

The Race for Sponsored Links: Bidding Patterns for Search Advertising

Zsolt Katona

Haas School of Business, University of California at Berkeley, Berkeley, California 94720, zskatona@haas.berkeley.edu

Miklos Sarvary

INSEAD, Fontainebleau, 77305 France, miklos.sarvary@insead.edu

Paid placements on search engines reached sales of nearly \$11 billion in the United States last year and represent the most rapidly growing form of online advertising today. In its classic form, a search engine sets up an auction for each search word in which competing websites bid for their sponsored links to be displayed next to the search results. We model this advertising market, focusing on two of its key characteristics: (1) the interaction between the list of search results and the list of sponsored links on the search page and (2) the inherent differences in attractiveness between sites. We find that both of these special aspects of search advertising have a significant effect on sites' bidding behavior and the equilibrium prices of sponsored links. Often, sites that are not among the most popular ones obtain the sponsored links, especially if the marginal return of sites on clicks is quickly decreasing and if consumers do not trust sponsored links. In three extensions, we also explore (1) heterogeneous valuations across bidding sites, (2) the endogenous choice of the number of sponsored links that the search engine sells, and (3) a dynamic model where websites' bidding behavior is a function of their previous positions on the sponsored list. Our results shed light on the seemingly random order of sites on search engines' list of sponsored links and their variation over time. They also provide normative insights for both buyers and sellers of search advertising.

Key words: Internet marketing; position auctions; game theory

History: Received: September 18, 2007; accepted: May 13, 2009; processed by Yuxin Chen. Published online in *Articles in Advance* August 14, 2009.

1. Introduction

Search advertising is becoming one of the dominant forms of online advertising. Potential advertisers bid for a place on the list of sponsored links that appears on a search engine's "results" page for a specific search word. In 2008, the revenues from such paid placements have increased by over 20% compared to 2007, reaching nearly \$11 billion in the United States despite the economic crisis.¹ This fast-growing market is increasingly dominated by Google, which today controls some 70% of Internet searches.² How such advertising is priced and what purchase behavior advertisers will follow for this new form of advertising is the subject of the present paper.

Previous research studying search advertising has focused on the problem of multi-item (or position)

auctions and examined the optimal bidding behavior of advertisers (Edelman et al. 2007, Varian 2007). However, a key characteristic of paid placements is that the consumer is facing two "competing" lists of sites that are both relevant in the context of the particular search: (1) the *organic links* and (2) the list of *sponsored links*. Furthermore, membership and position on the organic list is generally exogenous to the site and typically represents the site's general popularity or inherent value. The search engine cannot use this list strategically without losing credibility from users. Thus, the existence of this organic list cannot be ignored when one evaluates sites' bidding behavior for sponsored links appearing on the same page.

Another key characteristic of the problem is that the search engine can take into account advertisers' expected traffic when awarding paid links. Because the bids correspond to payments per click, this information is important in determining the search engine's total revenue from a given sponsored link. Therefore, search engines take sites' expected click-through rates into account in addition to their per-click bids when awarding paid placements. Furthermore, the search engine can also decide how many spon-

¹ See *Emarketer* (2009). Worldwide revenues from paid placements are expected to reach \$45 billion by 2011 (see *BusinessWeek* 2007).

² See *Emarketer* (2009). Furthermore, other major search engines use similar methods to target searching consumers. AOL uses Google's search, whereas Yahoo!'s search page is almost identical to Google's. Other popular sites like Amazon and eBay (also powered by Yahoo!) sell only a few sponsored links on their search pages, and many times these are linked to their own content.

sored links it offers for a particular search word. Again, advertisers' incentives for bidding and, in turn, the search engine's revenue will depend on this decision.

Finally, a third important characteristic of paid placements is that bidding for sponsored links happens frequently over time. This has two implications. First, it means that repeated bidding by the same players reveals their valuations for the different advertising links. Second, if the advertising effect of sponsored links has a lagged component—as is often the case with advertising—then bidding strategies should be dynamic rather than optimized for a single time period.

We develop a model that takes into account these key aspects of search advertising. Specifically, in our base model we explicitly describe consumers' clicking behavior on the search page as a function of sites' presence and order on the organic links list and/or among the sponsored links. Then we derive sites' optimal bidding strategies for sponsored links and the search engine's optimal behavior, taking into account consumers' clicking patterns.

Our results shed light on the advertising patterns observed on different search pages. Specifically, search pages can be characterized by a variety of patterns in terms of the identity and position of sponsored links. In particular, there does not seem to be a clear relationship between the organic list of a search and the list of sponsored links. Sometimes a site may appear in both or in only one (either one) of the lists. For example, at the time of writing this paper, on Google's search results page for the word "travel," the two lists were entirely different. However, on the results page for the search word "airlines," United Airlines appeared as the first organic result and second on the sponsored links list. One can also observe significant fluctuations in the sites' order in the sponsored-links list. Finally, the number of items listed in the sponsored list is also changing over time. Our model proposes a number of testable hypotheses that account for these variations. In particular, sites that are among the top organic links win the sponsored links if they are much more popular than their rivals and if marginal return on clicks is close to constant. Sites that are not on the organic list tend to win the sponsored links when the marginal return on clicks is quickly decreasing and when consumers are averse to sponsored links. Our results also generate normative guidelines to both advertisers and the search engine on how to buy and sell sponsored links. For instance, our analysis suggests that, in some cases, a search engine can attain higher revenues by displaying fewer sponsored links.

In a second step, we provide three extensions to the base model. First, we explore the case when

clicks are valued heterogeneously across sites. We find that the basic competitive dynamics do not change, although the actual outcomes are influenced by sites' specific valuations. In a second extension, we allow the search engine to choose the number of sponsored links to auction away. We show under what conditions it is worth it for the search engine to increase or decrease the number of links. Finally, we also explore a dynamic model where sites bid repeatedly and consumer clicks have a lagged effect (e.g., because of a loyalty factor). Here, we find conditions under which sites either alternate in winning the auction or their order remains relatively stable. In particular, we show that an alternating equilibrium is better for all the players.

The rest of the paper is organized as follows. The next section summarizes the relevant literature. This is followed by the basic model description in §3 and equilibrium analysis in §4. Section 5 explores the three extensions outlined above. We end with a summary of the main findings and model limitations in §6. All proofs and technical details appear in the appendix.

2. Relevant Literature

Because search advertising is mostly responsible for the growth of the online advertising business, it has attracted significant interest in the economics literature.³ Edelman et al. (2007) analyze the generalized second-price auction that is used by most search engines to allocate sponsored links on search pages.⁴ The paper focuses on equilibrium properties and compares these to other auction mechanisms. Varian (2007) studies a similar problem but assumes away uncertainty and shows that the equilibrium behavior matches empirical pricing patterns for sponsored links. More recent papers (Feng 2008, Feng et al. 2007, Athey and Ellison 2008) further elaborate on optimal auction design by considering reserve prices.

A separate set of papers explore the important issue of fraudulent behavior in the context of search advertising. Wilbur and Zhu (2009) study click fraud and its nontrivial effect on the distribution of surplus between advertisers and the search engine. In a related study, Bhargava et al. (2005) explore shill bidding in a consumer auction context where bidders can establish multiple identities.

Although the previous streams add considerably to our understanding of how to efficiently allocate search

³ The other dominant advertising model—sites buying ads on each other's pages—is analyzed in Katona and Sarvary (2008). That paper studies equilibrium advertising prices and the endogenous network structure determined by the advertising links.

⁴ This literature builds on an established stream of research on mechanism design represented by classic papers such as those of Myerson (1981) and Maskin and Riley (2000).

advertising, they neglect the behavior of searching consumers. Chen and He (2006) also study competitive bidding for paid placements but assume differentiated advertisers and explicitly consider consumers who are initially uncertain about their valuations for products. They show how the auction mechanism improves the efficiency of consumer search and results in possible price dispersions for advertising. Athey and Ellison (2008) extend this approach and further explore the implications of the results for optimal auction design.

Our work is different from these literature streams. We assume away fraud and are less focused on optimal auction design, but are interested in capturing relevant behaviors from searching consumers. In particular, our focus is on the *interaction* between the search engine's basic service of finding relevant sites in a given search context and its private objective to sell sponsored links on search pages. We model the inherent competition between the output of these two processes and evaluate its effect on advertisers' behavior. In terms of modeling the allocation of sponsored links, our paper is closest to Varian (2007), but our focus is elsewhere. Rather than characterizing the optimal auction "rules" for allocating multiple items, we are interested in how different sites compete for the sponsored links in the auction and what the role of the organic links is.

Beyond the explicit modeling of consumers' clicking behavior, our modeling approach is different in many other ways. We assume a weakly concave response function to advertising that is well documented in marketing. As opposed to the existing literature, we also explore the endogenous choice of the number of sponsored links offered, which can be an important decision variable for the search engine. Furthermore, we study a dynamic model in which advertisers repeatedly bid for sponsored links and consumer visits have a lagged effect. This dynamic advertising model is related to previous work on the dynamic setting of marketing variables in a competitive context using a Markovian game. For an application to advertising, see Villas-Boas (1993); an application for dynamic research and development competition can be found in Ofek and Sarvary (2003). Our work uses a similar framework and relates to the results of both papers. The possibility of an alternating advertising pattern is similar to Villas-Boas (1993) and is largely driven by decreasing returns on advertising. However, in our model, as in Ofek and Sarvary (2003), we have a contest as advertisers bid for each position on the list with only one winner. Our dynamic model is also somewhat related to the dynamic auction model of Zeithammer (2006). However, in our case this is a repeated auction for a period prize, whereas his paper considers dynamic bidding for a single item.

