A Theoretical and Empirical Analysis of Alternate Auction Policies for Search Advertisements

Subramanian Balachander
sbalacha@purdue.edu
Krannert School of Management
Purdue University
West Lafayette, IN 47907
Phone: 765-494-4466
Fax: 765-494-9658

Karthik Kannan
kkarthik@purdue.edu
Krannert School of Management
Purdue University
West Lafayette, IN 47907
Phone: 765-494-3414
Fax: 765-494-9658

David G. Schwartz
dschwar@mail.biu.ac.il
Faculty of Social Sciences
School of Business Administration
Bar-Ilan University
Ramat-Gan 52900,Israel
Phone: 972-3-531-8909
Fax: 972-3-535-3182

October 2009

Acknowledgments: The authors wish to thank Namita Balachander, Alex Kim and Mohammad Rahman for their assistance with data collection.

A Theoretical and Empirical Analysis of Alternate Auction Policies for Search Advertisements

Abstract

In the online world, publishers place ads from advertisers adjacent to internet search results for a given keyword. To sell such advertising, web publishers auction multiple ad slots using a generalized second-price auction. In this paper, we compare two auction policies that publishers can use to determine the rank and payments of bidding advertisers. The first policy, the highest bid policy, ranks ads based on the bids submitted while the second policy, the highest profit policy, ranks ads based on the expected profit generated to the publisher. Interestingly, we find that the highest profit policy may generate lower publisher profits per keyword even though it uses more information. Subsequently, we use data from a search engine and empirically establish that the correlation between valuations and click through rates are positive, an important assumption in our theoretical model. This finding provides significant support for the theoretical results.

Keywords: Advertising and Media, Pricing, Internet Marketing, Game Theory, E-commerce

1 Introduction

The Internet is increasingly becoming an attractive channel for advertising. For example, in the US, spending on online ads is expected to grow 18% in 2009 while the overall advertising market is expected to decrease 6.2% (ZenithOptimedia, 2008). Further, the market for internet ads in the US, pegged at $41 billion in 2007, is expected to grow to $84 billion by 2011 (ZenithOptimedia, 2008). The beneficiaries of this growth in online advertising revenue are publishers who are either content providers such as Yahoo, NYTimes.com, or search engines such as Google and Yahoo (formerly, Overture). While internet users visit content providers such as NYTimes.com for information, they visit websites of search engines to search for the internet addresses of content providers such as NYTimes.com or for the addresses of online sellers of specific products.

Search engines, such as Google, have particularly given a new impetus to online advertising by selling search-related ads that are placed alongside search results delivered to users. In gauging the usefulness of search-related ads, advertisers have more information than an average demographic characteristic of the search engine user. Because the search engine user reveals a specific interest in the product or even the brand by entering a search keyword such as "wall clock" or "Howard Miller wall clock", he becomes a valuable prospect to sellers of the product. This advertising innovation has enabled rapid growth of search-related advertising, accounting for 40% of all online advertising in 2008 (ClickThrough, 2008). The idea of search-related advertising has been extended to placing context-related ads at the website of content providers such as NYTimes.com (Delaney and Guth, 2005). For example, an internet user reading an article on skiing at a newspaper website may be shown an ad from a vendor of ski equipment.

In search advertising, the most common form of pricing is to charge an advertiser only when a consumer clicks on the displayed ad. Such a pricing model, known as cost-per-click, capitalizes on the ease of measuring consumer response on the Internet.[1] In another innovation facilitated by the Internet medium, the cost-per-click that a particular advertiser incurs for search-related ads is usually determined by a special auction method, which Edelman, Ostrovsky, and Schwarz (2007) describe as a generalized second-price (GSP) auction. Under this model, millions of search keywords such as "wall clock" are auctioned, and, for each keyword, advertisers may submit their bids per click for the right to have their ad shown alongside search results (or adjacent to an article containing the keyword for context-related ads). Generally, multiple ads can be shown adjacent to search results, each in a different slot. These advertising slots typically differ in terms of their attractiveness to the

[1] A pricing model that is even more closely tied to performance is one where the publisher is paid only when the click results in a purchase. This model is also used in Internet advertising.

advertiser, with the more attractive slots at the top of a webpage potentially catching the website viewer's eye or getting clicked by the viewer with a higher probability. Based on a bidder ranking and payment policy (which we simply refer to as an auction policy), the publisher ranks the advertisers and then assigns the advertisers' ads to the slots such that the higher-ranked advertiser gets the more attractive slot. Similar to a second-price auction, each advertiser is charged an amount based on the bid submitted by the next lower-ranked advertiser. Because the concept of a second-price auction is extended to the entire rank order, the auction is called a "generalized" second-price auction.

In this paper, we consider a publisher who uses a GSP auction to sell online search-related ad spaces, and analyze two alternate auction policies that have been employed in the online search advertising industry for ranking a bidder (advertiser) and determining his payment. In one policy, which has been used by some search engines, including Kanoodle, Mamma, and, until recently, Yahoo, the advertisers are ranked by their submitted bids per click and assigned to the ad slots in order of their attractiveness such that the highest-ranked bidder (highest bid) gets the most attractive ad slot. Consistent with a second-price auction, each advertiser is charged an amount per click equal to the bid submitted by the advertiser immediately below in rank. We refer to this policy as the *highest-bid-per-click policy* in the paper. On the other hand, an alternate policy computes the expected publisher profit from an advertiser by multiplying the advertiser's bid per click with the ad's click-through rate (percentage of impressions that result in a click-through), and ranks the advertisers based on the expected publisher profit. The higher-ranked (higher expected-profit generating) bidders are then assigned to the more attractive slots and each bidder is charged an amount per click that yields the same expected profit to the publisher as the next lower-ranked bidder. We refer to this alternate policy as the *highest-expected-profit* policy. Google's policy for choosing the winning advertiser for search-related ads is similar to this latter policy. Yahoo also has recently switched to the highest-expected-profit policy. In the rest of the paper, we use the shorthand notation of the *highest bid* policy and the *highest profit* policy to refer to the highest bid-per-click policy and highest expected-profit policy, respectively. A comparison of these two policies is valuable given the rapid growth of search-related advertising, and renewed interest in the search market among search engines. Although Yahoo's recent switch to a highest profit policy has resulted in two major search engine firms, Google and Yahoo, now using the highest profit policy, our analysis raises the issue of whether the highest profit policy is necessarily more profitable for search engines and for publishers in general.

While the above two policies are studied in this paper because of their use in industry, they are also appealing options to analyze from a theoretical standpoint (McAfee and McMillan, 1987). In particular, these two policies encompass the two main indicators available to the publisher concerning the value of the search keyword to the advertiser (value that is appropriable by the publisher). These indicators are the advertiser's bid for the keyword and the ad's click-through rate. Because the highest profit policy incorporates both the indicators and the highest bid policy uses only one of them (i.e., the bids), it might be expected *a priori* that the highest profit policy would be more advantageous to the publisher.[2] However, when we consider bidding competition among differentiated advertisers, we find surprisingly that the highest bid policy performs at least as well as, if not better than, the highest profit policy from a publisher's point of view. A distinguishing feature of our analysis is that it considers the effect of product-market positioning of online advertisers on their bids for online ads and the consequent implications for a publisher's auction policies.

The base model that we initially consider has a publisher or search engine auctioning a single ad slot. We assume that there are two advertisers who are bidding to advertise in this space. The competing products that they sell are related to the search keyword or the context where the ad will be placed. Given our interest in analyzing the effect of the product-market positioning of the bidders, we assume that one of these advertisers is better positioned to match the consumer using the search keyword. As indicated above, our analysis finds the surprising result that, between the two policies, the one that uses the expected publisher profit for ranking bidders and specifying their payments generates lower publisher profits, in general. We subsequently consider a generalized second price auction with multiple ad slots and show that our main result about the comparative profitability of the two policies extends to this case as well. The interesting aspect of the generalization is that the result holds not only for the total publisher profits but also for profits from every ad slot. Through a sensitivity analysis we show that a model assumption that is important to our results is that advertiser valuations and click-through rates are positively related for a consumer. This assumption is conceptually similar to the positive correlation observed between frequency (similar to click-through rates) and monetary value (similar to advertiser valuations) measures in the direct marketing literature (Bult and Wansbeek, 1995) and thus appears intuitively reasonable. Moreover, we show that the highest bid policy can be more profitable than the highest profit policy even when advertiser valuations and click-through rates are negatively related, although it is not always the case. Next, we use data from an anonymous search

[2]For example, consider two advertisers with advertiser A having a value of \$4 per click for a keyword, and whose ad has a click-through rate of 0.5% while the second advertiser B has a value of \$3 per click for the same keyword and generates a click-through rate of 0.75%. A policy that focuses on value alone may favor advertiser A while a policy that focuses on expected profit may favor advertiser B (the maximum expected profit generated by advertiser B is \$$0.0225$ which exceeds that of advertiser A at \$$0.02$).

engine and conduct an empirical study to ascertain the support for our theoretical model. This study finds that the correlation between advertiser valuation and click-through rates is positive, an assumption that is important for our theoretical results. Thus, this finding supports our theoretical results.

The rest of the paper is organized as follows. Section 2 reviews the nascent literature on the pricing of online ads and positions our contribution. Section 3 presents the basic auction model, which is analyzed in Section 4. In Section 5, we generalize the basic auction model to multiple ad slots. Section 6 reports our empirical results and Section 7 concludes the paper. Proofs and additional details are available in a separate technical appendix.

2 Literature Review

The literature on the pricing of online ads is limited in view of its fairly recent origins. A related paper that looks at the auctioning of online ad space is Edelman et al. (2007) (see also Varian, 2007, for a study of similar issues). Their focus is on understanding the generic properties of the GSP model. In particular, they show that the GSP auction does not lead to a truth-telling equilibrium unlike a standard second-price auction (Vickrey, 1961). They then characterize the bidders' payments in equilibrium. Unlike Edelman et al. (2007), we consider the implication of product-market positioning of advertisers (bidders) on their profit and that of the publisher (auctioneer). This added structure placed on the bidding by the product-market positioning of bidders allows us to compare the efficacy of the highest-bid and highest-profit auction policies. In contrast, these issues are not studied in Edelman et al. (2007).

Another paper that considers advertising auction models is Feng, Bhargava, and Pennock (2005). They use simulations to compare search engine revenues from alternate policies, including the highest bid and the highest profit policies, for choosing the winners and their payments. They conclude that the highest profit policy weakly dominates highest bid policy in terms of revenue to the search engine. In contrast, we find that search engine revenues under the highest bid policy can equal or exceed those under the highest profit policy. A few other papers consider related issues. Weber and Zheng (2005) considers use of information about an advertiser's product performance in ranking bidders. Hu (2004) considers the relative attractiveness of basing an advertiser's payment for ads on clicks versus impressions. Yao and Mela (2009) develop a dynamic structural model of consumer clicking behavior and advertiser bidding behavior and estimate it using empirical data from a search engine. Further, they use the structural model to compare alternative auction policies that the search engine could deploy, such as first-price and second-price auction. In contrast, we use an analytical model to compare the alternative auction policies of highest-bid and highest-profit, while using an empirical consumer clicking model to validate our critical model assumption. A similarity between their empirical work and ours is that we both use a discrete choice model in estimating consumer clicking behavior.

In a procurement context, Engelbrecht-Wiggans, Haruvy, and Katok (2007) compare two policies, one based solely on the price bid, and the other which also considers the "quality" of the product. While there are similarities between the issues examined in their paper and ours, because of institutional differences between their setting and ours, the insights from their model do not always hold in our context. In particular, search advertising is characterized by a GSP while they consider only a standard second-price auction. Moreover, we examine the implications of product-market positioning of advertisers on the comparison between the two auction policies while bidders' characteristics are exogenous in their paper.

While we consider pricing of online ads, Dukes and Gal-Or (2003) consider the effect of price competition in a product market on negotiated rates for offline advertising. Apart from the difference in the pricing mechanism, online ads differ from offline ads in that the former offers the consumer an opportunity to make an immediate purchase. However, there may be scope for future research using a single modeling framework that capture the idiosyncratic aspects of online and offline advertising.

There is also an extensive literature on what can be considered as online price ads in the websites of intermediaries such as price comparison sites or shopbots (Chen, Iyer, and Padmanabhan, 2002, Iyer and Pazgal, 2003, Baye and Morgan, 2005). Unlike these papers, our analysis considers ads that inform or remind about existence of a seller than about price, as is generally true for search ads. Apart from this key difference, there are other significant institutional differences between these two types of advertisements that warrant a separate analysis for each.

3 Model

3.1 Prelude to Model Development

We consider a publisher or search engine P, who sells search-based ad space at his website. In this paper, we do not consider strategic interactions between publishers, allowing us to focus on the pricing problem of a single publisher. In pricing the ad space, the publisher could potentially post a price for the space. In order to set prices in an optimal way, standard economic theory requires the publisher to know the distribution of valuations of the ad space among potential advertisers. However, this knowledge requirement is quite steep for search-related advertisements since advertisers' valuations of the ad space changes with the specific search keyword used by the website user, and the number of possible keywords runs into the millions. Thus, not knowing advertisers' valuations for these numerous keywords, it becomes difficult for the publisher to post prices and update them as advertisers' demand for keywords change. On the other hand, an auction mechanism may offer a more practical alternative to posting prices. According to McAfee and McMillan (1987, pg. 704) "the reason a monopolist chooses to sell by auction rather than, say, simply posting a price is that he does not know the bidders' valuations." This may be a reason why pricing of search-related advertising has largely embraced an auction model. Moreover, the automated environment of the Internet makes it possible to hold auctions for these numerous keywords in an effective manner. For all of the above reasons, we assume in the rest of the paper that the publisher uses an auction model rather than post prices.

Our model assumes that the publisher P uses a generalized second-price (GSP) auction, described in Section 1, to sell search-based ads at his website. This assumption allows our results to be more applicable to the search advertising industry, where the GSP auction is widely used to sell search-based ads (Edelman et al., 2007).[3] A general model for an analysis would involve M advertisers competing to get their ads placed in one of K available slots. If $K \geq M$, the advertiser assigned to the least attractive slot has no lower-ranked advertiser; in this case, the amount to be charged to the last advertiser is typically set exogenously. Modeling this case requires an additional parameter but the results are qualitatively similar to the case when $K < M$. Therefore, for simplicity, we assume $0 < K < M$ throughout the paper. For expository reasons, we initially consider a base model with two ($M = 2$) advertisers competing for one ($K = 1$) ad slot. Later, we show that our results from the base case generalize for any arbitrary K and M such that $0 < K < M$. Note that, with $K = 1$, the GSP auction reduces to a standard second-price auction.

3.2 The Base Auction Model

Consumers and Differentiated Advertisers: In the base auction model, we consider two advertisers, X and Y, who sell in a product market. The advertisers compete for their ad to be placed in the single advertising slot at the website of the publisher or search engine P. Thus, $M = 2$ and $K = 1$. There is a mass of consumers who visit the publisher's website and who are also in this product market. The ad in the slot can be dynamically changed for each consumer visiting the website. We assume that a consumer clicking an advertiser's ad at the publisher's website is taken directly through to the website of the advertiser (and hence the term click-through

[3]Edelman et al. (2007) consider generalized first price auctions in which advertisers whose ads are shown make payments equal to their bid amount. Such a mechanism was initially used by Overture before it switched to the GSP auction. Edelman et al. (2007) argue that a generalized first-price auction is unstable which may explain why it has not found significant use. More generally, the issue of alternate mechanisms for auctioning multiple advertising slots may be of interest but is beyond the scope of the current paper.

for such clicks). The consumers are of two types, i and j, and so is the case for the advertisers, with advertiser X of type i, and advertiser Y of type j. (Hereafter, we refer to the advertisers simply by their types.) Since we consider search-related advertising, we assume that the publisher P and advertisers can identify the type (i or j) of the website user (through the specific search keyword used by the user). For example, consumers interested in wall clocks may be of two types: those interested in the Howard Miller brand and those interested in the Seth Thomas brand. The consumer interested in the Howard Miller wall clock may use the Howard Miller brand name in their online search keyword at firm P's website, and thus reveal his type to firm P. In the context of this example, our advertiser assumptions imply that advertisers may either be of the 'Howard Miller' type or the 'Seth Thomas' type. An advertiser of the Howard Miller type may be interpreted as a vendor who carries a larger selection of Howard Miller wall clocks and vice versa.

