}$$

The cost of raising the price by one unit is the lost margin times the number of consumers that no longer buy; this is the left-hand side of this expression. The benefit is the additional unit earned on the fraction of consumers that purchase the good, given on the right.

Note that this condition does not depend upon the price that firm 2 is offering; consumers do not threaten to continue onward in search of a better deal if they are willing to pay p_1 . Additionally, it does not depend upon q_1 .

If F is the uniform distribution on $[0, 1]$, this condition becomes

$$\begin{aligned} (p_1 - c_1) &= 1 - p_1 \\ p_1 &= \frac{1 + c_1}{2}. \end{aligned} \tag{12}$$

7.2 Firm 2's maximization problem

For firm 2, this problem is slightly more difficult. This firm faces demand from two groups: those whose needs were not met by the first product and those whose needs were met, but the price was

too high. The maximization condition for firm 2 is¹³

$$\max_{p_2} (p_2 - c_2) q_2 [(1 - q_1) \Pr(v_i \geq p_2) + q_1 \Pr(p_1 > v_i \geq p_2)].$$

7.2.1 If $p_2 \leq p_1$

If $p_2 \leq p_1$, then this can be written as

$$\max_{p_2} (p_2 - c_2) q_2 [(1 - q_1) (1 - F(p_2)) + q_1 (F(p_1) - F(p_2))].$$

The first-order condition is

$$\begin{aligned} q_2 [(1 - q_1) (1 - F(p_2)) + q_1 (F(p_1) - F(p_2))] - (p_2 - c_2) q_2 f(p_2) &= 0 \\ (1 - q_1) (1 - F(p_2)) + q_1 (F(p_1) - F(p_2)) &= (p_2 - c_2) f(p_2). \end{aligned} \quad (13)$$

Here, the benefit of raising the price by one unit is on the left and the cost is on the right. As in the previous case, the optimal price is not a function of q_2 .

For the case of the uniform, this condition becomes

$$\begin{aligned} (1 - q_1) (1 - p_2) + q_1 (p_1 - p_2) &= (p_2 - c_2) \\ p_2 &= \frac{1 - q_1(1 - p_1) + c_2}{2}. \end{aligned}$$

Using the p_1 given by Equation 11, we have

$$p_2 = \frac{1 - q_1 \left(\frac{1}{2} - \frac{c_1}{2} \right) + c_2}{2}. \quad (14)$$

If both firms have the same costs, then $p_2 < p_1$; firm 2 cuts back its price because firm 1 has taken a fraction of the high-value consumers ($q_1(1 - F(p_1))$) and the average valuation of a consumer visiting site 2 is lower than that for site 1.

¹³If the firm was maximizing profit given that it was clicked on, then the term in brackets would be normalized by the probability of being clicked, $(1 - q_1) + q_1 \Pr(v_i > p_1)$. Since this is a positive constant with respect to p_2 , such a change does not impact the maximization problem. In our context, firms equivalently maximize expected profits or expected profits per click.

Recall that this maximization condition holds if $p_2 < p_1$. A necessary condition in the uniform case is that

$$c_2 \leq c_1 \left(1 - \frac{q_1}{2}\right) + \frac{q_1}{2}.$$

Note that the right-hand side is larger than c_1 , implying that a higher cost firm 2 may still price below the low cost firm that precedes it on the list. For example, if $c_1 = 0$, then $c_2 \leq \frac{q_1}{2}$. The greater q_1 , the more important consumers whose needs were met but had low valuations become in this maximization problem. Despite having higher costs than firm 1, these consumers induce firm 2 to set a lower price in order to make sales.

7.2.2 If $p_2 > p_1$

Instead, suppose that $p_2 > p_1$. Here, the maximization condition becomes

$$\max_{p_2} (p_2 - c_2) q_2 (1 - q_1) (1 - F(p_2))$$

as firm 2 does not sell to any consumer whose needs were met at site 1, but the price was too high (since, here, the price is higher still). We saw that firm 1 prices as a monopolist over the entire population; here, firm 2 prices as a monopolist over a fraction $1 - q_1$ of the population.

The first-order condition is

$$\begin{aligned} q_2 (1 - q_1) (1 - F(p_2)) - (p_2 - c_2) q_2 (1 - q_1) f(p_2) &= 0 \\ 1 - F(p_2) &= (p_2 - c_2) f(p_2). \end{aligned}$$

Again, the optimal price is not a function of q_2 . In the case of the uniform, we find

$$\begin{aligned} 1 - p_2 &= p_2 - c_2 \\ p_2 &= \frac{1 + c_2}{2}. \end{aligned} \tag{15}$$

This condition holds if $p_2 > p_1$, which requires that $c_2 > c_1$.

7.2.3 Intermediate cases

We now see that $c_1 < c_2 \leq c_1 \left(1 - \frac{q_1}{2}\right) + \frac{q_1}{2}$ are consistent with both maximization approaches for firm 2. For lower values of c_2 it is optimal to follow the approach that considers the low-value consumers coming from the first site and for higher cost, it is optimal to ignore these consumers and price for the high-value consumers alone. [[To be pursued further]]

7.3 Firms sort in decreasing order of q

Consider the expected value per click. Since all consumers visit firm 1, the expected value per click is just $(p_1 - c_1) q_1 (1 - F(p_1))$. Assume that all firms have the same cost c . Then, turning to the maximization result of Equation 11, any firm in slot 1 would charge the same price, as the optimal price is independent of q . If the expected value of a click for firm j in slot 1 is greater than that of firm j' , then

$$\begin{aligned} (p_1 - c) q_j (1 - F(p_1)) &> (p_1 - c) q_{j'} (1 - F(p_1)) \\ \implies q_j &> q_{j'}; \end{aligned}$$

the firm with the highest expected value per click of being in slot 1 is the firm with the highest relevance for consumers.