Finally, recent empirical work on search advertising (Rutz and Bucklin 2007a, b) studies the effectiveness of paid placements with particular attention devoted to spillover and lagged effects, as well as contexts when multiple search words are used. In another paper, Goldfarb and Tucker (2007) show that the auction mechanism allows search engines to discriminate between bidding firms with different inherent valuations for advertising. Similarly, in a recent study, Yao and Mela (2008) assess how different auction mechanisms affect advertiser and search engine profitability. They also explore the effect of information asymmetry between the search engine and advertisers and the possibility for advertisers to bid by consumers' search histories and demographics for more-targeted advertising. Our model extensions are largely motivated by these papers (see our dynamic model and our examination of heterogeneous firm valuations), although the present paper admittedly has a more normative focus.

3. The Model

We assume n websites that are indexed with respect to their exogenously given, inherent attractiveness levels, $1 > \gamma_1 > \gamma_2 > \dots > \gamma_n$. These rates represent the inherent attractiveness (value or popularity) of the sites in the eyes of the *average* consumer in the context of a given search word. Notice that the attractiveness of a site varies across consumers, but we assume that γ , i.e., the average across consumers, is fixed and exogenous. One could think of these γ values as objective measures of site quality (in the given search context).

The $(n + 1)$ th player is a search engine (SE), a special website.⁵ Although consumers do not know the γ values (because they do not observe other consumers' valuation of the sites), the SE's job is precisely to evaluate sites' inherent attractiveness. As such, we assume that the SE ranks the sites according to their inherent attractiveness in a given search context; that is, the γ values determine sites' ranking on the organic list. This is consistent with the idea that the SE's basic service lies in finding sites in which consumers (on average) are most interested.

We also assume that inherent attractiveness levels are also known by the sites because the search engine lists the sites in the order of their attractiveness and also because proxy statistics (e.g. click-through rates) are available from independent research firms. One could argue, however, that although it is reasonable to

⁵ We assume that the SE is a monopolist. Although this is not entirely true in practice, Google dominates the search industry with over 70% of all searches, a proportion that is growing (Emarketer 2009).

assume that sites know their own γ , they do not necessarily know their competitors' γ s. In an extension (see Appendix C), we explore such a scenario and show that our results are robust to information asymmetry across sites. Nevertheless, the assumption of quite well-known valuations across sites is consistent with reality. On one hand, sites can infer attractiveness levels from the search engine's organic listings, which rank the sites according to their attractiveness levels and sometime even provides popularity scores. On the other hand, sites have the possibility of conducting market research or even experimentation on their own sites by measuring the clicks on an experimental link to one of their competitor's sites.

In our model, the SE returns the r highest-ranked sites as the organic links (Sites $1, 2, \dots, r$). Next to these organic links, the SE also displays s number of sponsored links. The order of these links can be chosen by the SE, and this choice is based on the bids submitted by the websites. Let l_1, l_2, \dots, l_s denote the sites winning the sponsored links, in order of appearance. Thus, the output of the SE is modeled as a page with two lists: an organic list and a sponsored ads list. Google's search page is exactly like this, and other search engines have a similar format.

3.1. Consumers' Behavior on the Search Page

We assume that the SE attracts a unit traffic of consumers that is distributed in the following way. When a consumer arrives to the SE's page generated by the search, he or she either clicks on one of the regular results, one of the sponsored links, or leaves the page without clicking. We assume that consumers' clicking behavior is affected by the following four factors:

1. The order in which sites are listed,
2. Differences in click probabilities between the sponsored list and the search result list,
3. Individual differences between sites in their inherent attractiveness, and
4. Whether the site appears in both the organic and the sponsored lists or only one of the lists.

For the first factor, assume that $\alpha_1, \alpha_2, \dots > 0$ denote the psychological order constants that determine how the possible clicks are distributed through an ordered list of items. That is, whenever someone sees an ordered list of *equally interesting* items, he or she chooses the i th item with probability proportional to α_i .⁶ Generally, we can say that $\alpha_1 > \alpha_2 > \dots$, but there might be exceptions. For example, the last item in a list may be more appealing than one in the middle. For the second factor, let $\beta > 0$ denote how many times more/less attractive a sponsored link is than

an organic link—that is, how many times more/fewer consumers click on a sponsored link over an equally interesting link in the same position on the organic search list. Because consumers are likely to exhibit some level of aversion to advertising (see, e.g., Lutz 1985 for a classic reference; Edwards et al. 2002 and Schlosser et al. 1999 for empirical evidence in the Internet context), we expect β to be less than 1, although we do not need to assume this. Combining the two factors, the distribution of consumers among the links, not taking into account individual differences between sites' popularity, is determined by the parameters: $\alpha_1, \alpha_2, \dots, \alpha_r$ and $\beta\alpha_1, \beta\alpha_2, \dots, \beta\alpha_s$. Specifically, $M = \sum_{i=1}^r \alpha_i + \beta \sum_{i=1}^s \alpha_i$ represents the maximum potential traffic that can flow through all the links on the search page.⁷ Because the SE has a unit incoming traffic for each search word, we normalize M to 1. The traffic that flows through the links is less than M , however, because it also depends on the sites' attractiveness. This is taken into account in the third factor, which we explore next.

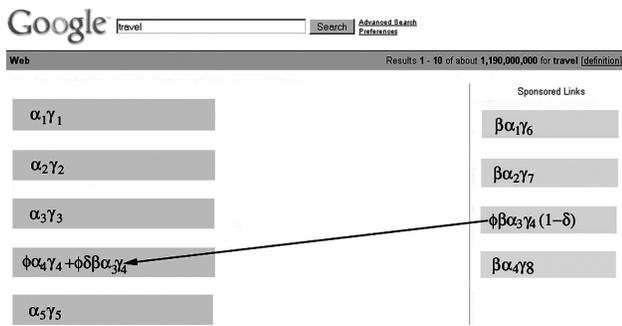
For the third factor, which takes sites' individual differences into account, we multiply the α and β parameters with the inherent attractiveness levels of the sites (γ_i). In any particular position, a site with a higher attractiveness level is more likely to receive a click than another site in the same position having a lower inherent attractiveness. For example, Site 1 will receive $\alpha_1\gamma_1$ clicks on the first organic link, whereas Site j in the second position on the sponsored list will receive $\beta\alpha_2\gamma_j$ clicks on its sponsored link. We use a multiplicative model because we think of γ_i as a proportion. It determines what proportion of the maximum clicks that are attainable in a certain position will be received by a site based on its inherent attractiveness. Naturally, we assume that γ_i is independent from the α values and β .

Finally, for the fourth factor, we consider the interaction effects between the two lists. First, if a link appears both in the organic and sponsored lists, it is possible that the total number of clicks that these two receive is smaller than just the sum of the possible clicks received for the two. This can be the result of decreasing returns to exposure on the search page as well as aversion to advertising in the sense that the presence of an advertising link decreases the attractiveness of the site. Specifically, if a site is listed both among the regular search results and the sponsored links, then the number of clicks on both links is ϕ times the number that it would have received had it not appeared on both lists, where $0 < \phi \leq 1$. Furthermore, we also assume that if a site has both an

⁶ The α_i s might also capture the extent of trust consumers have in the relevance of the order of links on top of their own judgement of the sites.

⁷ We assume M to be constant for the time horizon of the bidders, but clearly, M could change over the long run as more or fewer consumers use the search engine depending on the quality of the results.

Figure 1 The Distribution of Clicks on a Search Page



Notes. The left list contains the organic search results, where sites appear in the order of their inherent attractiveness, γ_i . The right list represents paid placements where the order depends on sites' bids. In this example, the third sponsored link also appears on the organic search list in position 4; therefore, $1 - \phi$ proportion of clicks on this site's links are lost and $\phi\delta\beta\alpha_3\gamma_4$ clicks are redirected to the organic link.

organic and a sponsored link, then some of the possible clicks on the sponsored link will transfer to the organic link. In particular, let δ denote the proportion of clicks that would be made on the sponsored link but are indeed made on the same site's organic link.⁸ Note that, as opposed to ϕ , the parameter δ does not have an effect on the total traffic that a site receives from the search engine because it simply changes the origin of this traffic. However, it affects the number of clicks on sponsored links, which will be important in the bidding process and will also affect the SE's revenue. Figure 1 summarizes the four factors of our model describing the clicking behavior of consumers.

Given these factors, we now determine how the traffic of the search engine is distributed through the websites. Let $A(i)$ denote the function that takes a value of 0 if Site i does not win a sponsored link; that is, $i \notin \{l_1, l_2, \dots, l_s\}$ and α_j if Site i wins the j th sponsored link. With this, the total traffic that Site i gets from the search engine is

$$t_i = t_i^R + t_i^S, \quad (1)$$

where the two types of links receive the following number of clicks depending on whether the site appears on both lists or only one. If the site appears on both lists, the traffic is

$$\begin{aligned} t_i &= t_i^R + t_i^S = \phi\gamma_i\alpha_i + \phi\gamma_i\beta A(i) \\ &= \phi(\gamma_i\alpha_i + \gamma_i\beta A(i)). \end{aligned} \quad (2)$$

If site i has an organic link only, then the traffic is

$$t_i = t_i^R = \gamma_i\alpha_i \quad \text{and} \quad t_i^S = 0. \quad (3)$$

⁸ Similarly to the case of β and ϕ , we speculate that because of aversion to advertising, consumers prefer organic results to sponsored links; that is, $\delta > 0$. However, our results also hold for negative values of δ .

Finally, if i appears only on the sponsored list, then it is

$$t_i = t_i^S = \gamma_i\beta A(i) \quad \text{and} \quad t_i^R = 0. \quad (4)$$

Note that these quantities largely depend on γ_i , the site's inherent attractiveness level, which determines how many clicks a site's link receives given its position.