Given our assumption that a consumer's type can be identified through his search keywords, we focus our analysis on only one consumer type, say i, without loss of generality. Accordingly, we now present our model assumptions for advertiser valuations and click-through rates for consumer type i and note that these parameters and segment size can be different for the other consumer type j without affecting our results. Specifically, we assume that advertisers i and j have a value of v_h and v_l respectively for a click-through from a consumer of type i, with $v_l < v_h$. This assumption implies that an advertiser attaches a higher value to a consumer of his own type. An advertiser's value for a click-through from a consumer reflects the expected payoff (such as from a purchase) from the consumer visiting the advertiser's website.[4] In the context of our above example, the advertiser of the Howard Miller type may have a higher expected payoff from a consumer of the Howard Miller type than the other advertiser because of his better selection of Howard Miller wall clocks. The difference between the two advertisers' valuations of consumer type i is a measure of the differentiation of the advertisers in the product market. A smaller difference indicates lesser differentiation.

Click-Through Rates for Ads: A consumer's click-through rate for an ad depends on the ad content and on whether the consumer is *ex ante* informed or uninformed about the advertiser's type. Note that although a consumer may be quite knowledgeable about his own type, e.g. his preferred brand of wall clock, he may not know an advertiser's type in the sense of whether the advertiser carries a good selection of his preferred brand. Let λ and $(1 - \lambda)$ represent the fraction of informed

[4] Another perspective borrowed from sales management is to view a consumer click as a lead and value it as such. This perspective not only considers the potential payoff from a consumer purchase but also considers the possibility that a consumer may be acquired using other media or channels. We believe that this alternative perspective requires modeling of the consumer's search process and as such leave this issue to future research.

and uninformed consumers, independent of the consumer type.[5]

When a consumer is informed about an advertiser's type, the online ad serves mainly as a reminder or to provide a link to the advertiser's website. It is reasonable to assume that the click-through rate for such consumers depends only on the advertiser's type. Specifically, we assume that for an informed customer of type i, the click-through rates for ads from advertisers i and j are c_h and c_l respectively, where $c_h > c_l$. This assumption is reasonable because an informed customer is more likely to click an advertiser who matches his type. Overall, our assumptions about valuations and click-through rates combine to imply that an informed consumer's click-through rate of an ad increases with the advertiser's valuation of the consumer, resulting in a positive correlation between valuations and click-through rates across advertisers.[6] Later, we show that this assumption is important for our main result as we explore the sensitivity of our results to this assumption.

Conversely, if a consumer is uninformed of the advertiser's type, his click-through rate depends only on the ad content or design. With uninformed consumers, an advertiser may design a tailored ad that appeals to a particular type (say i) by claiming that the advertiser is of that type (i). The consumer type to which the advertiser tailors his ads could even be his competitor's type. Thus, we do not preclude advertisers from making false appeals to a consumer type for which his competitor is better positioned. Consistent with the click-through rates of informed consumers, if an ad is tailored to a consumer type (say i), we assume that the ad obtains a click-through rate of c_h from that type (type i); and a click-through rate of c_l if the ad is tailored to the other consumer type (type j).[7] In the context of our running example, it seems natural to assume that an uninformed consumer interested in Howard Miller wall clocks would click ads featuring Howard Miller clocks at a higher rate than ads featuring Seth Thomas clocks. Note that our assumptions imply that if advertiser j tailors his ad to an uninformed consumer of type i, he obtains a click-through rate of c_h. However, the value from attracting that consumer remains v_l.

Sequence of Moves: The sequence of moves in our model is as follows. At the beginning of the game, the publisher P announces his choice of the auction policy

[5] Even with prior knowledge about an online seller, 41% of web users used a search engine in 2003 to locate the seller on the Internet instead of typing the seller's address into their browser (Oser, 2006).

[6] As pointed out earlier, the positive relation between click-through rates and advertiser valuations is similar to the positive correlation between frequency and monetary value in a direct marketing context (Bult and Wansbeek, 1995).

[7] It could be argued that advertisers may have higher valuations for informed consumers than for uninformed consumers, contrary to our assumptions. Our main results on the comparison between the two policies are not sensitive to this alternative assumption. A proof is available upon request.

from among the two described below. The policy details a mechanism by which the advertisers' bids are ranked, and specifies how the advertiser will be charged for ads displayed. Given our above assumption that the publisher can identify the type of the consumer, we assume that the publisher invites bids for each consumer type (similar to search engines accepting bids for each keyword). We also assume that while a consumer's type can be identified through the search keyword, the consumer's knowledge of advertisers cannot be ascertained beforehand. Furthermore, we assume that the publisher invites bids on a per-click basis for the clicks generated from displaying an advertiser's ad at the publisher's website. This assumption implies that advertisers, whose ads get displayed, eventually make payments on a per-click basis.[8]

After P's announcement of the auction policy, consumers of either type arrive sequentially at P's website. Advertisers simultaneously submit their bids to the publisher for their ad to be displayed in the ad slot to the consumer, with the bids depending on the consumer's type. For a given consumer type t, the strategy space for advertiser i involves a bid price, $b_i^t \in \Re^+$, and the content of the ad to be displayed. As discussed earlier, advertiser i can tailor his ad to appeal either to his own consumer type (type i) or to his competitor's type (type $j \neq i$). On receiving the bids, the publisher picks the winner in a manner consistent with the auction policy specified at the beginning of the game, and airs the ad to the consumer accordingly. If there is a tie between two advertisers, we assume that the publisher randomly chooses one of the advertisers for the slot. Let p denote the payment per click made by the advertiser whose ad is shown. We denote the publisher's profit by π_P and advertiser i's profit by π_i. Note that since consumers arrive in sequence, an advertiser can change his bid for a consumer type over time. As mentioned earlier, we focus our analysis on consumer type i without loss of generality.

Auction Policy: In choosing the policy, the publisher can make the bidder's rank and payment contingent on any variable that is observable to both the seller (here, publisher) and the bidder (advertiser) (McAfee and McMillan, 1987).[9] In a search-related advertising context, apart from the advertisers' bids, the publisher observes the click-through rates of advertisers' ads since website technology allows such information to be easily collected. Thus, depending on whether one or both of these observed variables are used, *prima facie* there can be two options for an auction policy:

[8] It is possible to make advertisers' payments contingent on impressions rather than clicks. However, the change does not affect our results since we do not model uncertainty in click-through rates or moral hazard issues (Hu, 2004).

[9] As an example, payment in mineral-rights auctions sometimes depend both on the winning bid and the amount of oil ultimately extracted by the winner (McAfee and McMillan, 1987).

1. The first option is a policy that uses only the advertisers' bids. Under this policy, which we call the highest bid policy, firm P airs the ad with the highest bid per click. For every click by a consumer, the winning advertiser is charged the bid submitted by the next highest bidder.

2. A second option is a policy that combines the advertisers' bids with the observed click-through rates of their ads for the consumer type (search keyword used). There are potentially many ways of combining these two pieces of information. However, a sensible way to do so is to take the product of the advertiser's bid and the click-through rate of the advertiser's ad, since this measure represents the expected profit from displaying the advertiser's ad. Thus, a policy based on this measure would entail airing of the ad with the highest expected profit for the publisher. The second-price payment implementation calls for the winning advertiser to make a payment per click that would yield the same profit as the losing advertiser, at the winning advertisers' click-through rate. Thus, if b_2 and c_2 represent the losing advertiser's bid and click-through rate, and if c_1 is the click-through rate of the winning advertiser, the winning advertiser's payment, p is given by b_2c_2/c_1.[10] We call this policy option the highest profit policy.[11] Such a policy is consistent with that studied by Feng et al. (2005).

The comparison of these two policies is a key objective of the analysis.

Information Structure and Equilibrium Concept: With advertisers able to change bids over time for newly arriving consumers, the game described above can be characterized as an infinitely repeated game for each consumer type (keyword). In general, such repeated games allow for a large set of equilibria sustained by complex punishment strategies. Rather than study the bidding strategies over time, we follow Edelman et al. (2007) in studying the long-run stationary equilibrium of this repeated game by restricting ourselves to simple strategies (strategies that do not depend on past strategies). We adopt the following assumptions from Edelman et al. (2007) about the characteristics of such a long-run stationary equilibrium: (i) Advertisers know each other's valuation because they learn about them over time,[12]

[10]This type of payment is instituted by Google with its highest profit policy. Note that one can conceive of a different payment policy whereby the winning advertiser is charged the bid of the next-highest bidder. However, such a policy may not necessarily satisfy the participation constraint of the winning advertiser. For instance, in the example presented in Footnote 2, the winning advertiser B under the highest profit policy would not want to participate if charged the losing bid of $4 per click because this amount exceeds his value per click.

[11]A third possible option is to rank advertisers purely on their click-through rates. However, since the ranking is then independent of the monetary bids of the advertisers, it is easy to show that such a policy does not in general lead to maximum profit for the publisher.

[12]See also Snir, Monderer, and Sela (1998).

and (ii) advertisers' bids must be best responses (in the sense of Nash) to each other. Additionally, we assume that click-through rates and the fraction of informed consumers are common knowledge among advertisers. The assumptions imply that the long-run stationary bids of the advertisers form an equilibrium in a simultaneous move, single-period bidding game of complete information. We characterize the dominant-strategy equilibrium of this game.[13] (Recall that in the base auction model, GSP reduces to a second-price auction.) Fudenberg and Tirole (1994) suggest that dominant-strategy equilibrium is a reasonable prediction of the outcome of a second-price auction. Moreover, the standard solution to the second-price auction in the literature satisfies the dominant-strategy equilibrium criterion (e.g., McAfee and McMillan, 1987).

4 Analysis of the Base Auction Model

We now compare the two auction policies under the base auction model, which has two advertisers competing for a single advertising slot at firm P's website. The proposition below characterizes the equilibrium under each policy.

Proposition 1 *The following strategies and payoffs constitute a dominant strategy equilibrium for consumer type i:*

- *Under the highest bid policy, $b_i^i = v_h$; $b_j^i = v_l$. Publisher P displays ads from advertiser i to consumer type i for a per-click payment, $p = v_l$. The publisher and advertiser profits per consumer of type i are $\pi_P = v_l c_h$, $\pi_i = c_h(v_h - v_l)$, and $\pi_j = 0$, $j \neq i$.*

- *Under the highest-expected-profit policy, $b_i^i = v_h$; $b_j^i = v_l$. Publisher P displays ads from advertiser i to consumer type i for a per-click payment, $p = v_l(\frac{\lambda c_l + (1-\lambda)c_h}{c_h})$. The publisher and advertiser profits per consumer of type i are $\pi_P = v_l(\lambda c_l + (1-\lambda)c_h)$, $\pi_i = v_h c_h - v_l(\lambda c_l + (1-\lambda)c_h)$, and $\pi_j = 0$, $j \neq i$.*

The most interesting result from the proposition is that the publisher profit of $v_l c_h$ under the highest bid policy is in general higher than the profit of $v_l(\lambda c_l + (1-\lambda)c_h)$ under the highest profit policy, with the exception that the profits are equal when $\lambda = 0$. The result is particularly surprising because, in picking the winning advertiser, the highest profit policy employs more information (the click-through

[13] Let s_i represent player i's strategy and s_{-i} represent the strategies of the players other than i. A pure strategy s_i is weakly dominated for player i if there exists a strategy σ_i such that the payoff $\pi_i(\sigma_i, s_{-i}) \geq \pi_i(s_i, s_{-i})$ for all feasible s_{-i} with the inequality being strict for at least one s_{-i}. We obtain a dominant-strategy equilibrium by iteratively eliminating dominated strategies for every player.

rate in addition to the advertiser's bid) than the highest bid policy. Next, we provide the intuition for this result.

For the general case of $\lambda > 0$, the publisher's lower profit under the highest profit policy is a result of the lower payment per click made by the winning advertiser i for consumer type i. The intuition is that the competing advertiser j obtains a lower click-through rate of $\lambda c_l + (1 - \lambda)c_h$ from consumers of type i, in comparison to advertiser i who obtains a click-through rate of c_h from this consumer. Note that in equilibrium, the competing advertiser j would tailor his ad to appeal to consumer type i. If the consumer is informed, his click-through rate depends only on the advertiser's type and not on the ad itself. Thus, the informed consumer would click advertiser j's ad at a rate of c_l. Conversely, the uninformed consumer is persuaded by the tailoring of the ad and would click at the rate of c_h. Thus, the effective click-through rate for advertiser j from consumer i is the weighted average of the click-through rates of informed and uninformed consumers. On the other hand, advertiser i is not penalized by informed consumers and thus enjoys an overall click-through rate of c_h. Thus, under the highest profit policy, advertiser i, who wins the bidding for consumer type i because of a higher expected profit per consumer, benefits from a lower payment per click because of his advantage over advertiser j in click-through rates as well as in valuation. In contrast, although advertiser i also wins the auction for consumer type i under the highest bid policy, his payment is higher in comparison to the highest profit policy because it does not incorporate his advantage in click-though rates. In the special case of $\lambda = 0$, advertiser j suffers no penalty in click-through rate from the consumer of type i. Consequently, both policies yield identical payments and profits in this case.[14] Note also that the publisher profit under the highest profit policy is a maximum when $\lambda = 0$ and decreases as λ increases.

Notwithstanding the difference in profits under the two policies, the proposition shows that the equilibrium strategies for the advertisers operating under the highest bid or the highest profit policies are identical (although the payment of the winning advertiser is higher under the highest profit policy). In each case, the bid submitted is the advertiser's true valuation for that consumer type. This result is consistent with the second-price auction literature which has shown that bidding his valuation is a dominant strategy for each bidder. When multiple slots are considered, however, we have a generalized second-place auction, and bidding one's true valuation is not always a dominant strategy (Edelman et al., 2007). Next, the proposition shows that both policies lead to the desirable result of a consumer being shown

[14]In order to avoid subsequent consumer dissatisfaction or because of regulatory reasons, it is possible that advertisers may refrain from tailoring ads that falsely suggest to a consumer that an advertiser is of a particular type. In such a case, the highest bid policy yields higher payments and profits even when $\lambda = 0$.

ads from an advertiser who represents the best possible match with the consumer type. Lastly, the proposition shows that the advertiser profits increase with advertisers' differentiation (as measured by the difference between their valuations). An interesting property of the above result is that the bids and payoffs do not depend on the number of consumers of a given type (segment). Thus, advertisers whose products are positioned toward small, niche consumer segments can outbid larger advertisers for this particular consumer segment. The key to this property is the ability of search advertising to identify the niche consumers through the keyword used by these consumers. In contrast, because mass media are not targeted, their pricing is driven by the number of mainstream consumers reached by them. These reasons may explain the tremendous popularity of search advertising, particularly among small businesses (Steel, 2007).

It is useful to compare the publisher profit from the auction model to the case when the publisher has complete information regarding advertisers' valuations. In the latter case, the publisher can capture all of the advertisers' valuations for consumer type i by setting a price of v_h per click for the slot, resulting in the maximum possible profit for the publisher. The resulting publisher's profit per consumer is $v_h c_h$ and the advertiser profits are zero. These results represent the first-best outcome for the publisher and the publisher's profit in the base auction model is lower as expected. As discussed earlier, the auction model is necessitated by the publisher's lack of knowledge about advertiser's valuations.[15]

5 Generalization: Multiple Ad Slots and Multiple Advertisers

5.1 Model and Equilibrium Analysis

In this section, we show that our main results on the comparison between the highest bid and highest profit auction policies carry over to a generalized second price auction involving multiple ad slots. First, we sketch out the general model. Then, we analyze a generalization to multiple slots of the equilibrium concept used in the base auction model.

Let there be M advertisers in the product market, competing for K advertising slots at the publisher's website. As discussed earlier, we assume $K < M$ without loss of generality. Each advertiser's product is differentiated and appeals most to one of M consumer types or segments. As in the previous section, we derive the equilibrium results for one consumer type, say t, since similar results carry over to

[15]It is possible that the publisher learns about advertisers' valuations, rescinds the auction, and then uses the information he learned to post prices to his advantage. However, to successfully implement an auction in the first place, there should be a perceived commitment on the part of the publisher not to renege *ex post* from the auction model (McAfee and McMillan, 1987) and post prices.

other consumer types as well. Advertisers are indexed by i, with $i \in \{1, 2, ..., M\}$. It is important to note that the meaning of an advertiser's index number is somewhat different from that in the base auction model. An advertiser is assigned a lower index number if his product is closer to the needs of consumer type t. We denote advertiser i's valuation of a consumer of type t by v_i, $v_i > 0$. Assuming as before that v_i reflects advertiser i's expected payoff from a consumer of type t visiting his website, and assuming that this expected payoff increases with his closeness to consumer type t, we have $v_i > v_{i+1}$ for every i.