Similarly, in slot 2, the price that a firm charges in slot 2 is independent of its relevance. Given a price and relevance of the slot 1 firm and the same costs across firms, the optimal price for a firm in slot 2 is the same across firms. Then, the expected value per click for firm 2 is

$$(p_2 - c) q_2 \frac{[(1 - q_1) (1 - F(p_2)) + q_1 (F(p_1) - F(p_2))]}{(1 - q_1) + q_1 F(p_1)} = (p_2 - c) q_2 \left[1 - \frac{F(p_2)}{(1 - q_1) + q_1 F(p_1)} \right]$$

If firm j has a higher expected value per click in slot 2 than firm j' , then $q_j > q_{j'}$.

Previous results indicate the firms sort themselves in decreasing order of the expected value per click. In the case of homogeneous prices, this amounts to sorting in decreasing order of relevance. This is good for consumers, as the most likely matches would be at the top of the list. With endogenous prices, firms could sort according to relevance or price or some combination of these factors. Here we see that the price charged is independent of the firm (assuming homogeneous

costs) and is a feature of the slot instead. In this case, firms bid in decreasing order of relevance just as in the homogeneous price case.

7.4 Firms Sort in Decreasing Order of Costs

Alternatively, consider the case where all firms have the same relevance q , but different costs of production. Here, the expected value of a click is decreasing if the margin is decreasing. Since the margin is $p_j - c_j$, the margin is decreasing in cost if the derivative of p_j with respect to c_j is less than 1. By inspecting Equations 12, 14, and 15, we find this to be the case (this derivative is $1/2$ in all three cases). As a result, firms bid such that they are ordered in increasing order of cost. Prices may or may not be increasing, as determined by the results of Section 7.2.3.

8 The Optimality of Searching Down the List

Given these results, we consider whether it is optimal for consumers to visit site 1 first. The answer is obvious if, following Section 6, firms all have the same relevance and sort in increasing order of price. The answer is far less obvious for the situation of Section 7, where firms sort in order of relevance, but have decreasing prices.

Consider the case of a consumer that would purchase from either firm if that firm's product met his needs. Then it is optimal to visit site 1 first if

$$q_1(v - p_1) + (1 - q_1)[q_2(v - p_2) - s] \geq q_2(v - p_2) + (1 - q_2)[q_1(v - p_1) - s].$$

This is the expected surplus from visiting site 1 then site 2 compared to the reverse order. This inequality holds if

$$q_1 q_2 (p_2 - p_1) \geq -s(q_1 - q_2).$$

First, note that this inequality certainly holds if price is increasing down the list. In the case that price is decreasing, as found in Section 7, the lefthand side reflects the fact that the savings are only achieved for consumers satisfied by both sites, which occurs with probability $q_1 q_2$. The cost of going to site 2 first to find the savings is the increased chance of not finding a relevant product,

reflected by the righthand side (recall that firms sort in decreasing order of relevance if consumers search from top to bottom in this competitive framework and thus $q_1 - q_2$ is positive).

If consumers have a uniform distribution and the firms have the same costs c , then

$$p_2 - p_1 = -\frac{q_1}{4}(1 - c)$$

and, for the inequality to hold, the search costs must satisfy

$$s \geq \frac{q_1^2 q_2}{4(q_1 - q_2)}(1 - c).$$

9 Extensions and Conclusions

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. This paper considers the former question and provides a foundation for studying the latter.

Few researchers have considered these questions. This paper goes beyond the work that has been done in considering the impact of heterogeneity among consumers and firms on their behavior, specifically by determining click-through rates, expected values-per-click, and firm bidding strategies. We show that, in many cases, heterogeneity changes firm bidding behavior in intuitive ways. Attrition by low-value consumers, however, reduces firms' incentives in providing a signal of the market price and the top-listed firm must have relatively high relevance to consumers to be willing to provide this signal.

The models of bidding behavior and associated auction revenues can inform important debates about competition among ad servers to provide the ads that appear alongside online content. For example, do content providers have an incentive to all choose the single ad server that produces the best consumer-firm matches, *i.e.*, the firm that offers firms with the highest relevances? Does an ad server have an optimal incentive to increase the relevance of consumer-firm matches by producing a better matching algorithm? How does competition affect this incentive? How does sharing revenue with content providers impact it?

If a tendency toward monopolization by the best ad server exists, how well are other firms able to challenge the incumbent's quality? Does a monopolistic ad server have an incentive to "pick winners" among the firms advertising a product in order to get a cut of the resulting high monopoly profits? Or would it prefer a competitive market and use the auction like a tax that extracts producer surplus? Relatedly, why might an ad server want to exclude particular firms from its listing?

Lastly, the analysis of this paper can be used to better understand the Web 2.0 business model. How do ad servers and content providers split the revenue of the ad auctions? How is this impacted by competition? Would monopolization of ad serving lead to lower quality content and less content altogether?

There are many unanswered questions related to contextual advertising auctions. This paper provides a framework for approaching these questions. While most of the work in this area focuses on creating optimal auction mechanisms, no papers consider how these auctions actually impact online content provision. We offer many important concluding questions—we hope that future work will provide the answers.

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.
- Ellison, Glenn and Sara Ellison. 2004. "Search, Obfuscation, and Price Elasticities on the Internet." MIT Working Paper 04-27.
- Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25(6):1163–1178.