3.2. Websites

Websites make profits from the traffic that arrives to their sites from the search engine.⁹ Let us assume that there is a common $R(t)$ function for all sites that determines the revenue associated with t amount of traffic for a given search word. As such, we assume that for each word there exists a common function determining how clicks can be converted into revenues. In §5.1, we relax this assumption and allow for individual differences in sites' valuations. Here, we naturally assume that $R(t)$ is increasing and weakly concave.¹⁰ To obtain sponsored links, sites have to submit bids to the search engine. The bid that Site i submits, b_i , is the maximum amount that it is willing to pay for unit traffic (per-click). If the search engine decides to include Site i among the sponsored links, Site i has to pay an advertising fee of $p_i t_i^S$, where $p_i \leq b_i$ is set by the search engine. Therefore, Site i 's utility is

$$u_i = R(t_i^R + t_i^S) - p_i t_i^S, \quad (5)$$

where t_i^R and t_i^S depend on which sites win the sponsored links as defined in (1).

Thus, in our model, the SE uses an auction to allocate the sponsored links. This is consistent with what search engines do in reality. However, an auction may not be necessary for such allocation because all players have common knowledge about all valuations; i.e., the game is one of complete information. As we mentioned before, we assume this because inherent attractiveness levels are common knowledge, and repeated bidding gives ample time and data for all players to discover the valuations of other parties. A similar argument is advanced in Edelman et al. (2007) and Varian (2007). As such, the auction mechanism is used as an efficient pricing mechanism. Although

⁹ Thus, we ignore the fact that sites could already have different amounts of incoming traffic from other sources. If we naturally assume that more-attractive sites also have higher outside traffic, then the results follow the same patterns.

¹⁰ We allow $R(t)$ to be linear, as we only require weak concavity (see Rutz and Bucklin 2007a for a detailed analysis on how $R(t)$ could be estimated in practice). Also, the concavity assumption captures some of the dependence of the different visits of the same person at the same site. If the same person visits the same site multiple times, the marginal value of subsequent clicks is probably decreasing. We would like to thank the area editor for drawing our attention to this point.

in theory the SE could calculate and offer the optimal price for each sponsored link, allowing sites to self-select for each position, pricing the links through an auction is easier and more robust to variations of participants over time.¹¹ There are also costs associated with setting prices individually, which might be overwhelming for the enormous number of possible keywords (see Zeithammer and Liu 2008 for a study of the trade-offs between using fixed prices or auctions).

At this point, the SE is completely free to determine the order of winners and the advertising fee it charges for a click, $p_i \leq b_i$. First, we will show that in a one-period game the SE sets $p_i = b_i$ corresponding to a first-price auction; then we will discuss the different types of auctions that search engines use in practice. Based on this discussion, in §3.3 we will restrict the SE's strategies and define the types of equilibria we use in the subsequent analysis.

The timing of the game is the following. First, websites simultaneously submit their b_i bids, knowing all the attractiveness levels and $R(t)$. Then, the search engine decides which sites it will include among the sponsored links and in what order. Finally, sites pay the advertising fee to the search engine and realize profits from the traffic they receive.

3.3. The Search Engine

First, we determine the SE's best response to given bids b_1, b_2, \dots, b_n in the second stage of the game. Although it would seem so, the best strategy is not to simply assign the sponsored links to websites in the order of their bids. The SE has to consider the sites' expected click-through rates (CTRs), because the total traffic it sells to them, and thus its revenue, depends on these rates. Therefore, a site with a high expected CTR may pay a higher total fee even if its bid is low. An opposite effect is that the most attractive sites will also appear on the regular search list. As a result, they will attain fewer clicks on the sponsored link because a $1 - \phi$ proportion of clicks is lost and a δ proportion of the consumers will click on the organic link instead.¹² Formally, the SE maximizes its profit,

$$\Pi_{SE} = \sum_{i=1}^s t_i^S p_i. \quad (6)$$

¹¹ The auction may also resolve problems related to some level of information asymmetry about valuations. Our perspective follows Varian (2007) in that asymmetric information is not the key issue in the pricing of sponsored links. We would like to thank the area editor for drawing our attention to this issue. Also, as mentioned before, in Appendix C, we explore the case where sites have imperfect information about their competitors' inherent attractiveness.

¹² In the exceptional case of $\delta < 0$, it is the sponsored link that receives more clicks.

The SE's decision can be described by the series l_1, l_2, \dots, l_n , where Site l_i will get sponsored link i . Sites $l_{s+1}, l_{s+2}, \dots, l_n$ will not get a sponsored link. Let $I(i)$ denote the function that takes the value 1 if $i \leq r$ and 0 otherwise and let $J(i) = 1 - \phi I(i)$. Using $I()$ and $J()$ helps us simply capture the different cases in which sites can appear on one or both of the list.

CLAIM 1. *In equilibrium,*

$$\gamma_i J(l_i)(1 - \delta I(l_i))b_i \geq \gamma_j J(l_j)(1 - \delta I(l_j))b_j \quad (7)$$

holds for $i < j$, where $i \leq s$ and the SE sets $p_i = b_i$.

In other words, the search engine ranks the sites according to their $\gamma_i J(i)(1 - \delta I(i))b_i$ and charges each site's bid. That is, for sites that are not in the top r among the organic results, their position among the sponsored links is determined by their inherent attractiveness level multiplied by their bid. For top sites, this value is multiplied by $\phi(1 - \delta)$, accounting for consumers who lose interest in the site because of the sponsored link and those who choose to click on the organic link instead of the sponsored link.

As a result of Claim 1, in a nonrepeated game, the search engine's best strategy is to charge the highest CTR-corrected bid. The reason is that in this simple case in which sites only bid once, the search engine does not have to consider influencing sites' subsequent bidding strategies. This corresponds to a first-price auction. However, in reality most search engines use second-price auctions (most of them correcting for differences in CTRs) to avoid the problem that when multiple items with different values are auctioned, then the first-price auction typically does not have an equilibrium. This is because bids in a first-price auction always converge towards each other, which makes it impossible to reflect the differences in valuations for the different items.¹³ Thus, for our analysis, it is important to discuss the different types of auctions and equilibria that can be used in our models. We do this in detail in Appendix A. Based on that analysis, in what follows, we will use two equilibrium concepts depending on the auction mechanisms considered: (1) the first-price Nash equilibrium (FNE) concept, and (2) the symmetric second-price Nash equilibrium (SSNE) concept. In brief, in a symmetric Nash equilibrium in a second-price auction, the player in position k is better off paying the bid of the player in position $k + 1$ than he or she would be in position l paying the bid of player $l + 1$ (see Varian

¹³ The existence of an equilibrium may not be important to the SE, although it guarantees a certain level of price stability because sellers tend to converge to it over time. An additional reason to use a second-price auction is that, if valuations are uncertain, then the second-price auction is a mechanism that leads to truth telling in a single-item auction.

2007 and Appendix A). In case of sites bidding for a single sponsored link, we use the FNE and SSNE interchangeably because they provide the same result. In the case of multiple sponsored links, we use the SSNE because the FNE typically does not exist.

We always correct for expected CTRs as it is established in Claim 1. Player i 's bid is multiplied by $\gamma_i J(i)(1 - \delta I(i))$, and the search engine ranks the

$$F_i = \gamma_i J(i)(1 - \delta I(i))b_i \quad (8)$$

values when determining the order of sites and the prices. In a first-price auction, Site i has to pay $p_i = b_i$ for a click, corresponding to a total fee of $\beta A(i)F_i$, where $A(i)$ reflects its position. In a second-price auction, if Site i is followed by Site j in the order, then Site i has to pay

$$p_i = \frac{F_j}{\gamma_j J(j)(1 - \delta I(j))} = b_j \frac{\gamma_j J(j)(1 - \delta I(j))}{\gamma_j J(i)(1 - \delta I(i))} \quad (9)$$

for a click, totaling to a fee of $\beta A(i)F_j$. The next section determines the equilibrium bids.

4. Equilibrium Analysis

4.1. Bidding Strategies for One Sponsored Link

To illustrate the primary forces that work in the game, we first consider the case in which there is only one sponsored link offered; that is, $s = 1$. Let

$$G(i) = R(J(i)(I(i)\gamma_i\alpha_i + \gamma_i\beta\alpha_1)) - R(I(i)\gamma_i\alpha_i) \quad (10)$$

denote the revenue gain for Site i of winning the sponsored link. That is,

$$G(i) = R(\gamma_i\beta\alpha_1) \quad (11)$$

when Site i does not have an organic link ($i > r$) and

$$G(i) = R(\phi\gamma_i(\alpha_i + \beta\alpha_1)) - R(\gamma_i\alpha_i) \quad (12)$$

when it does. Clearly, the total fee Site i will pay for the sponsored link cannot exceed $G(i)$. Let w_1, w_2, \dots, w_n be a permutation of sites such that $G(w_1) > G(w_2) \geq \dots \geq G(w_n)$ holds.¹⁴ Furthermore, let P_1 denote the total fee that the winner pays for the sponsored link,¹⁵ which is equal to the seller's revenue.

PROPOSITION 1. *In any FNE and SSNE, the winner of the sponsored link is Site w_1 , and the total fee it pays is $G(w_1) \geq P_1 \geq G(w_2)$.*

¹⁴ The assumption that there is a single highest value eases the presentation of the results but does not change them qualitatively.

¹⁵ In the case of a first-price auction, this is calculated from its own bid. In the case of a second-price auction, it is calculated from the second-highest bid, corrected for CTRs.