With multiple slots, the click-through rate of an ad is influenced by whether or not the consumer is informed about the advertisers, the identity of the advertiser, and the position of the ad slot on the publisher's webpage. As noted earlier, ad slots differ in their click-through rates with positions at the top of the webpage commanding inherently higher click-through rates. Therefore, consistent with prior work such as Feng et al. (2005) and Edelman et al. (2007), we characterize the effective click-through rate, γ_i^k, for an ad from advertiser i shown in slot k as the product of the consumer's click-through rate for that ad, χ_i, and the slot's "inherent" click-through rate factor, α_k. Thus, we assume that the inherent click-though rate of a slot is independent of the advertiser or the consumer type. Without loss of generality, we assume that the K slots are sorted in the descending order of their inherent click-through rate factors. Thus, $\alpha_k > \alpha_{k+1}$. By normalizing α_1 = 1, α_k for $k > 1$ can represent the factor by which the click-through rate of an ad is reduced because of being positioned in the k^{th} slot instead of the first slot. Consistent with prior work, we also treat $\alpha_k = 0$ for $k \geq K + 1$ since ads beyond the K slots are never displayed and, therefore, never clicked.

As in the base auction model, we assume that λ and $(1-\lambda)$ represent the fraction of informed and uninformed consumers. If consumers are informed about the advertisers, it is reasonable to expect that the consumer has a higher click-through rate for an advertiser whose product is closer to her need. Thus, we assume $c_i > c_{i+1} > 0$. As in the base model, our assumptions imply that click-through rates and advertiser valuations for an informed consumer are positively related. If a consumer is not informed about the advertisers, then he is unable to distinguish between advertisers who tailor their ad to appeal to his type. In this case, we assume that all advertisers who tailor their ad to consumer type t obtain a click through rate of c_1, i.e., the click-through rate corresponding to the advertiser that best fits consumer type t. Therefore, the average click-through rate for advertiser i is: $\chi_i = \lambda c_i + (1 - \lambda)c_1$.

The sequence of moves is the same as in the base auction model except that the publisher now assigns K ad slots to advertisers. To do so, the publisher ranks the M advertisers based on the announced auction policy. Let $g(j)$ represent the index number of the advertiser with the j^{th} highest rank so that $g(j) \in \{1, 2, \ldots, M\}$ for j = 1,2,...,M. The publisher then assigns the top K advertisers to the K ad slots such that the advertiser with the highest rank gets the first (most attractive) slot, the second-highest advertiser gets the next slot, and so on, with each advertiser getting at most one slot. If there is a tie between two advertisers, we assume that the publisher randomly chooses one of the advertisers for a slot.

Under the highest bid policy, let b_i^{HB} be the bid from advertiser i, p_k^{HB} be the payments per click made by the advertiser assigned to slot k, and π_k^{HB} be the publisher profit from slot k. Let the notations for the corresponding variables under the highest profit policy be b_i^{HP}, p_k^{HP}, and π_k^{HP}. In our model, $\pi_{P,k}^s = p_k^s \chi_{g(k)} \alpha_k$ for $s \in \{HB, HP\}$. The implementation of the two policies in the case of multiple slots is analogous to that in the base auction model. Under the highest bid policy, if the ad in slot k is clicked, the advertiser assigned to this slot pays the next-highest bid per click, i.e., $p_k^{HB} = b_{g(k+1)}^{HB}$, $k = 1, 2, \ldots, K$. Similarly, under the highest profit policy, the advertiser assigned to slot k makes a per-click payment that would yield the same profit as the losing advertiser, at the winning advertisers' click-through rate, i.e., $p_k^{HP} = b_{g(k+1)}^{HP} \chi_{g(k+1)} / \chi_{g(k)}$, $k = 1, 2, \ldots, K$. A summary of the notation used in the generalization is presented in Appendix D.

Information Structure and Equilibrium Concept: As in the base auction model, we follow Edelman et al. (2007) in studying the long-run stationary equilibrium of the repeated game induced by the GSP auction model. The information assumptions are identical to those in the base model. Consistent with Edelman et al. (2007), we assume (as discussed earlier) that the long-run stationary bids of the advertisers form an equilibrium in a simultaneous move, single-period bidding game of complete information. However, the requirement of best response is too weak and allows for some apparently implausible equilibria (Edelman et al., 2007). Therefore, Edelman et al. (2007) argue that, in a stable long-run equilibrium, an advertiser i assigned to slot k should not find it optimal to dislodge the next higher-ranked advertiser j who is assigned the more preferred slot $k-1$. In order to displace advertiser j, advertiser i can use the simple strategy of increasing his bid (thereby increasing advertiser j's payment without affecting advertiser i's own payoff) to the point where advertiser j prefers to bid lower just enough to be ranked below advertiser i and thus obtain the lower spot. Advertiser i would not employ such a strategy if the bids are such that he cannot improve his payoff through an exchange of positions with advertiser j. Edelman et al. (2007) characterize an equilibrium that satisfies this condition to be "locally envy-free," and we adopt this refinement in our paper. Under the highest bid policy, the above-described outcome, in which advertiser j bids lower just enough to exchange positions with advertiser i, results in a payoff to advertiser i that equals the payoff when advertiser j exchanges bids with advertiser i. This exchange of bids is reflected in the definition of a locally envy-free equilibrium for the highest bid policy. In the case of the highest profit policy, the exchange of positions is equivalent, in terms of the payoff to advertiser i, to an exchange of expected profit levels (to the publisher) between advertisers i and j, and this is used in the definition of the locally envy-free equilibrium for this policy. (See Appendix E for additional discussion of the locally envy-free equilibrium with examples.) Although the locally-envy free concept avoids some implausible equilibria, multiple locally envy-free equilibria exist. Among these equilibria, we consider a particular one to be the most reasonable and stable long-run equilibrium. We refer to this equilibrium as the "locally dominant envy-free equilibrium," defined as follows:

Definition 2 *An equilibrium of the simultaneous-move game induced by the GSP auction is locally dominant envy-free under the following conditions:*

(i) Under the highest bid policy, for any rank $j \leq K+1$, $\alpha_j(v_{g(j)} - p_j^{HB}) = \alpha_{j-1}(v_{g(j)} - p_{j-1}^{HB})$ *and for rank* $j > K+1$, $\alpha_j(v_{g(j)} - p_j^{HB}) \geq \alpha_K(v_{g(j)} - p_K^{HB})$; *and*

(ii) Under the highest profit policy, for any $j \leq K+1$, $\alpha_j(v_{g(j)} - p_j^{HP}) = \alpha_{j-1}(v_{g(j)} - p_{j-1}^{HP}(\frac{\chi_{g(j-1)}}{\chi_{g(j)}}))$ *and for any rank* $j > K+1$, $\alpha_j(v_{g(j)} - p_j^{HP}) \geq \alpha_K(v_{g(j)} - p_K^{HP}(\frac{\chi_{g(K)}}{\chi_{g(j)}}))$.

For rank $j \leq K$, the above definition requires that, under the highest bid policy (the highest profit policy), the payoff to the advertiser assigned to slot j does not change if that advertiser exchanges bids (expected profit levels) with the one in the previous slot.[16] For rank $j > K$, the definition ensures that every advertiser not assigned a slot cannot gain by similarly exchanging positions with the advertiser assigned to the last slot, K.

For the model with multiple slots, we characterize the locally dominant envy free equilibrium, which is compelling for the following reasons. First, in this equilibrium, advertiser i in slot k bids such that his marginal payment (i.e., the difference between his highest payment were he to win the $(k-1)^{\text{th}}$ slot and his payment for the k^{th} slot) equals the marginal value of extra clicks obtained from position $k-1$.[17,18] If the marginal payment exceeds the marginal value of extra clicks, advertiser i would earn less if advertiser j (who occupies position $k-1$) underbids him. Thus, it

[16] Note that advertiser $g(j)$'s payoff is generically given by $\alpha_j \chi_{g(j)}(v_{g(j)} - p_j)$. We omit $\chi_{g(j)}$ from the equilibrium conditions in Definition 2 as this term cancels out of both sides of the conditions.

[17] This can be seen by rearranging the condition for the highest bid policy: $\chi_{g(k)}(\alpha_{k-1}p_{k-1} - \alpha_k p_k) = v_{g(k)}\chi_{g(k)}(\alpha_{k-1} - \alpha_k)$; and also similarly for the highest profit policy: $\alpha_{k-1}\chi_{g(k-1)}p_{k-1} - \alpha_k \chi_{g(k)} p_k = v_{g(k)}\chi_{g(k)}(\alpha_{k-1} - \alpha_k)$.

[18] In other locally envy free equilibria, the marginal payment exceeds the marginal value of the extra clicks.

would appear the condition for the locally dominant envy free equilibrium is more consistent with stability. Second, the locally dominant envy-free equilibrium (and the locally envy-free equilibrium) results in a stable assignment in the sense that no advertiser wants to exchange positions with another in equilibrium. Third, advertiser i's bid satisfying Definition 2 weakly dominates other locally envy free bids in his "local" competition with advertiser j. Fourth, with one slot, when the GSP auction reduces to a standard second-price auction, the locally dominant envy-free equilibrium coincides with the dominant strategy equilibrium (Fudenberg and Tirole, 1994) of the second-price auction (the concept we have used in the previous section).

The locally dominant envy free equilibrium has other interesting properties. First, the locally dominant envy-free equilibrium results in the same assignments of advertisers to slots and identical advertiser payments as a Vickrey-Clarke-Groves auction mechanism for multiple slots (Edelman et al., 2007). Second, it can be shown that the locally dominant envy-free equilibrium is unique. Third, the advertisers' profits are higher in the locally dominant envy-free equilibrium than in any other locally envy-free equilibria, whereas the publisher's profit is lowest (cf. Edelman et al., 2007). We note that Varian (2007) finds empirical evidence in support of locally envy-free equilibria with GSP auction data.

Equilibrium Analysis: Next, we derive the locally dominant envy-free equilibrium for the multiple slot case. In the proposition below, we characterize the equilibrium results under each policy in terms of the payments per click for individual slots. Following the proposition, we discuss how the advertisers' equilibrium bids from the advertisers can be computed from the equilibrium payments.

Proposition 3 *In the unique locally dominant envy-free equilibrium, advertiser j is assigned a rank j, $1 \leq j \leq K+1$, under either the highest bid policy or the highest profit policy. Further,*

(i) For slot k, such that $1 \leq k \leq K$,

- *under the highest bid policy, $p_k^{HB} = \sum_{j=k}^{K} v_{j+1} \left(\frac{\alpha_j - \alpha_{j+1}}{\alpha_k}\right)$, and*
- *under the highest profit policy, $p_k^{HP} = \sum_{j=k}^{K} \left(\frac{\chi_{j+1}}{\chi_k}\right) v_{j+1} \left(\frac{\alpha_j - \alpha_{j+1}}{\alpha_k}\right)$.*

(ii) For every slot $1 \leq k \leq K$, $p_k^{HB} \geq p_k^{HP}$ and $\pi_{P,k}^{HB} \geq \pi_{P,k}^{HP}$. The payments and the profits are equal only when $\lambda = 0$.

Notably, the above proposition verifies that, even for multiple slots, the publisher profit is generally higher under the highest bid policy than under the highest profit policy, with the exception that the profits are equal in the special case of $\lambda = 0$. Thus, we demonstrate the robustness of our main result. An interesting aspect about Proposition 3 is that the profit and payment per click generated from every slot is at least as good in the highest bid policy as in the highest profit policy.

The locally dominant envy-free equilibrium presented in the above proposition results in an assortative match between slots and the advertisers, i.e., the higher an advertiser's valuation, the more attractive is the slot assigned to him (see also Lemma 1 and its proof in Edelman et al., 2007). Under the highest bid policy, an assortative match leads to the implication that the advertiser in slot k has a valuation of v_k, and that $v_k > v_{k+1}$ for $k \in \{1, 2, \ldots, K\}$. A similar argument also holds under the highest profit policy, where the assortative match causes $v_k \chi_k > v_{k+1} \chi_{k+1}$ for $k \in \{1, 2, \ldots, K\}$. Note that an assortative match in our model implies that the advertiser who represents the best possible match to the consumer (highest v_i) is assigned the most attractive slot in equilibrium. In general, this sorting of advertisers may be expected to have an influence on consumer's click-through rates for slots. However, our model assumes that click-through rates for slots are independent of the order in which advertisers are sorted in equilibrium. Relaxing this assumption is an interesting avenue for future research.

The equilibrium bids can be derived from the payments in the Proposition 3 based on the definition of the auction policies. Thus, the $(k+1)^{\text{th}}$ ranked bid is given by p_k^{HB} under the highest bid policy and by $p_k^{HP} \frac{\chi_k}{\chi_{k+1}}$ under the highest profit policy. Note that the payments in Proposition 3 only determine the equilibrium bids for ranks $2, 3, \ldots, K+1$. The equilibrium bid for the highest ranked advertiser could be any value greater than that of the second highest ranked advertiser. It can be seen by inspection of p_k^{HB} and p_k^{HP} that the payment per click increases for more attractive slots under either policy. However, the advertisers' bids while higher for the more attractive slots under the highest bid policy, need not be so under the highest profit policy (See Example 3 in the Technical Appendix). Lastly, it can be shown that $p_k^{HB} < v_{k+1}$ for $k \in \{1, 2, \ldots, K-1\}$.[19] Recalling that $p_k^{HB} = b_{k+1}^{HB}$, we see that advertisers bid below their valuations. This conclusion is also true under the highest profit policy.

5.2 Discussion

We have established in the foregoing analysis that the highest bid policy generates at least as much publisher profit as the highest profit policy. Thus, surprisingly, the

[19] By rearranging terms, we get $v_{k+1} - p_k^{HB} = [\sum_{j=k+1}^{K} \alpha_j (v_j - v_{j+1})]/\alpha_k > 0$ because $v_j > v_{j+1}$. Similarly, under the highest profit policy, we get $v_{k+1}\chi_{k+1} - p_k^{HP}\chi_k = [\sum_{j=k+1}^{K} \alpha_j (v_j\chi_j - v_{j+1}\chi_{j+1})]/\alpha_k > 0$ because $v_j\chi_j > v_{j+1}\chi_{j+1}$. This implies $v_{k+1} > p_k^{HP}\chi_k/\chi_{k+1} = b_{k+1}^{HP}$.

publisher's auction policy that relies on more pieces of information (valuation and click-through rates) does worse. Based on these considerations, it would appear that the highest bid policy is preferable for maximizing the profit per consumer. We now consider the sensitivity of this result to model assumptions.

5.2.1 Sensitivity Analysis to the Positive Correlation Assumption

We show below that the assumption that advertiser valuations and click-through rates are positively related for an informed consumer is important for our main result that the highest bid policy is more profitable than a highest profit policy. As argued earlier, it seems reasonable that the more valuable a consumer is to an advertiser in a product market, the more likely the consumer is to click at the advertiser's ad, resulting in a positive correlation between valuation and click-through rate. In our running example, an advertiser with the better selection of Howard Miller wall clocks attaches a higher value to a consumer of the Howard Miller type, and this consumer is also more likely to click his ad than an ad from the competing advertiser (of the Seth Thomas type). Later, we present empirical evidence in support of the positive correlation assumption. Nevertheless, it is useful to consider how the results may change if this assumption does not hold.

If click-through rates are independent of advertiser valuations (implying zero correlation between them) for informed consumers, both policies yield the same profit. The rationale is that informed consumers become similar to uninformed consumers in this case, so that the previous results reduce to the special case of $\lambda = 0$. Suppose, however, that advertiser valuations and click-through rates are negatively correlated across advertisers for a consumer type. Considering first the base auction model, such a negative correlation would result if, for consumer type i, advertisers i and j have click-through rates of c_l and c_h respectively with $c_l < c_h$, even though their respective valuations are v_h and v_l with $v_h > v_l$. If c_h exceeds c_l by a sufficient amount, we could have $v_l c_h > v_h c_l$. In this case, we can show that the profit is higher under the highest profit policy when $\lambda > 0$ even though the payment per click is higher for the highest bid policy. Moreover, the identity of the advertiser winning the auction under the two policies is also different. In the other case of negative correlation, namely when $v_l c_h < v_h c_l$ even though $c_h > c_l$, both the payment per click and profit are higher under the highest profit policy.