Given the assumption that $R(\cdot)$ is increasing and weakly concave, the winner can be any site from 1 to $r + 1$, depending on the parameters. For example, if $R(\cdot)$ were linear and $\phi = 1$, then the site with the highest $\gamma_i\beta\alpha_1$, that is, Site 1, would be the winner. However, if $R(\cdot)$ is very concave, ϕ is small, or the γ_i s are not too far from each other—that is, $\gamma_1 - \gamma_{r+1} \rightarrow 0$ —then the winner is Site $r + 1$. These two cases illustrate the two forces that work against each other in determining the outcome. On one hand, because $R(\cdot)$ is concave, sites who already receive traffic from the search engine through organic links have a lower benefit from winning the sponsored link.¹⁶ On the other hand, sites with a higher γ_i obtain more traffic from a sponsored link; therefore, they are willing to pay more for such a link, unless they are worried about hurting their image (ϕ is small) and receiving fewer clicks by displaying a sponsored link. If the first effect is stronger, then a regularly lower-ranked site wins; otherwise, a top site wins the sponsored link. In reality, these two cases translate to the distinct, observed scenarios we mentioned above. For the word “travel,” the sponsored links and search result are distinct. In contrast, for the word “airlines,” a site appearing among the top search results also obtains a (top) sponsored link. Search engines often claim that the top sponsored links are relevant to customers because the sites that are willing to bid high for them are presumably relevant and attractive to consumers. However, as the results show, this is not always the case. Indeed, in a typical scenario, sites that are not present among the organic links win the sponsored links. Interestingly, this effect is even stronger if ϕ is small; that is, consumers are averse to sponsored links, leading to fulfilled expectations about the low quality of sponsored links.

The following corollary describes the equilibrium bids.

COROLLARY 1. *The winning bid in an FNE is*

$$\frac{G(w_1)}{\beta\alpha_1\gamma_{w_1}J(w_1)(1 - \delta I(w_1))} \geq b_1 > \frac{G(w_2)}{\beta\alpha_1\gamma_{w_1}J(w_1)(1 - \delta I(w_1))}. \quad (13)$$

In an SSNE, the winning bid can be arbitrarily high, but the second-highest bid is

$$\frac{G(w_1)}{\beta\alpha_2\gamma_{w_2}J(w_2)(1 - \delta I(w_2))} \geq b_2 > \frac{G(w_2)}{\beta\alpha_2\gamma_{w_2}J(w_2)(1 - \delta I(w_2))}. \quad (14)$$

¹⁶ This force is even stronger if we assume that sites with a high attractiveness have a larger traffic independent from the SE.

Note that the bids largely depend on the parameters. Sites with similar valuations might submit significantly different bids based on their attractiveness levels or their position among the regular search results. In Appendix C, we analyze the case in which sites have imperfect information about each other and find that they reveal their attractiveness levels through their bids.

4.2. Bidding Strategies for Multiple Sponsored Links

We will now discuss the general case with multiple sponsored links ($s > 1$). As mentioned before, the first-price auction typically does not have an equilibrium in this case; thus, we analyze the SSNE only. Let

$$G_j(i) = R(J(i)(I(i)\gamma_i\alpha_i + \gamma_i\beta\alpha_j)) - R(I(i)\gamma_i\alpha_i) \quad (15)$$

denote the revenue gain for Site i of the winning sponsored link j ($j = 1, \dots, s$). Let w_1, w_2, \dots, w_n denote the sites in the order of their CTR-corrected bids (F_i s). Furthermore, let P_i denote the total fee that Site i pays for the advertising:

$$P_i = b_{w_{i+1}}\alpha_{w_i}\beta\gamma_{w_i}(1 - \delta I(w_i)). \quad (16)$$

The search engine ranks the sites according to their CTR-corrected bids; that is, if the order is w_1, w_2, \dots , then the following have to hold for $2 \leq i \leq s$:

$$\frac{P_{i-1}}{\alpha_{i-1}} > \frac{P_i}{\alpha_i}. \quad (17)$$

In any equilibrium, Site w_k does not have an incentive to bid less and get to a lower position. Therefore,

$$G_k(w_k) - P_k \geq G_l(w_k) - P_l. \quad (18)$$

Furthermore, according to the definition of a symmetric equilibrium, Site w_l does not want to get into position k even if it has to pay P_k (and not P_{k-1}). That is,

$$G_l(w_l) - P_l \geq G_k(w_l) - P_k. \quad (19)$$

Combining (18) and (19), we get the following inequalities, describing the equilibria of the auction:

$$G_k(w_k) - G_l(w_k) \geq P_k - P_l \geq G_k(w_l) - G_l(w_l). \quad (20)$$

The complexity of the problem does not allow us to characterize all the SSNEs in this general case. Multiple equilibria may exist where the order of winners is different. The following example illustrates the complexity of the problem even in a simple case.

EXAMPLE 1. Assume $s = 2$ and $n = 3$, with the following valuations:

$$\begin{array}{lll} G_1(1) = 10, & G_2(1) = 8, & G_1(2) = 9, \\ G_2(2) = 6, & G_1(3) = 8, & G_2(3) = 7. \end{array}$$

These gains can be derived from a suitable $R()$ function, γ -s and α -s. Note that with prices $P_1 = 9$ and $P_2 = 7$, the equilibrium order of sites can be either ($w_1 = 1, w_2 = 3, w_3 = 2$) or ($w_1 = 2, w_2 = 1, w_3 = 3$).

To solve for the maximum and minimum revenue equilibria in the general problem, we would have to solve the linear program defined by (17) and (20) for every i, k , and l . Although this problem is still very complex, with a minor restriction we can easily solve it.

DEFINITION 1. We say that the preferences of sites i and j are aligned if $G_1(i) > G_1(j)$ implies $G_k(i) - G_l(i) > G_k(j) - G_l(j)$ for every $1 \leq k, l \leq s + 1$.

The assumption of aligned preferences is rather natural. It means that there is a consensus between players about the value of different positions. With this, we can determine the equilibrium ranking of sites.

LEMMA 1. In any SSNE, $G_k(w_1) \geq G_k(w_2) \geq \dots \geq G_k(w_{s+1})$ for any $1 \geq k \geq s + 1$.

To fully describe the equilibria, we also have to assume that sites' valuation for the position they are in is high enough relative to the next site's valuation of the next position. Specifically, we assume that

$$\frac{G_j(w_j) - G_{j+1}(w_j)}{\alpha_j - \alpha_{j+1}} > \frac{G_{j+1}(w_{j+1}) - G_{j+2}(w_{j+1})}{\alpha_{j+1} - \alpha_{j+2}} \quad (21)$$

holds for every $1 \leq j \leq s - 1$. The basic intuition is that the differences in valuations of different positions have to reflect the differences in the objective values of those positions determined by the α order parameters (see Appendix B for more details). With these assumptions, we can describe the SSNE following the path proposed by Varian (2007).

PROPOSITION 2. If all the sites' preferences are aligned and (21) holds, then an SSNE exists. Furthermore,

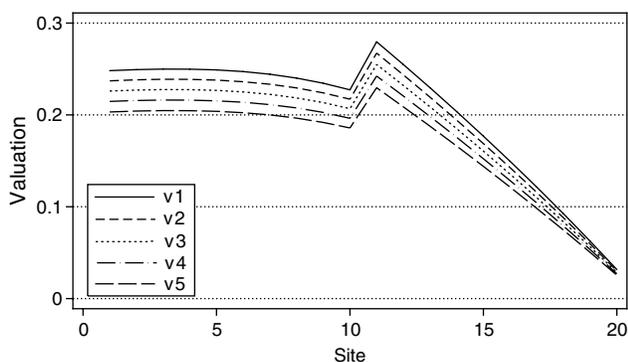
1. The maximum SSNE income of the seller is

$$M(s) = \sum_{j=1}^{s-1} [j(G_j(w_j) - G_{j+1}(w_j))] + sG_s(w_s). \quad (22)$$

2. The maximum SSNE income is equal to the maximum SNE income.

The results are similar to the case in which there is only one sponsored link to bid for. The set and order of winners is determined by two factors. Sites with higher traffic from other sources, such as regular search results, have a lower marginal valuation for traffic. However, sites with higher attractiveness levels value sponsored links higher. It is clear that the order among those sites that do not receive regular search results will be decreasing in γ , that is, $r + 1$,

Figure 2 Sites' Valuation of the Five Sponsored Links



Note. The parameters are $n = 20$, $r = 10$, $s = 5$, $\gamma_i = 0.5 - 0.025(i - 1)$, $\alpha_i = (20 - (i - 1))/232.5$, $\beta = 0.5$, $\phi = 1$, $\delta = 0.6$, and $R(x) = \log(1 + 30x)$.

$r + 2, \dots, n$. However, the top r sites may end up in any position depending on their parameters.

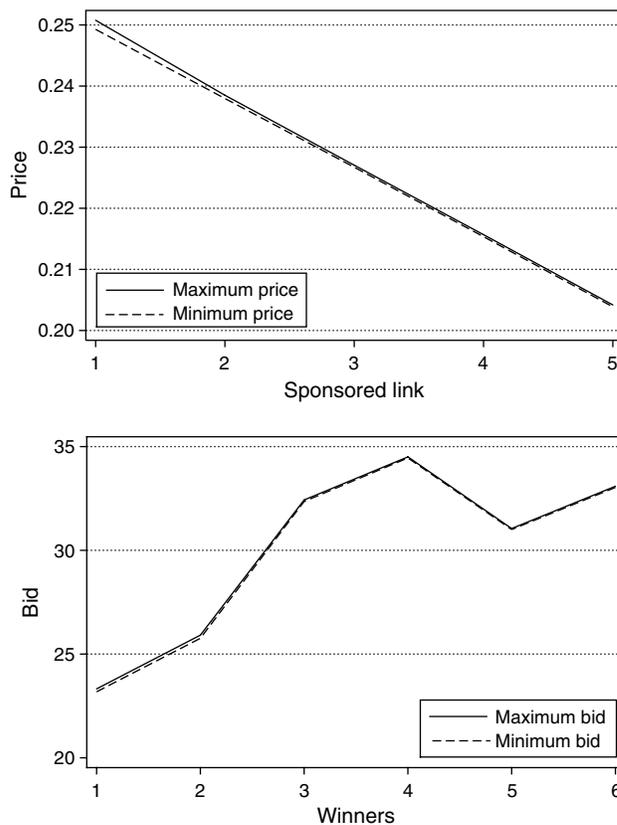
EXAMPLE 2. Let us consider an example of 20 sites competing for five sponsored links with the following parameters: $n = 20$, $r = 10$, $s = 5$, $\gamma_i = 0.5 - 0.025(i - 1)$, $\alpha_i = (20 - (i - 1))/232.5$, $\beta = 0.5$, $\phi = 1$, $\delta = 0.6$, and $R(x) = \log(1 + 30x)$. Then, Site 11 gets the top sponsored link, followed by Sites 12, 3, 4, and 2.