In the multi-slot model of Section 5, the results are less clear-cut in the case of negative correlation as seen below. First, it can be shown that the profit expressions in Proposition 3 (i) apply even when advertiser valuations and click-through rates are negatively correlated except that the payment per click, $\pi_{P,k}^{HP}$, for slot k under the highest profit policy is modified as follows:

$$p_k^{HP} = \frac{1}{\alpha_k} \sum_{j=k}^{K} \left(\frac{\chi_{g(j+1)}}{\chi_{g(k)}} \right) v_{g(j+1)} \left(\alpha_j - \alpha_{j+1} \right)$$

This revised expression for p_k^{HP} recognizes the different assignment of advertisers to slots under the highest profit policy for the case of negative correlation, i.e., it is no longer an assortative match, so that $g(j) \neq j$ in general for an advertiser whose rank is j. Because $\pi_{P,k}^{HB} = \alpha_k \chi_k p_k^{HB}$ and $\pi_{P,k}^{HP} = \alpha_k \chi_{g(k)} p_k^{HP}$, we can obtain the profit advantage in slot k for the highest-bid policy as the following:

$$\begin{aligned} \pi_{P,k}^{HB} - \pi_{P,k}^{HP} &= \sum_{j=k}^{K} [v_{j+1}\chi_k - v_{g(j+1)}\chi_{g(j+1)}] \left(\alpha_j - \alpha_{j+1} \right) \\ &= \sum_{j=k}^{K} [(v_{j+1}\chi_{j+1}) \frac{\chi_k}{\chi_{j+1}} - v_{g(j+1)}\chi_{g(j+1)}] \left(\alpha_j - \alpha_{j+1} \right) \end{aligned}$$

Note that the term $v_{j+1}\chi_k$ in the above expression reflects the profit from the highest bid policy. Because $\alpha_j > \alpha_{j+1}$, the sign of the above expression depends on the term in square brackets. Further, because advertisers are ranked based on the their expected profit under the highest profit policy, in general we would have $v_j\chi_j < v_{g(j)}\chi_{g(j)}$ for highly-ranked slots (low j) and vice versa. Therefore, the individual terms within the square brackets can have positive and negative values, rendering the overall sign ambiguous. However, it is interesting that the highest bid policy can be more profitable than the highest profit policy even when advertiser valuations and click-through rates are negatively correlated.[20]

We have argued that a positive correlation between advertiser valuations and click-through rates is reasonable. However, it may appear that a negative correlation becomes possible in the presence of niche advertisers who may have high valuations but may experience fewer clicks by virtue of appealing to a narrow segment. This argument is not necessarily valid in a search advertising context because a critical feature that distinguishes search advertising from traditional mass-marketing is that niche needs served by the niche advertisers can generally be defined in terms of

[20]Consider the following example with 2 slots and 3 advertisers with $v_1 = 2.55$, $c_1 = 0.515$, $v_2 = 2.5$, $c_2 = 0.5$, $v_3 = 2.4$, $c_3 = 0.52$, $\lambda = 1$, $\alpha_1 = 1$ and $\alpha_2 = \frac{1}{6}$. The correlation between the valuations and the click-through rates is negative (-0.4193). In this case, using the formulas in Proposition 3 (i) modified by equation above, we can establish that the publisher's total profit under the highest-bid and highest-profit policies respectively are 1.479 and 1.458. In the above example, if $c_3 = 0.6$ and all other parameters are retained the same, the publisher profit is higher with the highest profit policy.

niche keywords. In such a case, the niche consumer would use a niche keyword for which the ads from niche advertisers become more relevant and more likely to be clicked (resulting in a positive correlation between advertiser valuation and click-through rate).[21] It is for precisely this reason that our profit results in Propositions 1 and 3 are independent of the size of the consumer segments who are the best match for an advertiser type. Then, what might be conditions that produce a negative correlation between advertiser valuations and click-through rates? Perhaps a negative correlation is possible if the niche needs are not definable using keywords or if some advertisers with low valuations are exogenously endowed with higher awareness resulting in higher click-through rates for these advertisers.

Note that the positive correlation that we assume is only between valuations and click-through rates *across advertisers* for a consumer type. This positive correlation should not be confused with how average consumer valuations would change in the aggregate for *a given advertiser* as he expands his ad's appeal to include consumer types other than his own consumer type. In fact, our model assumes that the average (and marginal) consumer valuation would drop even as aggregate click-through rates increase suggesting a negative correlation between these two entities. This negative correlation between click-through rates and click valuations for *a given advertiser* is consistent with the general notion in marketing that broader targeting increases customer acquisition rate but leads to acquisition of less valuable customers. However, this negative correlation is distinct from the positive correlation between valuations and click-through rates *across advertisers* for a consumer type.

5.2.2 Other Model Assumptions

A limitation of our model is that we do not consider competition between publishers. With publisher competition, it is possible that if a publisher using the highest bid policy competes for advertisers with another publisher using the highest profit policy, advertisers may prefer the latter publisher because its auction policy allows them to make more profit. In such a case, it is possible that the former publisher may be unable to fill his ad slots and may, therefore, switch to a highest profit policy. However, if the publisher using the highest bid policy enjoys some differentiation advantages such as a loyal consumer base, it may be optimal for the publisher to continue with its auction policy. Because of the publisher's loyal consumers, ad-

[21] As an example, Hentschel is a niche, premium-priced brand in the wall-clock category while Howard Miller is a more mainstream brand. Therefore, an average consumer in the market is likely to click at an ad featuring Hentschel clocks at a much lower frequency in comparison to one featuring a Howard Miller clock. However, in a search-advertising context, a consumer preferring the niche Hentschel brand may use a search keyword of Hentschel wall clock, thereby revealing his preference or type. For this individual consumer (or keyword), it would be reasonable to expect that both click-through rates and advertiser valuations are higher for an ad featuring the niche Hentschel brand than the mainstream Howard Miller brand.

vertisers are likely to consider bidding for ad space at that publisher's website as well (much like traditional advertisers choose to advertise in multiple media like newspapers and televisions) even though its auction policy is not as favorable to them.

For the publisher, the highest bid policy does better with any non-zero fraction of informed consumers. Thus, in the scenario where all the consumers are uninformed, if advertisers use offline advertising to create brand awareness among consumers, it is plausible that there are at least some informed consumers. Therefore, the use of offline advertising would favor the highest bid policy from the publisher's point of view.

Our model and results have been agnostic about whether the publisher is a content provider (e.g. NYTimes.com) or a search engine such as Yahoo. However, in the case of search engines, an advertiser has an opportunity to appear on the search engine's organic listings and this may influence his bidding for search advertisements (sponsored listings). By allowing advertiser click-through rates to be less than one, our model implicitly recognizes, in the case of search engines, the competition from organic search listings, along which advertisers might be listed. Our results would continue to hold if we assume that some consumers choose to click on sponsored advertiser listings while others choose to click on the organic listings, so that sponsored advertiser listings can be analyzed in isolation. There is some evidence that consumer clicking behavior is consistent with such an assumption (see Yao and Mela, 2009). However, if some informed consumers who typically click on sponsored listings continue on to the organic listings in search of their favorite advertiser, this advertiser may be induced to bid less for sponsored listings.[22] In Appendix C, we analyze the effect of such consumer behavior in the presence of organic listings using an extension to our base model. Specifically, recall that in the base model, the click-through rate of informed consumers of type i decreases by $(c_h - c_l)$ on the sponsored listing when advertiser j's ad is shown instead of advertiser i's. Our base model assumes that this deficit in clicks from the sponsored listing is entirely lost to either advertiser. In the extended model, we assume that when the sponsored listing features advertiser j for consumer type i, a fraction η of the deficit in clicks of the sponsored listing from informed consumers gets transferred to advertiser i's organic listing. The implicit assumption is that because advertiser i is the best match for consumer i, this advertiser will be featured prominently in the organic listings.

With such an extended model, we show in the appendix that firm i's dominant strategy is to bid below his valuation. Further, we find that the highest bid policy generates more publisher profit than the highest profit policy if $\eta\lambda \leq \frac{v_h - v_l}{v_h}\frac{c_h}{c_h - c_l}$.

[22]We thank an anonymous reviewer for this insight.

Thus, in the presence of organic listings, the highest bid policy is preferable when (i) informed consumers prone to click on sponsored listings are less likely to click on organic listings when the perceived match of the set of sponsored advertisers is low (i.e., η is small); (ii) advertisers are well differentiated so that valuations are significantly different in comparison to the difference in click-through rates (which can be minimized by advertisers effectively tailoring the ad to the consumer); and (iii) the number of informed consumers is not too high (but non-zero).[23] As noted earlier, empirical research has found that sponsored listings are almost exclusively clicked by a small segment of consumers (Yao and Mela, 2009). Further, the attractiveness of search advertising to advertisers indicates that advertisers are typically well differentiated as otherwise all their incremental profits from search advertising would be competed away by bidding (see Proposition 1). Moreover, it seems reasonable to assume that consumers who search are predominantly likely to be uninformed (otherwise, they would not search). Therefore, the conditions required to render the highest bid policy more profitable for a search engine may seem reasonable enough to be satisfied in many situations. Nevertheless, the analysis shows that the highest bid policy may not outperform the highest profit policy for a search engine (with organic listings) if these conditions do not hold even though advertiser valuations and click-through rates are positively correlated. Thus, this analysis points to some limitations of our conclusions when applied to search engines.

6 Empirical Study

We collected data to ascertain empirical support for our theoretical conclusion that the highest-bid policy is more profitable than the highest-profit policy for a publisher. Since comparative profitability data from two publishers practicing the two alternative auction policies were not available, we analyze other data and establish indirect support to our theoretical conclusions. In particular, we test for the correlation between click-through rates and valuations using advertising click data from a search engine using the highest bid policy. The results from this study suggest that advertiser valuations and click-through rates are positively correlated validating an important assumption of our model and implying that the highest bid policy may be more profitable than the highest profit policy under these conditions. However, as discussed in section 5.2.2, positive correlation may not imply greater profitability of the highest bid policy for a search engine if other (weak) conditions are not satisfied. Because we do not verify these supporting conditions in our empirical study, our empirical conclusions are subject to this limitation. We describe below the study in greater detail.

[23] There exists industry research which shows that an advertiser enjoys higher click-through rates from its organic listings when his ad appears in the sponsored listings (iCrossing, 2007). It can be shown that such synergies favor the highest bid policy for a publisher.

6.1 Analysis

For this study, we use data on keyword-related advertising clicks recorded during the period, September 1-30, 2008, by a search engine that uses the highest bid policy. For each recorded click, the data contains the searched keyword, a date-time stamp, the identity of the advertiser whose ad was clicked, and the payment made by the advertiser for the click. The data consists of 9.2 million clicks related to 351,000 keyword phrases. The search engine deals with advertising intermediaries who place ads on behalf of their advertiser-clients, although we assume that the clients have the final say on the bid amounts. A limitation of the data is that it records the identity of the advertiser intermediary placing an ad rather than that of the advertiser client. In all, there are 31 unique advertiser/advertiser intermediaries in the data. We analyze the clicks for a random sample of 2000 keywords drawn from the set of 15,000 keywords with at least 15 clicks during the time period of the data. These 2000 keywords accounted for about 631,000 clicks by 25 advertiser intermediaries during the data period. As discussed above, the objective of the empirical analysis is to determine if advertiser valuations and click-through rates are positively correlated. Although we do not observe advertiser valuations in the data, the following proposition extends our analytical results to show that the payment per click is higher for more attractive slots under the highest-bid policy.

Proposition 4 *Under the highest bid policy, for slot k such that* $1 \leq k < K$, $p_k^{HB} > p_{k+1}^{HB}$.

From our discussion in section 5, we also know that the highest bid policy results in an assortative match so that $v_k > v_{k+1}$. Note that this result and Proposition 4 are independent of the correlation between advertiser valuations and click-through rates. Combining the results that $v_k > v_{k+1}$ and $p_k^{HB} > p_{k+1}^{HB}$ under the highest bid policy, we can conclude that advertiser valuations and their payments per click are strongly positively correlated under the highest bid policy. Thus, we can infer the correlation between advertiser valuations and click-through rates by using advertisers' payments per click as a proxy variable for their valuations. The use of a proxy for an explanatory variable may bias the estimated coefficient of the explanatory variable towards zero if the correlation between the proxy variable and the true variable is less than perfect (Levi, 1973). Therefore, a finding of a positive relationship between payments per click and click-through rates would be a conservative indicator of a positive relationship between advertiser valuations and click-through rates.

Proceeding to our model specification, let U_{iat} represent the latent propensity of the ad from advertiser a for keyword phrase i and time t to be clicked, for $a \in A_{it}$. A_{it} is the set of advertisers bidding for keyword i at time t. The value of U_{iat}

influences the observed dependent variable, y_{iat}, which takes on the value of 1 if advertiser a is clicked at time t for keyword i, and zero otherwise. Note that at least one of the y_{iat} is 1 for any i, a, and t, since our data only includes observations when an advertisement was clicked. We specify U_{iat} as follows:

$$U_{iat} = \theta_a + \beta_j x_{jiat} + \mu_i p_{iat} + \varepsilon_{iat}$$

In the above equation, θ_a is an advertiser fixed effect, x_{jiat} is a vector of explanatory variables with corresponding "fixed" parameter vector β_j, p_{iat} is the observed payment per click and μ_i is a "random effect" parameter that varies with keyword i. We assume that ε_{iat} follows an i.i.d. extreme value distribution and that $\mu_i = \overline{\mu} + \eta_\iota$, where $\overline{\mu}$ is the mean value of μ_i across keywords and $\eta_\iota \sim N(0, \sigma^2)$. Thus, we allow for heterogeneity across keywords in the relationship between payment per click (proxy variable for advertiser valuation) and the propensity to click, U_{iat}. If our theoretical assumption of positive correlation between advertiser valuation and click-through rates is true, we should observe $\overline{\mu} > 0$. Our assumptions about the error term, ε_{iat}, and the use of some random parameters result in a mixed logit model (Brownstone and Train, 1999). We estimate the mixed logit model using simulated maximum likelihood.

Our data does not record the set of losing advertisers (those who were not clicked) and their bids for each observation (click). Therefore, for each i and t, we construct the set A_{it} of bidding advertisers and the payments per click that the losing advertisers would have been charged had they been clicked. We develop this data as follows. We observed that there was significant variation in the payments per click charged to advertisers even within a day for many keywords. This observation suggested that the market was quite dynamic with advertisers varying bids frequently. Consequently, we included in set A_{it} those advertisers whose advertisements were clicked for keyword i within $\pm$ 2 hours of time t. (A few observations were dropped after this process because the competitive set A_{it} of advertisers was equal to one.) For each advertiser in the set A_{it}, we used the most recent payment per click charged to the advertiser as that advertiser's potential payment had his advertisement been clicked. (If a recent payment per click was not available, we found the payment per click charged to the advertiser at the most proximate time, going forward.) We changed the duration that we used to construct A_{it} to a day, and to three days and found that our results were robust to such alternate assumptions. Table 1 presents the descriptive statistics of the random keyword sample that we analyzed. The table shows that the set A_{it} comprised of two to nine advertisers across the observations.

Since the propensity to click an ad depends on an advertiser's position on the web page, we include variables that control for such positional effects in x_{jiat}. For

Table 1: Descriptive Statistics of Sample

Variable	Mean	Std. Dev.	Minimum	Maximum
CPC ($)	0.03954	0.07233	0.01	7.44
Advertisers/ Observation	3.028	1.112	2	9
Clicks Per Keyword	438.	1633.	1	38909

Table 2: Estimation results for Mixed Logit Model

Variables/ Parameters	Model I		Model II		Model III	
	Parameter Estimate	Std. Error	Parameter Estimate	Std. Error	Parameter Estimate	Std. Error
$POSTN$	0.260 (**)	0.0029	2.60(**)	0.0031	-0.049(**)	0.0031
$\bar{\mu}$	0.193(**)	0.0034	0.192(**)	(0.0036)	0.071(**)	0.0017
σ			0.0015	0.0016	0.0016	0.0017
Log-likelihood	-655,017		-655,017		-576,346	
BIC	1,310,061		1,310,074		1,152,919	

** Parameters are significant at 1% level

each observation, we rank advertisers based on the payment per click and set a variable named $POSTN_{iat}$ equal to the advertiser's rank, with a rank or $POSTN$ equal to 1 presumably representing the most attractive position on the webpage. Note that we do not observe the actual position occupied by an advertiser on the webpage but that we assign positions to advertisers based on their observed payments per click (which are an indication of their bids). Because positional effects are expected to be independent of keyword, we use "fixed" parameters β_j for this variable. Since the payment per click, p_{iat}, vary significantly across keywords (see Table 1), we standardize this variable for each observation to have a mean of zero and standard deviation of 1 to control for heterogeneity across keywords in the absolute values of payments per click. In estimating advertiser fixed effects, θ_a, we combine advertisers accounting for less than 1% of the observations into a single "Other Advertiser." Thus, we estimate advertiser fixed effects for 14 advertisers, with the fixed effect of one advertiser being normalized to zero. Recall that our data records the identity of the advertiser intermediaries rather than that of the advertising firm. This accounts for the low number of advertiser fixed effects needed in the model.

Table 2 presents the estimation results of the mixed logit model (Model III) in addition to simpler model specifications. Specifically, Model I does not allow for heterogeneity across keywords in μ_i by setting $\sigma = 0$. Model II allows for heterogeneity in μ_i but does not include advertiser fixed effects, while Model III incorporates both these effects. In Table 2 we do not present the advertiser fixed-effect parameters for Model III to conserve space. In all models, the parameter $\overline{\mu}$ is positive and significant suggesting that advertiser valuation, for which payment per click serves as a proxy variable, is positively correlated with click-through probability. The heterogeneity parameter σ is not statistically significant in both Models II and III. The positive value of $POSTN$ in Models I and II is a surprising result, as we expected that the probability of a click will reduce with increasing value of $POSTN$ (less attractive positions). However, when advertiser fixed effects are included in Model III, the $POSTN$ parameter takes on the expected negative sign. Both the likelihood ratio test and the Bayesian Information Criterion (BIC) (Schwarz, 1978) indicate that Model III is the preferred specification in comparison to Models I and II.

We note the following limitations of our empirical analysis. A limitation is that the positions of different advertisers were inferred from the data based on their recent payments per click and were not directly observed. This inference process used to generate the data may introduce measurement errors. This caveat could be addressed by future research with data that record positions on the webpage as assigned to advertisers. We also note that another limitation of our data is that the payments per click (bids) of losing advertisers had to be inferred rather than being made directly available. It is however difficult to predict whether these measurement issues are capable of biasing our results.

Overall, the empirical result on μ_i lends support to our theoretical assumption of a positive correlation between advertiser valuation and click-through rates. This observation of positive correlation is important because such a scenario is more likely to support our conclusion that the highest bid policy is more profitable to the publisher than the highest profit policy. However, as noted earlier, a positive correlation between advertiser valuations and click-through rates may not ensure that the highest bid policy is more profitable for a search engine unless other ancillary conditions are met. Another empirical issue that arises is the effect of the strength of the positive correlation between advertiser valuations and click-through probabilities on the relative profitability of the two auction policies. Note that our theoretical model assumes a monotonically increasing relationship between advertiser valuation and click-through rates implying a high positive correlation between these two entities. Although, our empirical results support a positive relationship between advertiser valuations and click-through rates, it may be useful to examine the sensitivity of our theoretical results to variation in the level of positive correlation between advertiser

Table 3: Results from our simulation

Correlation	Average Profits under the highest bid policy	Average Profits under the highest profit policy	Fraction of instances when $\pi^{HB} > \pi^{HP}$
0.86	0.581(0.005529)	0.417(0.005265)	99.8%
0.5	0.524(0.004176)	0.412(0.003826)	95.0%
0.14	0.462(0.004232)	0.390(0.003732)	80.5%

The value in parenthesis indicates the standard deviation of the mean.

valuations and click-through rates.

We conduct a simulation to examine the effect of the magnitude of the positive correlation between advertiser valuations and click-through rates on the relative profitability of the two auction policies to the publisher. First, for 10 advertisers, we generate 1000 datasets of advertiser valuation and click-through rates at each of several pre-determined positive correlation levels between these variables. Then, we apply Definition 2 to determine for each simulated dataset, the publisher profit under each policy for 9 advertising slots. Table 3 presents the mean and standard deviation of the publisher profit across the 1000 simulations for each of the pre-selected positive correlation levels. The results show that the highest bid policy is indeed more likely to be profitable than the highest profit policy at all the simulated levels of positive correlation.[24] The above simulation results, in conjunction with evidence of a positive relationship between advertiser valuations and click-through rates as presented in this study, provide significant support to our theoretical results.

7 Conclusion

The growth of search-related online advertising has made it imperative for marketing and advertising managers to understand the implications of the institutional features of the search-advertising pricing process. These institutional features include a well-entrenched generalized second price auction pricing model for search-related advertising together with alternative bidder ranking and payment policies. In this paper, we analyze the auction model for search-related advertisements to derive the implications of two alternate auction policies: the highest bid policy and the highest profit policy. These two policy alternatives encompass both instruments (advertiser's bid per click and click-through rate) available to the publisher for basing his policy, and are also two policies that are observed in practice. We

[24]The superiority of the highest bid policy reduces slightly as the number of advertisers increases, consistent with Engelbrecht-Wiggans et al. (2007).

also conduct an empirical analysis using data from an anonymous search engine to validate our theoretical results.

Our main results are as follows. *Prima facie*, it may appear that the highest bid policy is likely to be inferior to the highest profit policy for a publisher since the former ignores the click through rates of ads. However, considering a product-market setting with differentiated advertisers in an analytical framework, we find that the highest bid policy is generally more profitable for the publisher than the highest profit policy. Only in the special case where all consumers are uninformed about the advertisers' differentiation, the highest bid and the highest profit policies yield the same profits. Sensitivity analysis shows that our model assumption of a positive correlation between advertiser valuation and click-through rates for a consumer is important for our results. However, we find that while our main result may be reversed when the correlation is negative, the highest bid policy can be more profitable even in cases of negative correlation. Our results are also sensitive to the presence of competition to search advertisements from organic search listings in a search engine context. Our analysis shows that under the positive correlation assumption, a search engine that displays organic as well as sponsored search listings can gain more from the highest bid policy if consumers prone to click the sponsored listings seldom click on the organic listings or if advertisers are well differentiated or if the number of informed consumers is not too high. These conditions appear plausible enough to favor the highest bid policy for search engines although more research is needed.

We undertake an empirical study to ascertain the external validity of our theoretical results. In this study, we empirically establish that advertiser valuations and the click through rates are positively correlated in a search advertising context. This finding validates the assumption of our theoretical analysis that is important for the result that the highest bid policy is generally more profitable for the publisher than the highest profit policy. However, our empirical conclusions suffer from the caveat that a positive correlation may not be sufficient to ensure superiority of the highest bid policy for a search engine due to the presence of organic listings, as noted earlier. Overall, our empirical analysis establishes indirect support for our theoretical results. Given Yahoo's recent switch to the highest profit policy, our results point to the need for further analysis and re-examination of the relative profitability of these two policies for the publisher. The results also raise the question of whether publishers should consider adopting different auction policies for different categories of keywords, depending on the relationship between advertiser valuations and click-through rates, in order to maximize profit.

Our analysis shows that an auction model does not extract the maximum possible profit from advertisers for the publisher. In this sense, an auction model is second-best to the publisher posting prices for search keywords. However, as we argue earlier, the information needs required to post prices for millions of keywords is quite steep, perhaps necessitating the auction model.

We conclude with a discussion of the limitations of our study and with some additional suggestions for further research. While many of our assumptions are similar to the ones used by Edelman et al. (2007), further research is needed about the validity of the equilibrium concept as a representation of advertiser bidding behavior. Another limitation is that our modeling of consumers is fairly simplistic in the sense that consumers may click one of the ads shown, generating a payoff for the advertiser. A more comprehensive consumer model employing a utility-maximizing framework in which consumers may click on multiple ads as part of a sequential search strategy can provide additional insights. Such a model may also include competition with organic listings in a search engine context, price competition between advertisers in the product market as in Dukes and Gal-Or (2003) and Iyer, Soberman, and Villas-Boas (2005), and may consider dependence between the click-through rates for slots and the sorting order of advertisers in equilibrium (Chen and He, 2007). A limitation of our model is that the proportion of informed consumers is exogenous. It would be useful to study a more general model in which advertisers influence the number of informed consumers through other forms of advertising (such as mass media advertising) or publicity. Our model assumes that the keyword used by the consumer reveals the consumer's product or brand preference (horizontal attributes) as well as quality preference (vertical attributes). It may be useful to study a model where advertisers are explicitly positioned on horizontal and vertical attributes. Another limitation of our analytical study is that we have focused on the single issue of comparing the highest-bid policy and the highest-profit policy in a generalized second-price auction framework. There are many more interesting issues to be addressed in this nascent area of search advertising. For example, what other auction mechanisms may be used to sell several ad slots to bidders, and how might these mechanisms compare with currently used mechanisms? Also interesting is the issue of the optimal mechanism to achieve the same purpose. On the empirical side, more research is needed with better and different data to compare the profitability of the two auction policies.

Technical Appendix

A Variables Used in the Single Slot Case

See Table 4.

B Proof of Proposition 1

Proposition 1 : *The following strategies and payoffs constitute a dominant strategy equilibrium for consumer type i:*

Table 4: Variables used in the single slot case

Vars	Explanation
v_i	Valuation for advertiser i.
λ	Fraction of consumers ex ante informed.
hline c_h	Click-through rate for advertiser i from consumer type i when consumers are ex ante informed.
c_l	Click-through rate for advertiser i from consumer type j when consumers are ex ante uninformed.
b_i^t	bid from the advertiser i for consumer type t.
p	Payment for the slot.
π_P	Publisher profits for the slot.
π_i	Advertiser i's profit.

Proof. Consider first the case of the highest bid policy. Here, advertiser i's profit from customer type i is:

$$\pi_i^i(b_i^i, b_j^i) = \begin{cases} 0 & \text{if } b_j^i > b_i^i \\ \frac{1}{2}(v_h - b_j^i)c_h & \text{if } b_j^i = b_i^i \\ (v_h - b_j^i)c_h & \text{if } b_j^i < b_i^i, \end{cases}$$

and advertiser j's profit from customer type i is:

$$\pi_j^i(b_j^i, b_i^i) = \begin{cases} 0 & \text{if } b_i^i > b_j^i \\ \frac{1}{2}(v_l - b_i^i)(\lambda c_l + (1-\lambda)c_h) & \text{if } b_i^i = b_j^i \\ (v_l - b_i^i)(\lambda c_l + (1-\lambda)c_h) & \text{if } b_i^i < b_j^i. \end{cases}$$

From the definition in Footnote 13 in the paper, $b_i^i < v_h$ is weakly dominated for advertiser i by a bid of v_h since advertiser i's profits are strictly lower in the former case when $b_i^i \leq b_j^i < v_h$. Similarly, $b_i^i > v_h$ is weakly dominated for advertiser i by a bid of v_h since advertiser i's profits are strictly lower in the former case when $b_i^i \geq b_j^i > v_h$. Thus, $b_i^i = v_h$ is the dominant strategy for advertiser i. Likewise, $b_j^i = v_l$ is the dominant strategy for advertiser j. Therefore, $b_i^i = v_h$ and $b_j^i = v_l$ represent the dominant strategy equilibrium. The equilibrium payments per click and profits follow from the definition of the highest bid policy.

Under the highest profit policy, the advertiser profits are as follows:

$$\pi_i^i(b_i^i, b_j^i) = \begin{cases} 0 & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) > b_i^i c_h \\ \frac{1}{2}(v_h c_h - b_j^i(\lambda c_l + (1-\lambda)c_h)) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) = b_i^i c_h \\ (v_h c_h - b_j^i(\lambda c_l + (1-\lambda)c_h)) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) < b_i^i c_h. \end{cases}$$

and

$$\pi_j^i(b_j^i, b_i^i) = \begin{cases} 0 & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) < b_i^i c_h \\ \frac{1}{2}(v_l(\lambda c_l + (1-\lambda)c_h) - b_i^i c_h) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) = b_i^i c_h \\ (v_l(\lambda c_l + (1-\lambda)c_h) - b_i^i c_h) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) > b_i^i c_h. \end{cases}$$

In this case, $b_i^i < v_h$ is weakly dominated for advertiser i by a bid of v_h since advertiser i's profits are strictly lower in the former case when $b_i^i \leq b_j^i[\lambda c_l + (1-\lambda)c_h]/c_h < v_h$. Similarly, $b_i^i > v_h$ is weakly dominated for advertiser i by a bid of v_h since advertiser i's profits are strictly lower in the former case when $b_i^i \geq b_j^i[\lambda c_l + (1-\lambda)c_h]/c_h > v_h$. Thus, $b_i^i = v_h$ is the dominant strategy for advertiser i. Likewise, $b_j^i < v_l$ is is weakly dominated for advertiser j by a bid of v_l since advertiser j's profits are strictly lower in the former case when $b_j^i \leq b_i^i c_h/[\lambda c_l + (1-\lambda)c_h] < v_l$. Similarly, $b_j^i > v_l$ is weakly dominated for advertiser j by a bid of v_l since advertiser j's profits are strictly lower in the former case when $b_j^i \geq b_i^i c_h/[\lambda c_l + (1-\lambda)c_h] > v_l$. Therefore, $b_i^i = v_h$ and $b_j^i = v_l$ represent the dominant strategy equilibrium. Given the equilibrium strategies, the payments per click and the profits follow from the definition of the auction policy. ■

C Competition between Search Ads and Organic Search Listings

We consider an extension of the base model to incorporate competition from organic search listings. We focus only on consumer type i without loss of generality. For the sponsored listing, we retain the original model characterization. The organic search listing is modeled as follows. Assume that the per-click valuation for a firm on the organic listing side is the same as that on the sponsored listing. If firm i is the winner on the sponsored listing, without loss of generality, let the organic listing receive a base click-through rate of 0 for both firms. Note that in the original model, if firm j wins the sponsored listing instead of firm i, there is a total decrease of $\lambda(c_h - c_l)$ in the click-through rate on the sponsored listing side. This decrease in the click-through rate can be justified on the basis that informed consumers do not perceive a match. We now assume that a fraction η of this deficit in the number of clicks revert to advertiser i's organic listings. In other words, when firm j appears in the sponsored listing, firm i receives clicks at the rate of $\eta\lambda(c_h - c_l)$ on the organic listing, and does not pay anything to the publisher. Note that $\eta = 0$ corresponds to our original model.

In this extended base model, the equilibrium bids are the same under both policies: $b_i^i = v_h\big(1 - \eta\lambda(1 - \frac{c_l}{c_h})\big)$; $b_j^i = v_l$. The highest bid policy generates higher publisher profits when $\eta\lambda < \frac{v_h - v_l}{v_h}\frac{c_h}{c_h - c_l}$ and vice-versa.