Figure 2 shows the valuations of the 20 sites for the five sponsored links. The parameters are such that Sites 11 and 12 have the highest valuations for the sponsored links because they are the most attractive sites that are *not* listed among the organic results. Because the advertising response function is concave, these sites have a higher marginal valuation for a click. As a result, the winner of the first sponsored link is Site 11, followed by Sites 12, 3, 4, and 2. Figure 3 shows the equilibrium prices the sites pay and the bids they submit. Here, the sites are listed in their order of appearance. It is not surprising that the total fee they pay is decreasing with the position they are in. However, it is interesting to see that higher per-click bids do not automatically lead to a better position. Generally, sites with higher inherent attractiveness do not need to bid too high. However, top sites (such as 3, 4, and 2) still have to bid higher than others for the same position because their higher attractiveness level guarantees them a position on the SE's organic list that, in turn, directs traffic away from the sponsored link. In our example on Figure 3, the sixth site's bid is higher than that of the fifth site, but this site did not manage to fetch a sponsored link.

To illustrate the effect of aversion to advertising ($\phi < 1$), we change the parameters of the example as follows. The return function in this case is linear, but ϕ is lower than one.

EXAMPLE 3. In this example, 20 sites are competing for five sponsored links with the following param-

Figure 3 Fees Paid by and Bids Submitted by the Five Winners in Their Order of Appearance: (11, 12, 3, 4, 2)



Note. The parameters are $n = 20$, $r = 10$, $s = 5$, $\gamma_i = 0.5 - 0.025(i - 1)$, $\alpha_i = (20 - (i - 1))/232.5$, $\beta = 0.5$, $\phi = 1$, $\delta = 0.6$, and $R(x) = \log(1 + 30x)$.

eters: $n = 20$, $r = 10$, $s = 5$, $\gamma_i = 0.5 - 0.025(i - 1)$, $\alpha_i = (20 - (i - 1))/232.5$, $\beta = 0.5$, $\phi = 0.78$, $\delta = 0$, and $R(x) = x$. Then Site 11 gets the top sponsored link, followed by Sites 12, 13, 6, and 7.

The revenue structure is significantly different from that in the previous example, but we obtain a very similar pattern because of the decrease in click through when having both types of links. A linear revenue function and $\phi = 1$ would lead to Sites 1, 2, 3, 4, and 5 winning the sponsored links; however, when $\phi < 1$, top sites do not have an incentive to fight for the sponsored links because they risk losing some of their organic traffic (22% if winning a sponsored link).

In summary, our model explains why sponsored links may exhibit peculiar and seemingly unpredictable patterns on SEs' search pages. The most attractive sites will rank high on the SE's organic list and therefore are likely to benefit less from advertising links. Furthermore, consumers may be averse to sponsored links hurting the sites appearing on both lists who could receive fewer clicks even if they bid high for a sponsored link. These two effects may

cause sites with lower attractiveness levels to win the auction on the sponsored list. However, if the popularity of a site is large enough compared to secondary sites, then these effects are not enough to compensate for the inherent advantage of a site in directing traffic to itself, and top sites may still end up high on the list of sponsored links. Thus, the presence and order of sites on the sponsored-links list is a result of many interacting factors, including the sites' inherent attractiveness and—more importantly—consumers' clicking behavior on the search page. The model shows how behavioral measures of α , β , ϕ , and δ can help SEs, as well as websites, to better optimize their strategies.

5. Extensions

5.1. Heterogeneity in Sites' Valuations

In the main model, we assume that sites value incoming traffic similarly. The rationale behind this assumption is that for a given word, there is a standard rate of converting traffic to revenues, and most sites have the same $R(t)$ function. However, there might be cases in which sites are heterogeneous with respect to their valuation of traffic. For example, as a result of its branding strategy, a company may have an incentive attract more traffic to increase its brand recognition, resulting in higher long-term profits. Here, we examine the implications of heterogeneity in sites' valuation for traffic. Let us assume that Site i has the return function

$$R_i(t) = \vartheta_i R(t), \quad (23)$$

where ϑ_i denotes Site i 's traffic conversion parameter. That is, every site has a similarly shaped traffic return function, but there are individual differences in how sites can make revenues from one visitor. Then, using previous notation, the gain for Site i of winning the sponsored link j is

$$G_j(i) = \vartheta_i [R(J(i)(I(i)\gamma_i\alpha_i + \gamma_i\beta\alpha_j)) - R(I(i)\gamma_i\alpha_i)]. \quad (24)$$

With these modified gain functions, we can apply Proposition 2 (the conditions do not change). The results are similar, but we can observe a simple effect of higher valuation for traffic. It simply boosts sites' willingness to pay for sponsored links; thus, sites with a higher valuation get a better sponsored link. In the extreme case, when sites have similar inherent attractiveness levels ($\gamma_1 - \gamma_n \rightarrow 0$) and an extra visitor results in constant extra revenue ($R(t)$ is linear and ϕ is close to 1), this effect dominates and sites' valuation for traffic (ϑ_i) determines the order of sponsored links. In a typical case, however, this effect is combined with the other previously discussed factors. Let us consider Example 2 again and assume that

$\vartheta_i = 1$ for all i , except for Sites 2 and 3, for which $\vartheta_2 = \vartheta_3 = 1.1$. Then, the order of the five sponsored links changes to 11, 3, 2, 12, 4 from 11, 12, 3, 4, 2. That is, Sites 2 and 3 improved their position because they value an extra visitor relatively higher but still could not get in front of Site 11, which values visitors even higher because it does not receive traffic from organic search results. Also, the relative order of Sites 2 and 3 did not change because their valuations were increased to the same extent.

In summary, heterogeneity in sites' valuation for traffic does have an effect on the order of sponsored links and the bids. A higher valuation leads to higher bids, resulting in a better position in the sponsored-links list. An interesting aspect of these valuations is that, presumably, this is where sites may have more private information in the sense that sites do not perfectly know each other's ϑ_i s. This might result in imperfect information across sites, a case that we explore in Appendix C. For the case of a single sponsored link, we show that the qualitative results do not change as the bidding process leads to revealed valuations.

5.2. Endogenizing the Number of Sponsored Links

So far, we have considered the number of sponsored links displayed by the search engine given. In this section, we compare the search engine's revenue in cases of offering different numbers of links. For the sake of simplicity we assume a linear revenue function and no aversion to advertising; that is, $R(t) = at$ and $\phi = 1$. Then $G_k(i) = \beta\gamma_i\alpha_k$. We assume that the search engine makes a decision about the number of sponsored links and announces it prior to the auction. When it makes the decision, it has to take into consideration two forces. First, if it offers more links for sale, it will receive payments from more sites. However, when the number of links is increased, the traffic flowing through each one goes down. Let us compare the cases when the search engine offers s sponsored links and when it offers $t < s$ instead. If $\beta\alpha_j$ is the traffic going to sponsored link j in the first case, then in the second case, it increases to

$$\beta\alpha'_j = \beta\alpha_j(1 + \beta\alpha_{t+1} + \dots + \beta\alpha_s). \quad (25)$$

As we saw in the previous section, there are usually many equilibria, and the revenue of the SE cannot be determined. Here, we will only compare the maximum revenues the SE can attain by selling a different number of sponsored links.

PROPOSITION 3. *The SE can attain a higher maximum revenue by offering $t < s$ sponsored links instead of s if and only if*

$$\beta(\alpha_{t+1} + \dots + \alpha_s) \left(\sum_{j=1}^t j\gamma_j\alpha_j - \sum_{j=1}^{t-1} j\gamma_j\alpha_{j+1} \right) > \sum_{j=t+1}^s j\gamma_j\alpha_j - \sum_{j=t}^{s-1} j\gamma_j\alpha_{j+1}. \quad (26)$$

Decreasing the number of sponsored links increases the traffic on the remaining ones. Thus, the sites are willing to pay more for them. The left-hand side (LHS) of the inequality is equal to this benefit. However, by forgoing sponsored links $t + 1$ to s , the SE loses $s - t$ advertisers. The resulting loss is the right-hand side (RHS) of the inequality. Note that the RHS is sometimes negative; that is, even without the increased traffic on the remaining links, the SE may have an incentive to decrease the number of links. This is a result of the fact that the value of sponsored links increases in the advertisers' eyes and they are willing to pay more for them.

EXAMPLE 4. Assume that $s = 2$ and $t = 1$. The SE is better off offering one link only iff

$$\beta\alpha_1 > \frac{2\gamma_2 - \gamma_1}{\gamma_1}. \quad (27)$$

In essence, the SE should offer only one sponsored link when the second-highest attractiveness level is relatively low. In particular, if $\gamma_2 < \gamma_1/2$, then the SE is better off selling one link even if the second link still drains traffic. More generally, the SE should only add additional links as long as the attractiveness of an additional site getting that link is relatively high. In other words, if there is a sharp drop in the top attractiveness levels after the i th site, then selling more than i sponsored links may not make sense.

5.3. Dynamic Bidding for Sponsored Links

In the previous models, we assumed that the process through which the sponsored links are assigned is a one-shot game. However, in reality, the auctions for the links take place repeatedly. We cannot always ignore the effects that previous bids and results have on the current auction. An important effect is, for example, that when a site wins a sponsored link, the traffic that it receives through the link may have a lagged effect. Such lagged effects have been documented in Rutz and Bucklin (2007b). Some consumers who get to a website through advertising may become regular customers of the site. If they want to return to the site, they do not need the sponsored link again; they may remember or "bookmark" the site's address. This effect, however, decreases with time. For the sake of simplicity, we assume that it lasts only for one time

period and that there is only one sponsored link. Precisely, if a consumer arrives from the SE to the site in a given time period, then with probability q , she will return in the next period without the use of the search engine. Then, if Site i receives traffic t_i from the search engine in a given period, then the lagged effect of this traffic is qt_i in the next period.