Proof. Consider first the case of the highest bid policy. Here, advertiser i's profit from customer type i is:

$$\pi_i^i(b_i^i, b_j^i) = \begin{cases} \eta\lambda(c_h - c_l)v_h & \text{if } b_j^i > b_i^i \\ \frac{1}{2}\big((v_h - b_j^i)c_h + \eta\lambda(c_h - c_l)v_h\big) & \text{if } b_j^i = b_i^i \\ (v_h - b_j^i)c_h & \text{if } b_j^i < b_i^i, \end{cases}$$

and advertiser j's profit from customer type i is:

$$\pi_j^i(b_j^i, b_i^i) = \begin{cases} 0 & \text{if } b_i^i > b_j^i \\ \frac{1}{2}(v_l - b_i^i)(\lambda c_l + (1-\lambda)c_h)) & \text{if } b_i^i = b_j^i \\ (v_l - b_i^i)(\lambda c_l + (1-\lambda)c_h) & \text{if } b_i^i < b_j^i. \end{cases}$$

The dominant strategy for advertiser i is obtained by comparing the payoffs i receives from winning and losing the slot. That is, so long as b_j^i is such that $(v_h - b_j^i)c_h \geq \eta\lambda(c_h - c_l)v_h$, advertiser i desires the slot. The condition corresponds to $b_j^i \leq v_h\big(1 - \eta\lambda(1 - \frac{c_l}{c_h})\big)$. Therefore, the dominant strategy bid for advertiser i is $v_h\big(1 - \eta\lambda(1 - \frac{c_l}{c_h})\big)$. Prices both below and above this value are weakly dominated. Similarly, advertiser j's dominant strategy is to bid v_l. The publisher profits under this policy are as follows

$$\pi_P^{HB} = \begin{cases} v_l c_h & \text{if } v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) > v_l \\ \frac{1}{2}v_l\big(c_h + \lambda c_l + (1-\lambda)c_h)\big) & \text{if } v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) = v_l \\ (\lambda c_l + (1-\lambda)c_h)\frac{v_h}{c_h}(1 - \eta\lambda(1 - \frac{c_l}{c_h})) & \text{if } v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) < v_l, \end{cases}$$

Under the highest profit policy, the advertiser profits are as follows:

$$\pi_i^i(b_i^i, b_j^i) = \begin{cases} \eta\lambda(c_h - c_l)v_h & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) > b_i^i c_h \\ \frac{1}{2}\big(v_h c_h - b_i^i c_h + \eta\lambda(c_h - c_l)v_h\big) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) = b_i^i c_h \\ (v_h c_h - b_j^i(\lambda c_l + (1-\lambda)c_h)) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) < b_i^i c_h, \end{cases}$$

and

$$\pi_j^i(b_j^i, b_i^i) = \begin{cases} 0 & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) < b_i^i c_h \\ \frac{1}{2}(v_l(\lambda c_l + (1-\lambda)c_h) - b_i^i c_h) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) = b_i^i c_h \\ (v_l(\lambda c_l + (1-\lambda)c_h) - b_i^i c_h) & \text{if } b_j^i(\lambda c_l + (1-\lambda)c_h) > b_i^i c_h. \end{cases}$$

In this case, the advertiser will have an incentive to increase his bid so long as $v_h c_h - b_j^i(\lambda c_l + (1-\lambda)c_h) \geq v_h \eta\lambda(c_h - c_l)$. This condition translates into a dominant strategy bid of $v_h\big(1 - \eta\lambda(1 - \frac{c_l}{c_h})\big)$. Other prices are weakly dominated. Likewise, we can determine that the dominant strategy for advertiser j to bid v_l. The publisher profits are:

$$\pi_P^{HP} = \begin{cases} v_l(\lambda c_l + (1-\lambda)c_h) & \text{if } v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) \geq v_l(1 - \lambda(1 - \frac{c_l}{c_h})) \\ v_h(c_h - \lambda(c_h - c_l)) & \text{if } v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) < v_l(1 - \lambda(1 - \frac{c_l}{c_h})), \end{cases}$$

Note that the condition $v_h(1 - \eta\lambda(1 - \frac{c_l}{c_h})) < v_l(1 - \lambda(1 - \frac{c_l}{c_h}))$ is never valid since $v_l < v_h$.

The publisher profit comparison are accomplished under the following two conditions:

- $v_l c_h \leq v_h(c_h + (c_l - c_h)\eta\lambda)$: $\pi_P^{HB} > \pi_P^{HP}$. This result is obvious.
- $v_l c_h > v_h(c_h + (c_l - c_h)\eta\lambda)$: Multiply both sides of the condition by $(\lambda c_l + (1-\lambda)c_h)$ and it leads to $\pi_P^{HB} \leq \pi_P^{HP}$.

■

From these results, we can conclude that the highest-bid policy dominates the highest profit policy from publisher profit standpoint if $\eta\lambda \leq \frac{v_h - v_l}{v_h}\frac{c_h}{c_h - c_l}$, and vice-versa.

D Variables Used in the Multiple Slot Case

See Table 5.

E Additional Discussion of the Equilibrium Concept for Multiple Slots

We begin with the definition of a locally envy-free equilibrium based on Edelman et al. (2007). Subsequently, we motivate this equilibrium as well as the locally dominant envy-free equilibrium through discussions and an example. As discussed in the paper, in a locally envy-free equilibrium, an advertiser with a given rank cannot gain by exchanging bids (expected profit levels) with the advertiser ranked immediately above him under the highest bid policy (highest profit policy). This requirement leads to the following definition of a locally envy-free equilibrium. The locally dominant envy-free equilibrium as defined in the paper is a refinement of this equilibrium.

Definition 5 (Based on Edelman et al. (2007)) *An equilibrium of a simultaneous move game induced by the GSP auction is locally envy-free under the following conditions:*

Table 5: Variables used in the multiple slot generalization.

Vars	Explanation
M	Number of Advertisers.
K	Number of slots.
v_i	Valuation for advertiser i.
λ	Fraction of consumers ex ante informed.
c_i	Click-through rate for advertiser i if consumers are ex ante informed.
χ_i	Advertiser-specific click through rate; accounts for consumer knowledge.
α_k	"Inherent" click through rate factor for slot k.
γ_i^k	Effective click through rate for an ad from advertiser i in slot k.
$\pi_{P,k}$	Publisher profits for slot k
b_i	bid from advertiser i.
p_k	Payment for slot k.
$g(k)$	Identity of the advertiser in slot k.
Superscripts HB and HP refer to the highest bid and the highest profit policies.	

(i) Under the highest bid policy, for any rank $j \leq K+1$,

$$\alpha_j(v_{g(j)} - p_j^{HB}) \geq \alpha_{j-1}(v_{g(j)} - p_{j-1}^{HB}) \tag{E-1}$$

and for rank $j > K+1$,

$$\alpha_j(v_{g(j)} - p_j^{HB}) \geq \alpha_K(v_{g(j)} - p_K^{HB}); \tag{E-2}$$

(ii) Under the highest profit policy, for any $j \leq K+1$,

$$\alpha_j(v_{g(j)} - p_j^{HP}) \geq \alpha_{j-1}(v_{g(j)} - p_{j-1}^{HP}(\frac{\chi_{g(j-1)}}{\chi_{g(j)}})) \tag{E-3}$$

and for any rank $j > K+1$,

$$\alpha_j(v_{g(j)} - p_j^{HP}) \geq \alpha_K(v_{g(j)} - p_K^{HP}(\frac{\chi_{g(K)}}{\chi_{g(j)}})). \tag{E-4}$$

Note that there are two sets of conditions for the locally envy-free equilibrium under each policy: one for $j \leq K+1$, and a second for $j > K+1$. The rationale is that for $j > K + 1$, exchange of positions with the advertiser ranked immediately above an advertiser produces no change in the payoff since both payoffs are zero. Thus, constraints E-1 or E-3 have no bite for these positions and have to be modified suitably. It seems reasonable that advertisers with a rank j such that $j > K + 1$, would bid to displace the advertiser in the last slot K rather than the advertiser ranked immediately above them. Thus, the locally envy-free condition for $j > K + 1$ ensures that the advertiser in the jth position does not gain by exchanging positions with the advertiser assigned to slot K. Example 2 below provides additional motivation for this condition. We begin with Example 1, which motivates the locally envy-free equilibrium and the locally dominant envy-free equilibrium in a case where the number of bidders is exactly one more than the number of slots.

Example 1: Consider three advertisers 1, 2, and 3 competing for two ad slots in a GSP auction, and having valuations of \$1, \$2, and \$3 respectively for a consumer's click. All advertisers have the same click-through rates, and these rates (probabilities) for the first and second slot are 0.05 and 0.04 respectively. Assume that the publisher uses a highest bid-per-click policy. In spite of the second-price nature of the auction, each advertiser bidding his valuation cannot be a long-run equilibrium for our example, and it is consistent with Edelman et al. (2007).[25] We use the ordered triple $\{b_1, b_2, b_3\}$ to denote the bids of advertisers 1, 2, and 3 respectively. To show that requiring that bids only be a best response to each other results in some non-intuitive equilibria, consider the bid triple $\{\$1, \$1.9, \$1.1\}$. In this case, advertisers 2 and 3 get the first and second slot for a payment of $\$1.1$ and $\$1$ per click respectively. Their expected profits per consumer are $\$0.045$ and $\$0.08$ respectively. It can be checked that the above bids are best responses to each other. However, it would appear that advertiser 3 could gain by increasing his bid in order to raise advertiser 2's cost per click, with a view to dislodging that advertiser from the top spot. For example, advertiser 3 can increase his bid to a little over $\$1.2$ to reduce advertiser 2's profit to a little less than $\$0.04$. This will, in turn, induce advertiser 2 to to concede the first spot to advertiser 3 by bidding $\$1.2$, since he will then make a profit of $\$0.04$. By thus getting the first slot, advertiser 3 has increased his profit to $\$0.09$. Therefore, the bids $\{\$1, \$1.9, \$1.1\}$ cannot be considered a plausible long-run stationary equilibrium of the simultaneous move game induced by the GSP auction. It can be checked that these bids do not satisfy the criterion of being locally envy-free although they constitute a Nash equilibrium. On the other hand, the bids $\{\$1, \$1.2, \$1.2 + \varepsilon\}$ with any $\varepsilon > 0$ form a locally envy-free equilibrium.

[25] Suppose advertisers bid their valuations in our example. Then, 3 and 2 get the first and second slots respectively and make payments per click of \$2 and \$1 respectively. Their expected profits per consumer are \$0.05 and \$0.04 respectively. However, this cannot be a long-run stationary equilibrium as advertiser 3 can gain by underbidding advertiser 2 to get the second slot and an expected profit of \$0.08.

Note that the highest bid does not affect any advertiser's payment and is thus not uniquely determined.

As discussed above, a locally envy-free equilibrium ensures that an advertiser does not want to exchange positions with the next higher-ranked advertiser. Another intuitive interpretation of the locally envy-free equilibrium is that the advertiser, who gets position k, competes for position $k-1$ by making a bid such that his potential payment for position $k-1$ exceeds his payment for position k by at least the marginal value for the extra clicks he would receive in position $k-1$. This can be seen by rearranging Equation E-1 for the highest bid policy: $\chi_{g(k)}(\alpha_{k-1}p_{k-1} - \alpha_k p_k) = v_{g(k)}\chi_{g(k)}(\alpha_{k-1} - \alpha_k)$; and equation (E-3) for the highest profit policy: $\alpha_{k-1}\chi_{g(k-1)}p_{k-1} - \alpha_k\chi_{g(k)}p_k = v_{g(k)}\chi_{g(k)}(\alpha_{k-1} - \alpha_k)$. (See also footnote 17, p. 249 of Edelman et al., 2007.) Equations E-1 and E-3 still allows for multiple envy-free equilibria, but we argue that the most reasonable stable long-run equilibrium is the locally envy-free equilibrium in which the equations are satisfied with equality. One reason is that a strict inequality in both the equations imply that the advertiser i occupying position k is bidding more than the marginal value for the extra clicks obtainable from position $k-1$. Under these circumstances, advertiser i could earn less if he wins position k due to advertiser j (occupying position $k-1$) slightly underbidding him. Thus, it would appear that advertiser i's bid, when Equation E-1 or Equation E-3 (as applicable) is an equality, is more consistent with stability than other locally envy-free equilibria. Indeed, under each policy, advertiser i's bid corresponding to the equality condition weakly dominates other locally envy-free bids in his "local" competition with advertiser j. Thus, we call advertiser i's bid that satisfy equality conditions in Equations E-1 and E-3 (and as defined in the paper as Definition 2 on Page 17) as a "locally dominant envy-free" equilibrium. In our example, both $\{\$1, \$1.2, \$3\}$ and $\{\$1, \$1.3, \$3\}$ are locally envy-free equilibria while only the former set of bids is a locally dominant envy-free equilibrium (Here we have set the highest bid equal to the valuation).

The following example illustrates the need for the additional condition for positions $j > K+1$ in the definitions of the locally envy-free equilibrium and the locally dominant envy-free equilibrium.

Example 2: To the above example, add a fourth advertiser, 0, with a valuation of $\$0.5$ for a consumer's click. We use the ordered quartet $\{b_0, b_1, b_2, b_3\}$ to denote the bids of advertisers 0, 1, 2, and 3 respectively. If condition (E-1) applied to all positions j, then the bid quartet $\{\$0.5, \$0.2, \$0.8, \$3\}$ is a locally envy-free equilibrium. However, this equilibrium is not intuitively appealing because it is advertiser 0, and not advertiser 1, who is competing for the last slot with advertiser 2, although advertiser 1 has a higher valuation than advertiser 0. This equilibrium is not ruled out by condition (E-1) because advertiser 0 makes a zero payoff regardless of whether he gets the third or last rank. However, this non-intuitive equilibrium is ruled out by condition (E-2).

The following example extends Example 1 by allowing for differences in click-through rates across advertisers. This example illustrates the locally dominant envy-free equilibrium under the highest profit policy.

Example 3: We modify Example 1 to assume different click-through rates for the three advertisers. Specifically, advertisers 1, 2, and 3 have click-through rates of $1/6$, $1/3$, and $1/2$ respectively. By Proposition 3, advertisers 3 and 2 are assigned to the first and second slot respectively in the unique locally dominant envy-free equilibrium. The equilibrium payments per click for the first and second slot are $\$0.533$ and $\$0.5$ respectively. Using the rules of the highest profit policy, these payments correspond to the bid triple $\{1, 0.8, 3\}$, where we have arbitrarily set the bid of the highest-ranked advertiser (advertiser 3) to his valuation, since his bid is not uniquely determined. Note that while payments increase for the more attractive slots even under the highest profit policy, the winning bids need not be higher for such slots under this policy.

F Proof of Proposition 3

Proposition 3:

(i) *In equilibrium, the advertiser's payment per click for slot k, such that $1 \leq k \leq K$, is as follows:*

- *under the highest bid policy,* $p_k^{HB} = \sum_{j=k}^{K} v_{j+1} \left(\frac{\alpha_j - \alpha_{j+1}}{\alpha_k}\right)$, *and*
- *under the highest profit policy,* $p_k^{HP} = \sum_{j=k}^{K} \left(\frac{\chi_{j+1}}{\chi_k}\right) v_{j+1} \left(\frac{\alpha_j - \alpha_{j+1}}{\alpha_k}\right)$.

(ii) *For every $1 \leq k \leq K$, $p_k^{HB} \geq p_k^{HP}$ and $\pi_{P,k}^{HB} \geq \pi_{P,k}^{HP}$. The payments and the profits are equal only when $\lambda = 0$.*

Proof. ***Part (i).*** Recall that $g(j)$ represents the index number of the advertiser with the j^{th} highest rank and, in addition, we use $g^{-1}(i)$ to refer to the rank of advertiser i.

Proof for Highest Bid Policy

Step 1: We claim that $v_{g(j)} \geq v_{g(j+1)}$ for $1 \leq j \leq K$. First, the requirement of advertisers' bids being best response to others' bids implies that $\alpha_j(v_{g(j)} - p_j) \geq \alpha_{j+1}(v_{g(j)} - p_{j+1})$. In other words, the advertiser assigned to position j does not gain by moving one position down. Second, from the requirement of a locally dominant envy-free equilibrium, we have $\alpha_{j+1}(v_{g(j+1)} - p_{j+1}) = \alpha_j(v_{g(j+1)} - p_j)$. Combining the above two inequalities yields $(\alpha_j - \alpha_{j+1})v_{g(j)} \geq (\alpha_j - \alpha_{j+1})v_{g(j+1)}$, which implies $v_{g(j)} \geq v_{g(j+1)}$ as $\alpha_j > \alpha_{j+1}$.

Step 2: We claim that there is no advertiser whose index is $i \leq K+1$, but whose rank $g^{-1}(i) > K + 1$. Suppose to the contrary that there are such advertisers in equilibrium. Let i' be the smallest index among such advertisers, i.e. $i' = \min\{i | i \leq K + 1$ and $g^{-1}(i) > K + 1\}$. Then, applying the definition of the locally dominant envy-free equilibrium for advertiser $g(K + 1)$, we have $p_K^{HB} = v_{g(K+1)}$. Further, applying the definition of the locally dominant envy-free equilibrium for advertiser i', we have $v_{i'} \leq p_K^{HB}$. From the last two inequalities, we have $v_{i'} \leq v_{g(K+1)}$. However, since $v_{g(j)} \geq v_{g(j+1)}$ for $1 \leq j \leq K$ by Step 1, we must have $g(K + 1) = \max\{i | g^{-1}(i) \leq K + 1\}$. Since at least one advertiser with index $i \leq K + 1$ has rank $g^{-1}(i) > K + 1$ by assumption, we must have $g(K + 1) > i'$. Then, from our model assumption about advertiser valuation, $v_{g(K+1)} < v_{i'}$ which contradicts the earlier conclusion that $v_{i'} \leq v_{g(K+1)}$. This proves our claim for this step.