Now let us examine how this effect changes sites' valuations of the sponsored link. If a site did not win the sponsored link in the previous period, then the gain associated with winning it is

$$G_l(i) = R((J(i) + q)I(i)\gamma_i\alpha_i + J(i)\gamma_i\beta\alpha_1) - R((1 + q)I(i)\gamma_i\alpha_i), \quad (28)$$

where we also deal with the lagged effect of regular search results. On the other hand, if the site did win the sponsored link in the previous period, then its gain is

$$G_w(i) = R(J(i)((1 + q)I(i)\gamma_i\alpha_i + (1 + q)\gamma_i\beta\alpha_1)) - R((1 + J(i)q)I(i)\gamma_i\alpha_i + J(i)q\gamma_i\beta\alpha_1). \quad (29)$$

Therefore, if $R(\cdot)$ is strictly concave, then $G_w(i) < G_l(i)$, given that the site had a traffic increase by winning last period.¹⁷ The intuition is that because of decreasing marginal returns, the site values the sponsored link less if it has already won the link in the previous period. To solve the repeated game we use the concept of Markov-perfect equilibrium, where players' actions depend only on the states of the world. In this case, the states represent the possible winners of the auction, and when a site wins the auction, the world moves to that state. In such an equilibrium, forward-looking players choose their strategies to maximize their profits over time using the discount factor d . Let $V_i^{(i)}$ denote Site i 's discounted equilibrium profits counted from a period, when the previous winner is Site j . Sites' payoffs in the current period will be determined by their bids. If Site i does not win the auction, it does not make any profit in the current period; that is, its overall discounted profit will be

$$dV_i^{(w)}, \quad (30)$$

where w is the winner of the current auction. On the other hand, if Site i wins the auction, then it will make a profit of $v_i = G_w(i) - P$ if $i = j$ and $v_i = G_l(i) - P'$ if $i \neq j$, where P and P' are the prices the winner has to pay (these depend on the bids). Therefore, its overall discounted profit will be

$$v_i + dV_i^{(i)}. \quad (31)$$

¹⁷ Although it is technically possible that a site's traffic decreases because of winning a sponsored link (as a result of $\phi < 1$), we ignore this possibility because the site would not submit a positive bid in this case.

In equilibrium, player i chooses its bid to maximize this quantity; that is,

$$V_i^{(j)} = \max_{b_i} (dV_i^{(w)}, v_i + \delta V_i^{(i)}), \quad (32)$$

where w and v_i both depend on b_i .

Because there is only one sponsored link, we can use the first-price auction’s equilibrium and the second-price auction’s symmetric equilibrium concepts interchangeably. We will determine the Markov-perfect first-price Nash equilibria (MFNE) and Markov-perfect second price symmetric Nash equilibria (MSSNE) of the game. Regarding the valuations, let us assume that only the first two sites have a high-enough valuation to win the auction; that is, $G_i(j) < \min(G_w(1), G_w(2))$ for $j \geq 3$. Then, we only have to examine the auction where Sites 1 and 2 bid for the link.

PROPOSITION 4.

1. If $G_1(2) < G_w(1)$, then Site 1 is the winner in every period and

$$G_w(1) \geq P_1 \geq G_1(2). \quad (33)$$

2. If $G_1(1) < G_w(2)$, then Site 2 is the winner in every period and

$$G_w(2) \geq P_1 \geq G_1(1). \quad (34)$$

3. In every other case, the two sites alternate winning.

In essence, if a site values winning the link for a second time more highly than the other site does for the first time, then that bidder is always the winner. Otherwise, the two sites alternately win and lose the auction. The intuition is that when a site wins the link in one period, then its valuation goes down in the next period and the other site is willing to pay more for the link. Now that this other site wins the auction, the valuations will again cross each other, leading to the alternation. Therefore, the only way one site can win the auction in every period is if its valuation dominates the other site’s valuation in the sense that, even after winning it, the site is willing to bid more than its losing competitor.

Interestingly, and consistently with our focus, the organic links play an important role here. In addition to the natural results when the most attractive site always wins (the one which is the first in the organic lists), there are scenarios under which the site in the second position wins all the time. This can happen, for example, if the first hits in the lists are much better than the second positions. In this case, the site that is in lower position on the organic list competes very aggressively for the sponsored link. Also, for words for which the traffic return function $R(t)$ is very steep for the first few visitors and then becomes flat quickly (a very negative second derivative), the site that already has many visitors from organic search

has a low incentive to compete. Furthermore, it is possible that the site in the first position on the organic list fears losing clicks because of $\phi < 1$. The typical situation, however, is that neither site’s valuations dominate the other’s, and consequently, they alternately win the sponsored link. Next, we examine how the strength of the lagged effect (the value of q) affects the outcome.

COROLLARY 2. *There is a $q^* > 0$ such that*

- if $0 < q < q^*$, then the winner is always the same site; and
- if $q^* < q$, then the two sites alternate as winners.

In other words, as the ratio of returning customers increases, at one point the type of equilibrium changes and the two sites start winning alternately. This critical value is smaller if the marginal return on traffic decreases quickly.

Finally, let us compare the search engine’s income in the different cases. Let us assume that $G_1(1) > G_1(2)$; that is, either Site 1 always wins the link or sites alternate winning. Then the search engine’s maximum discounted income (in the two cases, respectively) is¹⁸

$$M_1 = \frac{G_w(1)}{1-d}, \quad (35)$$

$$M_2 = \frac{G_1(1) + dG_1(2)}{1-d^2}. \quad (36)$$

It is worth noting that M_1 and M_2 not only represent the SE’s maximum income in the two cases but also the total surplus of all players (SE and sites) in all the equilibria of the given type.¹⁹ This raises an interesting question: What happens to this surplus when the type of the outcome changes? Are the sites and the SE better off under an alternating winning scenario or with a fixed winner? We compare these values around the boundary of the two regions that separates the alternating and nonalternating equilibria, that is, where $G_w(1) = G_1(2)$.

COROLLARY 3.

$$\lim_{G_w(1)-G_1(2) \rightarrow 0^+} M_1 < \lim_{G_w(1)-G_1(2) \rightarrow 0^-} M_2, \quad (37)$$

and the difference increases in q and d .

We find a discontinuity in the total income at the boundary of the two regions because the SE and the sites are strictly better off in the case of an alternating equilibrium. The intuition is that the alternating assignment of the SE’s traffic is a more efficient allocation than when one site is the winner in every period.

¹⁸ The proof can be found in the proof of Proposition 4.

¹⁹ Individual incomes depend on how this surplus is divided in a given equilibrium. In their data, Yao and Mela (2008) find that the SE gets more than 80% of the total surplus.

This extra revenue is higher if the ratio of returning consumers is higher and if the discount rate is higher. Whether the SE or the sites appropriate this extra revenue depends on the actual bids. The key insight, however, is that *all* players are better off in an alternating equilibrium. Again, knowing consumers' behavior on the search page, the SE can influence the design of the auction to increase the likelihood of such an outcome.

6. Conclusion

In this paper, we have modeled the race for sponsored advertising links on a search engine's page between websites endowed with different attractiveness levels. We argue that the SE's problem cannot simply be described as a multi-item auction. The existence of the organic list on the SE's page represents an important externality for both types of players. In addition to exploring the effect of this externality on the allocation outcomes, we also study two other issues: the endogenous choice of the number of sponsored links and the dynamics of the bidding behavior.

Our key results explain the mechanism that may lead to wildly different patterns observed in the behavior of sponsored links. In particular, top sites who rank high on the SE's search results list are likely to benefit less from advertising links. Furthermore, if consumers are averse to sponsored links, top sites may be hurt by obtaining a sponsored link if visitors are less likely to click on links that appear on both lists. These two effects may cause secondary sites to end up winning the auction on the sponsored list. On the other hand, if the attractiveness of a site is high enough compared to secondary sites, then the above effects are not enough to compensate for the inherent advantage of a site in directing traffic to itself, and top sites may still end up high on the list of sponsored links.

Search engines often claim that not only their organic links but also their sponsored links, create value by displaying the most relevant and attractive links on the top. The reasoning is that the bidding process makes the top sites self-select to the top of the sponsored list. However, as we show, because of the interaction between organic and sponsored links, this is not always the case. Interestingly, when consumers' trust in the sponsored links is lower, it is more likely that the most attractive sites will give up on these links. Thus, low expectations about the quality of sponsored links are generally self-fulfilling. This might explain why the competition between search engines favors only one of the competitors, and why Google has near-monopoly power in this domain. A second mover offering sponsored links that faces somewhat lower consumer expectations on the value

of these links will attract less-relevant sites, and as a result, consumer expectations further deteriorate, increasing in this way the incumbent search engine's advantage.

We also explore three extensions. First, we relax the assumption that sites value traffic uniformly. We find that although sites valuations matter in terms of the actual bids, the basic competitive mechanisms remain the same. Second, endogenizing the number of sponsored links allocated by the SE, we show that the SE can increase traffic flowing through sponsored links by decreasing the number of these links. A decrease in this number increases the value of the links and may result in compensating for the loss associated with a smaller number of links. Finally, we examine a dynamic model in which online advertising has a lagged effect on the site that wins the sponsored link. We identify dynamic bidding patterns that lead to alternating or constant allocations of the sponsored links, depending on the strength of the lagged effect. Interestingly, we find that the SE and the sites together are strictly better off under an alternating equilibrium.

Our analytic results have interesting normative implications. Our core result may help search engines refine the weights attributed to sites' bids for sponsored links. By explicitly measuring the parameters describing consumers' behavior on a search page, the weights attributed to bids can be corrected beyond the sites' attractiveness levels. We also provide insights with respect to when an SE should add/subtract a sponsored link from the page. In particular, we find that this decision primarily depends on the distribution of attractiveness levels across sites. When a sharp drop occurs in this distribution, then the SE should stop adding sponsored links to the page. Finally, our analysis of the dynamic game suggests that the SE should try to promote an alternating bidding pattern between sites. Again, understanding consumers' behavior on the search page and maybe influencing it might help the SE to do so.

Similarly to the SE, bidding sites can also benefit from a deeper understanding of how consumers behave on the search page. Acquiring such information through experimentation can help sites assess the outcome of their bidding strategies and, in this way, lead to significant improvement in budget allocation.

6.1. Limitations

Our model also has a number of limitations. First, when modeling consumer's behavior on the search results page, we assumed that a person either clicks on one link or leaves the page without clicking. In reality, someone can click on one link and then return to the search page if he or she is not satisfied with that particular link and click on another link.

To account for multiple visits, we could apply our assumptions “per visit” and not “per person.” This way, we model each visit separately, and one person can make several visits. These visits to the search page, however, have to be independent in this setting. Note that to some extent, we capture the possibility of dependence between the visits to a particular site by assuming a weakly concave $R(t)$ function. Clearly, it is a limitation of our study that we do not explicitly take into account the possible connection between these visits.

Second, in reality, sites not only place bids according to what they are willing to pay for a click, but they can also set daily or monthly budgets. Then, there is an automatic system that submits the site’s bid continuously until the budget is reached, after which the system automatically withdraws their bids. We do not model this feature, because in our model sites can perfectly estimate how much traffic they get through a sponsored link. However, it is a limitation of our model that it does not consider uncertainty regarding the number of clicks on a link. Furthermore, in the last extension, we only model repeated bidding in a discrete setting in which sites submit bids for each consecutive time period, but not continuously.

Third, we assume that all sites are profit maximizing. However, in reality many noncommercial sites, such as Wikipedia for example, are not for profit. Many of these sites appear in one of the top positions in the organic list but not on the sponsored list. Although we do not fully explore how these sites behave, they would fit in the model by assuming a constant zero-revenue function for them. In this case, they would not participate in the auction for sponsored links. Also, noncommercial sites are presumably worried about the negative effects of advertising and having sponsored links displayed. We could account for this by assuming that these sites have a lower ϕ parameter than others, which will also result in no sponsored links for them according to the model. Third, as we already mentioned before, if top sites have high traffic from sources other than the SE, then the result of secondary sites winning the sponsored links is accentuated. This is also consistent with the phenomenon that top-rated noncommercial informational sites are not present in the sponsored list. A thorough examination of the behavior of noncommercial sites and their interaction with their commercial counterparts remains an interesting topic for future research.

Fourth, our model implicitly assumes a limited time horizon. We have assumed that the SE receives a fixed number of clicks per unit time and that the inherent attractiveness of sites is constant. Clearly, these quantities may evolve in the long run with the evolution

of the Internet and the sites themselves. For example, if a site spends a significant amount on sponsored links, its inherent attractiveness may eventually increase permanently. Our model ignores these long-term strategic effects.

Finally, throughout the paper we have assumed that every consumer is interested in the same topic and the results include the same pages for every query. Obviously, this is rather unrealistic, and the allocation of sponsored links in relation to a given search word changes when multiple *interacting* search words are considered. As reported in Rutz and Bucklin (2007a), most advertisers manage/bid for a bundle of keywords. Websites may offer content in every topic, although their relevance may vary from topic to topic. In other words, the inherent attractiveness may be different for the same site in different topics. For example, Travelocity.com may have a high attractiveness in the context of travel but most likely has a lower one when consumers are searching for home appliances.

The pricing of search advertising is a dynamic field that provides a fertile area for future research. Rather than focusing on various auction mechanisms, our goal was to concentrate on the interaction (conflict) between the SE’s core business as a reliable source of information and its business as an advertiser. Our results provide insights on how to minimize the conflict between these business objectives. Clearly, there are many possible ways in which the present analysis can be extended, including empirical work to test some of the analytic results.

Appendix A. Discussion of Equilibrium Concepts and Auction Mechanisms

As mentioned in the paper, it is important to discuss the various auction types and their equilibrium concepts to be able to provide a definite outcome of the game. Examining the auction types reveals that when competitors’ valuations are known, a first-price auction for a single item typically has an infinity of equilibria. For example, let $v_1 > v_2 > \dots > v_n$ be the valuations of n bidders for a single item. If a first-price auction is applied, then the winner pays his or her bid. In equilibrium, the winner is always player 1 and the winning bid, b_1 , can take any value in the $(v_2, v_1]$ interval. Thus, the auctioneer’s revenue is between v_2 and v_1 . We denote this type of equilibrium by FNE (first-price Nash equilibrium).

In the case of a second-price single-item auction, anyone can win the auction in a Nash equilibrium (SNE). If every player bids 0 except player i , who bids $v_0 > v_1$, then the winner is player i , who has to pay nothing. In general, the second-highest bid is always below v_1 , so the auctioneer’s revenue is somewhere between 0 and v_1 . To restrict the possible outcomes of a second-price auction, Varian (2007) introduced the notion of symmetric equilibria for multi-item second-price auctions, (SSNE), which is a subset of the pure-strategy Nash equilibria. In such an equilibrium, the player

in position k is better off paying the bid of the player in position $k + 1$ than he or she would be in position l paying the bid of player $l + 1$. This is a stronger restriction than in an SNE for moving up in the ranking because in an SNE a player is only supposed to be better off paying bid $k + 1$ for position k than paying bid l for position l . Because bid l is higher than bid $l + 1$, an SSNE is always an SNE, but the opposite is not true. According to Varian (2007), in an SSNE, the order of winners is always $1, 2, 3, \dots$; that is, in the case of a single item the winner is always player 1. Furthermore, the auctioneer's maximum SSNE revenue is the same as the maximum SNE revenue and is equal to v_1 in the case of a single item. Because the equilibria in a first-price single-item auction (FNE) and symmetric equilibria in a second-price single-item auction (SSNE) give the same results for the bid orders and *maximum* revenues of the seller, we can use the two concepts interchangeably for our analysis if there is only one sponsored link. For multiple links, the FNE usually does not exist; so in this case, we will always use the SSNE as the equilibrium concept. That is, we will restrict the search engine's strategy space to running second-price auctions.

Appendix B. Proofs

PROOF OF CLAIM 1. The search engine wishes to maximize the income from the s winners of the sponsored links. Given the order of sites, it is obviously optimal to set the p_i s to the maximum, that is, $p_i = b_i$, because it does not affect sites' bidding strategies because sites only place bids once, in the first stage of the game. Regarding the order of sites, if Site i acquires a sponsored link, the search engine will receive a total payment of $\beta A(i)F_i$ from that site, where $F_i = \gamma_i b_i J(i)(1 - \delta I(i))$. The F_i values are site specific and only depend on the site's parameters, whereas the $A(i)$ values are determined by the search engine when it assigns the sponsored links. To maximize $\beta \sum_{i=1}^n A(i)F_i$, the SE has to assign the α s in a decreasing order of the F_i values. \square

PROOF OF PROPOSITION 1. As we have discussed before, the winner—both in an FNE and an SSNE—is the site with highest valuation. The payment of the winner is between the first and second valuations. \square

PROOF OF LEMMA 1. If sites' preferences are aligned, then (20) yields $G_1(w_l) \geq G_1(w_m)$ for every $l < m$, proving the lemma. \square

PROOF OF PROPOSITION 2. To prove the existence of an SSNE, we have to show that there exist $P_1 \geq P_2 \geq \dots \geq P_s$, such that they satisfy inequalities (18) and (19) for every $1 \leq k < l \leq s$. We will show that if the sites' preferences are aligned, then it is enough to check that $P_1 \geq P_2 \geq \dots \geq P_s$ satisfy a subset of them, namely, the following inequalities, for every j :

$$G_j(w_j) - G_{j+1}(w_j) \geq P_j - P_{j+1} \geq G_j(w_{j+1}) - G_{j+1}(w_{j+1}). \quad (B1)$$

We have to show that all the inequalities in (18) and (19) follow from those in (B1). Let $1 \leq k < l \leq s$ be arbitrary indices. Summing (B1) for $j = k$ to l , we get

$$\begin{aligned} \sum_{j=k}^{l-1} [G_j(w_j) - G_{j+1}(w_j) \geq P_k - P_l] \\ \geq \sum_{j=k}^{l-1} [G_j(w_{j+1}) - G_{j+1}(w_{j+1})]. \quad (B2) \end{aligned}$$

Because the preferences are aligned, $G_j(w_k) - G_{j+1}(w_k) \geq G_j(w_j) - G_{j+1}(w_j)$ for $j > k$; therefore, we obtain

$$G_k(w_k) - G_l(w_k) \geq P_k - P_l, \quad (B3)$$

and similarly,

$$P_k - P_l \geq G_k(w_l) - G_l(w_l). \quad (B4)$$

We have shown that the system given by (18) and (19) is equivalent to that defined by (B1). That is, it is always enough to check whether a site wants to get to a position that is one higher or lower. Therefore, given that (17) holds, the values of $P_j - P_{j+1}$ can be chosen arbitrarily from the intervals given in (B1), fixing $P_{s+1} = 0$. In (21), we basically assume that selecting the maximum values does not violate (17). Thus, we get the second part of the proposition by summing the left-hand sides of (B1) in the following way.