Step 3: The combined import of Steps 1 and 2 is that $g(j) = j$ (and that $v_{g(j)} = v_j$) $\forall\, j$ such that $1 \leq j \leq K + 1$. Thus, the locally dominant envy-free equilibrium results in an assortative match (see also Edelman et al. (2007)). To determine the equilibrium payments as presented in the Proposition for slots k, $1 \leq k \leq K$, set $g(k) = k$ in Definition 2, and solve recursively. The recursive solution establishes that the equilibrium payments are unique.

Step 4: Finally, we verify in a couple of steps that the locally dominant envy-free equilibrium thus computed satisfies the basic Nash equilibrium condition that no advertiser wishes to move to a higher or lower position by changing his bid. For positions $j > K + 1$, no advertiser $g(j)$ wants to move to slot K or higher since $p_K^{HB} = v_{K+1} > v_{g(j)}$, and because p_k^{HB} increases as k decreases (as can be seen by inspection). For the same reason, the advertiser assigned to position $K + 1$ cannot increase his payoff by moving up to a higher-ranked position. Similarly, an advertiser assigned to any slot k does not want to move to a position $j > K+1$ (and make zero profit), since he currently makes a positive profit. To see this last point, the advertiser's profit π_k^{HB} an be written as follows, obtaining the third equality in the expression by rearranging terms:

$$\begin{aligned} \pi_k^{HB} &= \chi_k \alpha_k (v_k - p_k^{HB}) = \chi_k [\alpha_k v_k - \sum_{j=k}^{K} v_{j+1} (\alpha_j - \alpha_{j+1})] \\ &= \chi_k [\sum_{j=k}^{K} \alpha_j (v_j - v_{j+1})] > 0 \end{aligned} \quad \text{(F-5)}$$

Step 5: Next, consider the advertiser in a slot k who considers moving to slot $k + m$ for $1 \leq m \leq K - k$ by bidding slightly below the advertiser in slot $k + m$.

His payoff, π_k^d, from doing so is given by the following expression:

$$\pi_k^d = \chi_k[\alpha_{k+m}(v_k - p_{k+m}^{HB}) - \alpha_k(v_k - p_k^{HB})] = \alpha_{k+m}\chi_k(v_k - p_{k+m}^{HB}) - \alpha_k\chi_k(v_k - p_k^{HB})$$

The following chain of equalities and inequalities shows $\alpha_k\chi_k(v_k - p_k^{HB}) > \alpha_{k+m}\chi_k(v_k - p_{k+m}^{HB})$ implying that $\pi_k^d < 0$. Thus, the advertiser in slot k does not gain by moving to slot $k+m$.

$$\begin{aligned}
\alpha_k(v_k - p_k^{HP}) &> \alpha_{k+m}(v_k - v_{k+1}) + \alpha_k(v_{k+1} - p_k^{HB}) && \text{(F-6)}\\
&= \alpha_{k+m}(v_k - v_{k+1}) + \alpha_{k+1}(v_{k+1} - p_{k+1}^{HB}) && \text{(F-7)}\\
&> \alpha_{k+m}(v_k - v_{k+2}) + \alpha_{k+1}(v_{k+2} - p_{k+1}^{HB}) && \text{(F-8)}\\
&> \ldots\\
&= \alpha_{k+m}(v_k - v_{k+m-1}) + \alpha_{k+m-1}(v_{k+m-1} - p_{k+m-1}^{HB})\\
&> \alpha_{k+m}(v_k - v_{k+m}) + \alpha_{k+m-1}(v_{k+m} - p_{k+m-1}^{HB})\\
&= \alpha_{k+m}(v_k - v_{k+m}) + \alpha_{k+m}(v_{k+m} - p_{k+m}^{HB})\\
&= \alpha_{k+m}(v_k - p_{k+m}^{HB})
\end{aligned}$$

In the above chain, inequality (F-6) is due to the result from Step 3 that $v_j > v_{j+1}$ $\forall\ j$ such that $1 \leq j \leq K+1$, and because $\alpha_k > \alpha_{k+m}$. The equality (F-7) follows from the locally dominant envy-free equilibrium condition. The inequality (F-8) is obtained because $v_{k+1} > v_{k+2}$, $\alpha_{k+1} > \alpha_{k+m}$, and due to:

$$\begin{aligned}
&\alpha_{k+m}(v_k - v_{k+1}) + \alpha_{k+1}(v_{k+1} - p_{k+1}^{HB})\\
&> \alpha_{k+m}[(v_k - v_{k+1}) + (v_{k+1} - v_{k+2})] + \alpha_{k+1}(v_{k+2} - p_{k+1}^{HB})\\
&= \alpha_{k+m}(v_k - v_{k+2}) + \alpha_{k+1}(v_{k+2} - p_{k+1}^{HB})
\end{aligned}$$

The rest of the chain is obtained similarly. Now, consider the payoff π_k^u to advertiser in slot k ($k > 1$) from moving up to slot $k - m$, where $1 \leq m \leq k-1$, by bidding slightly above the advertiser assigned to that slot. Note that by doing so, the advertiser makes the payment assigned to slot $k - m - 1$.[26] π_k^u is given by the following expression:

$$\begin{aligned}
\pi_k^u &= \chi_k[\alpha_{k-m}(v_k - p_{k-m-1}^{HB}) - \alpha_k(v_k - p_k^{HB})]\\
&< \chi_k[\alpha_{k-m}(v_k - p_{k-m}^{HB}) - \alpha_k(v_k - p_k^{HB})]
\end{aligned}$$

[26] When k - m = 1, the advertiser should make the payment assigned to slot 0, which is not defined. However, this really means that the advertiser has to beat the highest-ranked bid (which is the payment to the hypothetical slot 0) to win the first lot. Since the highest-ranked bid is not uniquely determined in equilibrium, we define p_0^{HB} = p_1^{HB} + ε, for some arbitrary $\varepsilon > 0$, and this definition is without loss of generality.

In obtaining the inequality, in the above expression, we replace p^{HB}_{k-m-1} by p^{HB}_{k-m}. $p^{HB}_{k-m-1} > p^{HB}_{k-m}$, thus justifying the above inequality. The following chain of equalities and inequalities shows $\alpha_k(v_k - p^{HB}_k) > \alpha_{k-m}(v_k - p^{HB}_{k-m})$ implying that $\pi^u_k < 0$. Thus, the advertiser in slot k does not gain by moving to slot $k - m$.

$$
\begin{aligned}
\alpha_k(v_k - p^{HB}_k) &= \alpha_{k-1}(v_k - p^{HB}_{k-1}) && \text{(F-9)}\\
&> \alpha_{k-m}(v_k - v_{k-1}) + \alpha_{k-1}(v_{k-1} - p^{HB}_{k-1}) && \text{(F-10)}\\
&= \alpha_{k-m}(v_k - v_{k-1}) + \alpha_{k-2}(v_{k-1} - p^{HB}_{k-2}) && \text{(F-11)}\\
&> \alpha_{k-m}(v_k - v_{k-2}) + \alpha_{k-2}(v_{k-2} - p^{HB}_{k-2}) && \text{(F-12)}\\
&= \ldots\\
&= \alpha_{k-m}(v_k - v_{k-m+2}) + \alpha_{k-m+1}(v_{k-m+2} - p^{HB}_{k-m+1})\\
&> \alpha_{k-m}(v_k - v_{k-m+1}) + \alpha_{k-m+1}(v_{k-m+1} - p^{HB}_{k-m+1})\\
&= \alpha_{k-m}(v_k - v_{k-m+1}) + \alpha_{k-m}(v_{k-m+1} - p^{HB}_{k-m})\\
&> \alpha_{k-m}(v_k - v_{k-m}) + \alpha_{k-m}(v_{k-m} - p^{HB}_{k-m})\\
&= \alpha_{k-m}(v_k - p^{HB}_{k-m})
\end{aligned}
$$

In the above chain, equality (F-9) obtains because of the locally dominant envy-free equilibrium condition. Inequality (F-10) follows from the result in Step 3 that $v_j > v_{j+1}\ \forall\ j$ such that $1 \leq j \leq K+1$, and because $\alpha_{k-1} > \alpha_{k-m}$. The equality (F-11) again follows from the locally dominant envy-free equilibrium condition. The inequality (F-12) is from $v_{k+1} > v_{k+2}$, $\alpha_{k-2} < \alpha_{k-m}$, and because of the following:

$$
\begin{aligned}
&\alpha_{k-m}(v_k - v_{k-1}) + \alpha_{k-2}(v_{k-1} - p^{HB}_{k-2})\\
&> \alpha_{k-m}[(v_k - v_{k-1}) + (v_{k-1} - v_{k-2})] + \alpha_{k-2}(v_{k-2} - p^{HP}_{k-2})\\
&= \alpha_{k-m}(v_k - v_{k-2}) + \alpha_{k-2}(v_{k-2} - p^{HP}_{k-2})
\end{aligned}
$$

The rest of the chain is obtained in an analogous fashion. Thus, we have established the equilibrium presented in the Proposition. It can be further shown that the assignment $g(j) = j\ \forall\ j$ such that $1 \leq j \leq K+1$ is a stable assignment in the sense that no advertiser wants to exchange bids with another advertiser (cf. Edelman et al. (2007)). Our proof above showing that an advertiser does not want to move to a lower-ranked slot establishes that he does not want to exchange bids with an advertiser assigned to a lower-ranked slot. Similarly, in showing that the advertiser assigned to slot k does not want to position $k - m$, we established that $\alpha_{k-m}(v_k - p^{HB}_{k-m}) - \alpha_k(v_k - p^{HB}_k) < 0$. This establishes that the advertiser in slot k does not gain by exchanging bids with the advertiser in slot $k - m$. Thus, the equilibrium presented is a stable assignment.

Proof for Highest Profit Policy

Step 1: We claim that $v_{g(j)}\chi_{g(j)} \geq v_{g(j+1)}\chi_{g(j+1)}$ for $1 \leq j \leq K$. First, the requirement of advertisers' bids being best response to others' bids implies that $\alpha_j(v_{g(j)} - p_j) \geq \alpha_{j+1}(v_{g(j)} - p_{j+1}\chi_{g(j+1)}/\chi_{g(j)})$. In other words, the advertiser assigned to position j does not gain by moving one position down. Second, from the requirement of a locally dominant envy-free equilibrium, we have $\alpha_{j+1}(v_{g(j+1)} - p_{j+1}) = \alpha_j(v_{g(j+1)} - p_j\chi_{g(j)}/\chi_{g(j+1)})$. Combining the above two inequalities yields $(\alpha_j - \alpha_{j+1})v_{g(j)}\chi_{g(j)} \geq (\alpha_j - \alpha_{j+1})v_{g(j+1)}\chi_{g(j+1)}$, which implies $v_{g(j)}\chi_{g(j)} \geq v_{g(j+1)}\chi_{g(j+1)}$ as $\alpha_j > \alpha_{j+1}$.

Step 2: We claim that there is no advertiser whose index is $i \leq K + 1$, but whose rank $g^{-1}(i) > K + 1$. Suppose to the contrary that there are such advertisers in equilibrium. Let i' be the smallest index among such advertisers, i.e. $i' = \min\{i|i \leq K + 1$ and $g^{-1}(i) > K + 1\}$. Then, applying the definition of the locally dominant envy-free equilibrium for advertiser $g(K + 1)$, we have $p_K^{HP} = v_{g(K+1)}\chi_{g(K+1)}/\chi_{g(K)}$. Further, applying the definition of the locally dominant envy-free equilibrium for advertiser i', we have $v_{i'} \leq p_K^{HP}\chi_{g(K)}/\chi_{i'}$. From the last two inequalities, we have $v_{i'}\chi_{i'} \leq v_{g(K+1)}\chi_{g(K+1)}$. However, since $v_{g(j)}$ $\chi_{g(j)} \geq v_{g(j+1)}\chi_{g(j+1)}$ for $1 \leq j \leq K$ by Step 1, we must have $g(K + 1) = \max\{i|g^{-1}(i) \leq K+1\}$. Since at least one advertiser with index $i \leq K+1$ has rank $g^{-1}(i) > K + 1$ by assumption, we must have $g(K + 1) > i'$. The, from our model assumption about advertiser valuation and click-through rates, $v_{g(K+1)}\chi_{g(K+1)} < v_{i'}\chi_{i'}$ which contradicts the earlier conclusion that $v_{i'}\chi_{i'} \leq v_{g(K+1)}\chi_{g(K+1)}$. This proves our claim for this step.

Step 3: The combined import of Steps 1 and 2 is that $g(j) = j$ (and that $v_{g(j)} = v_j$) $\forall\, j$ such that $1 \leq j \leq K+1$. Thus, the locally dominant envy-free equilibrium results in an assortative match in the sense that $v_j > v_{j+1}$(Also, $v_j\chi_j > v_{j+1}\chi_{j+1}$) . To determine the equilibrium payments as presented in the proposition for slots k, $1 \leq k \leq K$, set $g(k) = k$ in Definition 2, and solve recursively. The recursive solution establishes that the equilibrium payments are unique.

Step 4: Finally, we verify in a couple of steps that the locally dominant envy-free equilibrium thus computed satisfies the basic Nash equilibrium condition that no advertiser wishes to move to a higher or lower position by changing his bid. For positions $j > K + 1$, an advertiser $g(j)$ wanting to move up to slot k should bid such that his expected profit to the publisher slightly exceeds that of the advertiser assigned to slot k. Consider, first, the payoff to advertiser $g(j)$ from moving to position K. His payoff, π_k^u, in doing so is given by the following expression:

$$\pi_k^u = \alpha_K \chi_{g(j)} (v_{g(j)} - p_{K-1}^{HP} \chi_{K-1} / \chi_{g(j)}) = \alpha_K (v_{g(j)} \chi_{g(j)} - p_{K-1}^{HP} \chi_{K-1})$$
$$= \alpha_K \{ v_{g(j)} \chi_{g(j)} - [v_K \chi_K (\alpha_{K-1} - \alpha_K) + v_{K+1} \chi_{K+1} \alpha_K] / \alpha_{K-1} \}$$
$$= \alpha_K \{ v_{g(j)} \chi_{g(j)} - [\alpha_K (v_{K+1} \chi_{K+1} - v_K \chi_K) + \alpha_{K-1} v_K \chi_K] / \alpha_{K-1} \}$$
$$< \alpha_K \{ v_{g(j)} \chi_{g(j)} - [\alpha_{K-1} (v_{K+1} \chi_{K+1} - v_K \chi_K) + \alpha_{K-1} v_K \chi_K] / \alpha_{K-1} \} \tag{F-13}$$
$$= \alpha_K (v_{g(j)} \chi_{g(j)} - v_{K+1} \chi_{K+1}) < 0. \tag{F-14}$$

Note that the expression for π_k^u incorporates the payment for the $K - 1^{th}$ slot, since the payment for this slot depends on the expected profit of the advertiser assigned to slot K, and this profit is also the level to be exceeded by advertiser $g(j)$ to gain slot K. The inequality in (F-13) is obtained by replacing α_K by α_{K-1}, and follows because $v_{K+1}\chi_{K+1} < v_K \chi_K$ and because $\alpha_{K-1} > \alpha_K$. The inequality in (F-14) is because $v_{g(j)}\chi_{g(j)} < v_{K+1}\chi_{K+1}$ by Step 3. Next, it can be seen by inspection that p_k^{HP} increases as k decreases. Therefore, advertiser $g(j)$ cannot gain by moving up to any higher-ranked slot $k < K + 1$. Using similar reasoning, we can show that the advertiser assigned to position $K + 1$ cannot increase his payoff by moving up to a higher-ranked position. Similarly, an advertiser assigned to any slot k does not want to move to a position $j > K + 1$ (and make zero profit), since he currently makes a positive profit. To see this last point, the advertiser's profit π_k^{HP}can be written as follows by rearranging terms:

$$\pi_k^{HP} = \chi_k \alpha_k (v_k - p_k^{HP}) = \alpha_k \chi_k v_k - \chi_k [\sum_{j=k}^{K} \left(\frac{\chi_{j+1}}{\chi_k} \right) v_{j+1} (\alpha_j - \alpha_{j+1})]$$
$$= \sum_{j=k}^{K} \alpha_j (v_j \chi_j - v_{j+1} \chi_{j+1}) > 0$$

Step 5: Next, consider the advertiser in a slot k who considers moving to slot $k+m$ for $1 \leq m \leq K - k$ by bidding lower such that his promised expected profit to the publisher is slightly below the advertiser in slot $k + m$. From the following chain of equalities and inequalities it can be seen that the advertiser in slot k does not gain by moving to slot $k + m$. His payoff, π_k^d, from doing so is given by the following expression:

$$\pi_k^d = \chi_k \alpha_{k+m} (v_k - p_{k+m}^{HP} \chi_{k+m} / \chi_k) - \chi_k \alpha_k (v_k - p_k^{HP})$$
$$= \alpha_{k+m} (\chi_k v_k - p_{k+m}^{HP} \chi_{k+m}) - \alpha_k (\chi_k v_k - \chi_k p_k^{HP})$$

The following chain of equalities and inequalities will show that $\alpha_k(\chi_k v_k - \chi_k p_k^{HP}) > \alpha_{k+m}(\chi_k v_k - \chi_{k+m} p_{k+m}^{HP})$ implying that $\pi_k^d < 0$. Thus, the advertiser in slot k does not gain by moving to slot $k+m$.

$$
\begin{aligned}
\alpha_k(\chi_k v_k - \chi_k p_k^{HP}) &> \alpha_{k+m}(\chi_k v_k - \chi_{k+1} v_{k+1}) \\
&\quad +\alpha_k(\chi_{k+1} v_{k+1} - \chi_k p_k^{HP}) && \text{(F-15)}\\
&= \alpha_{k+m}(\chi_k v_k - \chi_{k+1} v_{k+1}) \\
&\quad +\alpha_{k+1}(\chi_{k+1} v_{k+1} - \chi_{k+1} p_{k+1}^{HP}) && \text{(F-16)}\\
&> \alpha_{k+m}(\chi_k v_k - \chi_{k+2} v_{k+2}) \\
&\quad +\alpha_{k+1}(\chi_{k+2} v_{k+2} - \chi_{k+1} p_{k+1}^{HP}) && \text{(F-17)}\\
&> \ldots \\
&= \alpha_{k+m}(\chi_k v_k - \chi_{k+m-1} v_{k+m-1}) \\
&\quad +\alpha_{k+m-1}(\chi_{k+m-1} v_{k+m-1} - \chi_{k+m-1} p_{k+m-1}^{HP}) \\
&> \alpha_{k+m}(\chi_k v_k - \chi_{k+m} v_{k+m}) \\
&\quad +\alpha_{k+m-1}(\chi_{k+m} v_{k+m} - \chi_{k+m-1} p_{k+m-1}^{HP}) \\
&= \alpha_{k+m}(\chi_k v_k - \chi_{k+m} v_{k+m}) \\
&\quad +\alpha_{k+m}(\chi_{k+m} v_{k+m} - \chi_{k+m} p_{k+m}^{HP}) \\
&= \alpha_{k+m}(\chi_k v_k - \chi_{k+m} p_{k+m}^{HP})
\end{aligned}
$$

In the above chain, inequality (F-15) obtains because of the result from Step 3 that $v_j\chi_j > v_{j+1}\chi_{j+1}\ \forall\ j$ such that $1 \le j \le K+1$, and because $\alpha_k > \alpha_{k+m}$. The equality (F-16) follows from the locally dominant envy-free equilibrium condition. The inequality (F-17) obtains because $v_{k+1}\chi_{k+1} > v_{k+2}\chi_{j+2}$, $\alpha_{k+1} > \alpha_{k+m}$, and because of the following:

$$
\begin{aligned}
&\alpha_{k+m}(\chi_k v_k - \chi_{k+1} v_{k+1}) + \alpha_{k+1}(\chi_{k+1} v_{k+1} - \chi_{k+1} p_{k+1}^{HP}) \\
&> \alpha_{k+m}[(\chi_k v_k - \chi_{k+1} v_{k+1}) \\
&\quad + (\chi_{k+1} v_{k+1} - \chi_{k+2} v_{k+2})] + \alpha_{k+1}(\chi_{k+2} v_{k+2} - \chi_{k+1} p_{k+1}^{HP}) \\
&= \alpha_{k+m}(\chi_k v_k - \chi_{k+2} v_{k+2}) + \alpha_{k+1}(\chi_{k+2} v_{k+2} - \chi_{k+1} p_{k+1}^{HP})
\end{aligned}
$$

The rest of the chain is obtained in an analogous fashion. Now, consider the payoff π_k^u to advertiser in slot k ($k > 1$) from moving up to slot $k-m$, where $1 \le m \le k-1$, by bidding sufficiently higher such that his expected profit to the publisher is slightly above that of the advertiser assigned to that slot. Note that by doing so, the advertiser's payment depends on the payment assigned to slot

$k-m-1$.[27] π_k^u is given by the following expression:

$$\begin{aligned}\pi_k^u &= \chi_k[\alpha_{k-m}(v_k - p_{k-m-1}^{HP}\chi_{k-m-1}/\chi_k) - \alpha_k(v_k - p_k^{HP})] \\ &< \chi_k[\alpha_{k-m}(v_k - p_{k-m}^{HP}\chi_{k-m}/\chi_k) - \alpha_k(v_k - p_k^{HP})] \\ &= \alpha_{k-m}(v_k\chi_k - p_{k-m}^{HP}\chi_{k-m}) - \alpha_k(v_k\chi_k - p_k^{HP}\chi_k)\end{aligned}$$

In obtaining the inequality, in the above expression, we replace $p_{k-m-1}^{HP}\chi_{k-m-1}$ by $p_{k-m}^{HP}\chi_{k-m}$. $p_{k-m-1}^{HP} > p_{k-m}^{HP}$ and $\chi_{k-m-1} \geq \chi_{k-m}$, thus justifying the above inequality. The following chain of equalities and inequalities shows that $\alpha_k(v_k\chi_k - p_k^{HP}\chi_k) > \alpha_{k-m}(v_k\chi_k - p_{k-m}^{HP}\chi_{k-m})$ implying that $\pi_k^u < 0$. Thus, the advertiser in slot k does not gain by moving to slot $k-m$.

$$\begin{aligned}\alpha_k(v_k\chi_k - p_k^{HP}\chi_k) &= \alpha_{k-1}(v_k\chi_k - p_{k-1}^{HP}\chi_{k-1}) & \text{(F-18)} \\ &> \alpha_{k-m}(v_k\chi_k - v_{k-1}\chi_{k-1}) & \text{(F-19)} \\ &\quad +\alpha_{k-1}(v_{k-1}\chi_{k-1} - p_{k-1}^{HP}\chi_{k-1}) \\ &= \alpha_{k-m}(v_k\chi_k - v_{k-1}\chi_{k-1}) & \text{(F-20)} \\ &\quad +\alpha_{k-2}(v_{k-1}\chi_{k-1} - p_{k-2}^{HP}\chi_{k-2}) \\ &> \alpha_{k-m}(v_k\chi_k - v_{k-2}\chi_{k-2}) & \text{(F-21)} \\ &\quad +\alpha_{k-2}(v_{k-2}\chi_{k-2} - p_{k-2}^{HP}\chi_{k-2}) \\ &= \ldots \\ &= \alpha_{k-m}(v_k\chi_k - v_{k-m+2}\chi_{k-m+2}) \\ &\quad +\alpha_{k-m+1}(v_{k-m+2}\chi_{k-m+2} - p_{k-m+1}^{HP}\chi_{k-m+1}) \\ &> \alpha_{k-m}(v_k\chi_k - v_{k-m+1}\chi_{k-m+1}) \\ &\quad +\alpha_{k-m+1}(v_{k-m+1}\chi_{k-m+1} - p_{k-m+1}^{HP}\chi_{k-m+1}) \\ &= \alpha_{k-m}(v_k\chi_k - v_{k-m+1}\chi_{k-m+1}) \\ &\quad +\alpha_{k-m}(v_{k-m+1}\chi_{k-m+1} - p_{k-m}^{HP}\chi_{k-m}) \\ &> \alpha_{k-m}(v_k\chi_k - v_{k-m}\chi_{k-m}) \\ &\quad +\alpha_{k-m}(v_{k-m}\chi_{k-m} - p_{k-m}^{HP}\chi_{k-m}) \\ &= \alpha_{k-m}(v_k\chi_k - p_{k-m}^{HP}\chi_{k-m})\end{aligned}$$

In the above chain, equality (F-18) obtains because of the locally dominant envy-free equilibrium condition. Inequality (F-19) follows from the result in Step 3 that $v_j\chi_j > v_{j+1}\chi_{j+1}\ \forall\ j$ such that $1 \leq j \leq K+1$, and because $\alpha_{k-1} > \alpha_{k-m}$. The equality (F-20) again follows from the locally dominant envy-free equilibrium

[27] See footnote 26 on this point. Without loss of generality, we define $p_0^{HP} = p_1^{HP} + \varepsilon$, for some arbitrary $\varepsilon > 0$.

condition. The inequality (F-21) obtains because $v_{k-2}\chi_{k-2} > v_{k-1}\chi_{k-1}$, $\alpha_{k-2} > \alpha_{k-m}$, and because of the following:

$$\begin{aligned}
&\alpha_{k-m}(v_k\chi_k - v_{k-1}\chi_{k-1}) + \alpha_{k-2}(v_{k-1}\chi_{k-1} - p^{HP}_{k-2}\chi_{k-2}) \\
&> \alpha_{k-m}[(v_k\chi_k - v_{k-1}\chi_{k-1}) + (v_{k-1}\chi_{k-1} - v_{k-2}\chi_{k-2})] \\
&\quad + \alpha_{k-2}(v_{k-2}\chi_{k-2} - p^{HP}_{k-2}\chi_{k-2}) \\
&= \alpha_{k-m}(v_k\chi_k - v_{k-2}\chi_{k-2}) + \alpha_{k-2}(v_{k-2}\chi_{k-2} - p^{HP}_{k-2}\chi_{k-2})
\end{aligned}$$

The rest of the chain is obtained in an analogous fashion. Thus, we have established the equilibrium for the highest profit policy as presented in the proposition. As in the case of the highest bid policy, the equilibrium can shown to be a stable assignment as well.

Part (ii). Since $\chi_{j+1} \leq \chi_j$, inspection reveals that $p^{HB}_k \geq p^{HP}_k$. Moreover, since $\pi^{HB}_{P,k} = \alpha_k\chi_k p^{HB}_k$ and $\pi^{HP}_{P,k} = \alpha_k\chi_k p^{HP}_k$, $\pi^{HB}_{P,k} \geq \pi^{HP}_{P,k}$. Only for $\lambda = 0$, $\chi_{j+1} = \chi_j$; and so, $p^{HB}_k = p^{HP}_k$ and $\pi^{HB}_{P,k} = \pi^{HP}_{P,k}$ in this case. ■

G Proof of Proposition 4

Proposition 4 : *Under the highest bid policy, for slot k such that $1 \leq k < K$, $p^{HB}_k > p^{HB}_{k+1}$.*

Proof. From Proposition 3 (i), we have for k such that $1 \leq k < K$:

$$\begin{aligned}
p^{HB}_k - p^{HB}_{k+1} &= \frac{1}{\alpha_k}\sum_{j=k}^{K} v_{j+1}(\alpha_j - \alpha_{j+1}) - \frac{1}{\alpha_{k+1}}\sum_{j=k+1}^{K} v_{j+1}(\alpha_j - \alpha_{j+1}) \\
&= \left(v_{k+1} + \frac{\alpha_{k+1}}{\alpha_k}(v_{k+2} - v_{k+1}) + \sum_{j=k+2}^{K}\frac{\alpha_j}{\alpha_k}(v_{j+1} - v_j)\right) \\
&\quad - \frac{\alpha_{K+1}}{\alpha_k}v_{K+1} - \left(v_{k+2} + \sum_{j=k+2}^{K}\frac{\alpha_j}{\alpha_{k+1}}(v_{j+1} - v_j) - \frac{\alpha_{K+1}}{\alpha_{k+1}}v_{K+1}\right) \\
&= \frac{1}{\alpha_k}(\alpha_k - \alpha_{k+1})(v_{k+1} - v_{k+2}) + (\frac{1}{\alpha_k} - \frac{1}{\alpha_{k+1}})\sum_{j=k+2}^{K}\alpha_j(v_{j+1} - v_j) \\
&> 0.
\end{aligned}$$

We obtain the second equality in the above sequence by rearranging terms. The third equality is obtained by regrouping terms and noting that $\alpha_{K+1} = 0$. The last inequality follows because $\alpha_k > \alpha_{k+1}$, and $v_k > v_{k+1}$. ■

References

Baye, M. R. and J. Morgan (2005). Brand and Price Advertising in Online Markets. Working Paper, Indiana University.

Brownstone, D. and K. Train (1999). Forecasting New Product Penetration with Flexible Substitution Patterns. *Journal of Econometrics 89*, 109–129.

Bult, J. R. and T. Wansbeek (1995, Fall). Optimal Selection for Direct Mail. *Marketing Science 14*(4), 378–394.

Chen, Y. and C. He (2007). Paid Placement: Advertising and Search on the Internet. Working Paper, University of Colorado.

Chen, Y., G. Iyer, and V. Padmanabhan (2002). Referral Infomediaries. *Marketing Science 21*(4), 23–45.

ClickThrough (2008). PPC internet marketing growth due to reliability, says eMarketer. `http://www.clickthrough-marketing.com/PPC%20internet%20marketing%20growth%20due%20to%20reliability,%20says%20eMarketer/18907861/`.

Delaney, K. J. and R. A. Guth (2005, August 3). Yahoo, Microsoft Enhance Ad Focus. Wall Street Journal.

Dukes, A. and E. Gal-Or (2003). Negotiations and Exclusivity Contracts for Advertising. *Marketing Science 22*(2), 222–245.

Edelman, B., M. Ostrovsky, and M. Schwarz (2007). Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review 97*(1), 242–259.

Engelbrecht-Wiggans, R., E. Haruvy, and E. Katok (2007). A Comparison of Buyer-Determined and Price-Based Multiattribute Mechanisms. *Marketing Science 26*(5), 629–641.

Feng, J., H. Bhargava, and D. Pennock (2005). Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal of Computing*.

Fudenberg, D. and J. Tirole (1994). *Game Theory*. Cambridge, MA: MIT Press.

Hu, J. (2004). Performance Based Pricing Models in Online Advertising. Working Paper, MIT.

iCrossing (2007). iCrossing Search Synergy: Natual and Paid Search Symbiosis.

Iyer, G. and A. Pazgal (2003). Internet Shopping Agents: Virtual Colocation and Competition. *Marketing Science 22*(1), 85–106.

Iyer, G., D. Soberman, and J. M. Villas-Boas (2005). The Targeting of Advertising. *Marketing Science 24*, 461–476.

Levi, M. D. (1973). Errors in the Variables Bias in the Presence of Correctly Measured Variables. *Econometrica 41*(5), 985–986.

McAfee, P. and J. McMillan (1987). Auctions and Bidding. *Journal of Economic Literature 15*, 699–738.

Oser, K. (2006). Why Marketers Need to Take Search Seriously. `http://adage.com/searchmarketing/article?article_id=108777`.

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics 6*(2), 461–464.

Snir, H., D. Monderer, and A. Sela (1998). A Learning Approach to Auctions. *Journal of Economic Theory 82*, 65–88.

Steel, E. (2007, May 25). Advertising's Brave New World — Different Lineup of Players Emerges With Online's Rise. The Wall Street Journal.

Varian, H. R. (2007). Position Auctions. *International Journal of Industrial Organization*.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance 16*, 8–27.

Weber, T. and E. Zheng (2005). A Model of Search Intermediaries and Paid Referrals. Working Paper, Stanford University.

Yao and Mela (2009). A Dynamic Model of Sponsored Search Advertising. Working Paper, Duke University.

ZenithOptimedia (2008). No growth in global adspend in 2009. `http://www.optimedia.de/fileadmin/pdf/presse/2008-12-08_PM_AEF_englisch.pdf`.