$$\sum_{i=1}^s P_i = sP_s + \sum_{j=1}^{s-1} j(P_j - P_{j+1}). \quad (B5)$$

For the fourth part, let us note that every SSNE is an SNE; therefore, the maximum SNE income is at least as high as the maximum SSNE income. For the other direction, let P_i^N denote the expenditure of Site i in an SNE with maximum revenue and let P_i^S denote the same expenditure in a maximum revenue SSNE. From the previous part, we know that

$$P_j^S = P_{j+1}^S + G_j(w_j) - G_{j+1}(w_j). \quad (B6)$$

However, according to the definition of an SNE,

$$P_j^N \leq P_{j+1}^N + G_j(w_j) - G_{j+1}(w_j). \quad (B7)$$

Because $G_{s+1}(w_s) = 0$,

$$P_s^N \leq G_s(w_s) = P_s^S. \quad (B8)$$

Then, it is easy to show recursively that $P_i^N \leq P_i^S$, completing the proof. \square

PROOF OF PROPOSITION 3. According to Proposition 2, the maximum equilibrium revenue of the SE, in case of selling s links, is

$$M(s) = \beta \left(\sum_{j=1}^s j\gamma_j \alpha_j - \sum_{j=1}^{s-1} j\gamma_j \alpha_{j+1} \right). \quad (B9)$$

If the SE decides to instead sell only t links, the traffic on the remaining links will increase by a factor of $(1 + \beta(\alpha_{t+1} + \dots + \alpha_s))$. Therefore, the maximum equilibrium revenue will be

$$(1 + \beta(\alpha_{t+1} + \dots + \alpha_s)) \beta \left(\sum_{j=1}^t j\gamma_j \alpha_j - \sum_{j=1}^{t-1} j\gamma_j \alpha_{j+1} \right) \quad (B10)$$

in this case. Comparing the two quantities, we get the expression in the proposition. \square

PROOF OF PROPOSITION 4. We will assume without loss of generality that $G_1(1) > G_1(2)$. The proof of the opposite case is straightforward. First, we prove the third part of the proposition, that is, identify the conditions necessary for an alternating equilibrium. In such an equilibrium, bidding strategies are such that if Site i has won the previous auction, then Site $j = 3 - i$ is the current winner. Let $P^{(i)}$ denote

the fee that Site $j = 3 - i$ has to pay in the auction when Site i is the previous winner. Let $V_i^{(j)}$ denote the discounted equilibrium profits of Site i from a given period when Site j is the previous winner. In an alternating equilibrium,

$$\begin{aligned} V_1^{(1)} &= dV_1^{(2)}, \\ V_1^{(2)} &= G_I(1) - P^{(2)} + dV_1^{(1)}, \\ V_2^{(1)} &= G_I(2) - P^{(1)} + dV_2^{(2)}, \\ V_2^{(2)} &= dV_2^{(1)}. \end{aligned}$$

Therefore,

$$\begin{aligned} V_1^{(2)} &= \frac{G_I(1) - P^{(2)}}{1 - d^2}, \\ V_2^{(1)} &= \frac{G_I(2) - P^{(1)}}{1 - d^2}. \end{aligned}$$

The sufficient and necessary conditions these valuations and prices have to satisfy are that in a given auction, the winner has to have a higher valuation and the fee paid by the winner must fall between the two players' valuations (both in an MFNE and MSSNE). For example, if the previous winner is Site 1, then the current winner must be Site 2; therefore,

$$G_w(1) + d(V_1^{(1)} - V_1^{(2)}) \leq P^{(1)} \leq G_I(2) + d(V_1^{(1)} - V_1^{(2)}) \quad (\text{B11})$$

must hold. Plugging the corresponding formulas, we obtain

$$G_w(1) - \frac{1-d}{1-d^2}(G_I(1) - P^{(2)}) \leq P^{(1)} \leq G_I(2). \quad (\text{B12})$$

Comparing the valuations in a period when Site 2 is the previous winner, we get a similar inequality,

$$G_w(2) - \frac{1-d}{1-d^2}(G_I(2) - P^{(1)}) \leq P^{(2)} \leq G_I(1). \quad (\text{B13})$$

The set defined by (B12) and (B13) is a two-dimensional simplex. It is easy to see that it is nonempty iff $G_I(2) \geq G_w(1)$ (given the other restrictions on the parameters).

The maximum discounted income of the seller depends on the first period of the game. Let P_t denote its income in period t . If Site 1 is the first winner, then it would be

$$\sum_{t=1}^{\infty} d^{t-1} P_t = \frac{P^{(2)} + dP^{(1)}}{1 - d^2}. \quad (\text{B14})$$

If Site 2 is the first winner, then it is

$$\frac{P^{(1)} + dP^{(2)}}{1 - d^2}. \quad (\text{B15})$$

We determine the maximum for both and consider the higher value. Clearly, because Site 1 has higher valuations, the SE's income will be higher if Site 1 is the first winner. Maximizing $P^{(2)} + dP^{(1)}$ on the simplex defined by (B12) and (B13), we get

$$M_2 = \frac{G_I(1) + dG_I(2)}{1 - d^2}. \quad (\text{B16})$$

The first part of the proposition can be proven by following the same steps. However, it is obvious that because in both states Site 1 has a higher valuation, it is always the

winner. Then the price paid must be in the given range, yielding the stated maximum income. \square

PROOF OF COROLLARY 2. The values of $G_I(1) > G_I(2)$ are independent of q . When $q = 0$, $G_I(i) = G_w(i)$, and as q increases, $G_w(1)$ decreases. Let q^* be the unique solution of

$$\begin{aligned} R((1+q)I(1)\gamma_1\alpha_1 + (1+q)\gamma_1\beta\alpha_1) - R((1+q)I(1)\gamma_1\alpha_1 + q\gamma_1\beta\alpha_1) \\ = R((1+q)I(2)\gamma_2\alpha_2 + \gamma_2\beta\alpha_1) - R((1+q)I(2)\gamma_2\alpha_2). \end{aligned}$$

Then, for $0 < q < q^*$, we get the first case in Proposition 4, and for $q^* < q$, we get the second case. \square

PROOF OF COROLLARY 3. Fixing $G_I(2)$ in Proposition 4, we can establish

$$\lim_{G_w(1) - G_I(2) \rightarrow 0^+} M_1 = \frac{G_I(2)}{1 - d}, \quad (\text{B17})$$

$$\begin{aligned} \lim_{G_w(1) - G_I(2) \rightarrow 0^-} M_2 &= \frac{G_I(1) + dG_I(2)}{1 - d^2} \\ &= \frac{G_I(2)}{1 - d} + \frac{G_I(1) - G_I(2)}{1 - d^2}. \end{aligned} \quad (\text{B18})$$

Hence, the difference is

$$0 < \frac{G_I(1) - G_I(2)}{1 - d^2} = \frac{G_I(1) - G_w(1)}{1 - d^2}, \quad (\text{B19})$$

which clearly increases in q and d . \square

Appendix C. Bidding for a Single Sponsored Link Under Imperfect Information

Throughout the paper we assume that sites have perfect information about each other—that is, they know how attractive their competitors are. Although we justify this assumption by arguing that repeated bidding reveals attractiveness levels and that sites can run experiments or purchase market research, here we examine the case when Site i only knows γ_i , but not γ_j , for $i \neq j$. We assume that sites bid for a single sponsored link to make the problem tractable. Aside from incomplete information, the model is the same as the one we present in §4.1. Each site submits its bid (b_i) for the per-click price of the sponsored link. The SE then orders the sites according to their bids correcting for CTRs and awards the sponsored link to the first one. If site i is the winner and is followed by site j in the order, then it has to pay a per-click price of $p_i = b_j(\gamma_j F_j / (\gamma_i F_i))$. Then, using the notations of §4.1, the following proposition summarizes the results.

PROPOSITION 5. *The winner of the sponsored link is site w_1 , paying a total fee of*

$$P_1 = G(w_2), \quad (\text{C1})$$

and the bids are

$$b_i = \frac{G(i)}{\beta\alpha_1\gamma_i J(i)(1 - \delta I(i))}. \quad (\text{C2})$$

PROOF. The valuation of site i for the link, taking into account the expected traffic, is $G(i)$. Although the bids are submitted for price per click, the SE and the site itself can calculate how much it is willing to pay for the total traffic by multiplying the bid b_i by F_i . That is, $b_i\gamma_i F_i$ will represent the site's "bid" for the total value of obtaining the link. As in the case of second-price sealed-bid auctions, truth telling

is a dominant strategy; that is, $G(i) = b_i \gamma_i F_i$ for each site. Therefore, $b_i = G(i)/(\gamma_i F_i)$, and the winner is the site with the highest $G(i)$. \square

Because there is only one sponsored link, we were able to use the well-known idea that truth telling is an equilibrium outcome in second-price auctions. The results are almost identical to those in Proposition 1 and Corollary 1. The winner is the same, but in this case we get a unique bid for the winner in equilibrium, which is equal to the maximal winning bid in Proposition 1.

References

- Athey, S., G. Ellison. 2008. Position auctions with consumer search. Working paper, Harvard University, Cambridge, MA. <http://kuznets.harvard.edu/~athey/position.pdf>
- Bhargava, B., M. Jenamani, Y. Zhong. 2005. Counteracting shill bidding in online English auction. *Internat. J. Cooperative Inform. Systems* **14**(2/3) 245–263.
- BusinessWeek*. 2007. Where is Microsoft Search? (April 2) 30.
- Chen, Y., C. He. 2006. Paid placement: Advertising and search on the Internet. NET Institute Working Paper 06-02, <http://ssrn.com/abstract=936472>.
- Edelman, B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *Amer. Econom. Rev.* **97**(1) 242–259.
- Edwards, S. M., H. Li, J.-H. Lee. 2002. Forced exposure and psychological reactance: Antecedents and consequences of the perceived intrusiveness of pop-up ads. *J. Advertising* **31**(3) 83–95.
- Emarketer*. 2009. What happened to search spending in 2008? (January 12), <http://emarketer.com/Article.aspx?R=1006853>.
- Feng, J. 2008. Optimal mechanism for selling a set of commonly ranked objects. *Marketing Sci.* **27**(3) 501–512.
- Feng, J., Z.-J. M. Shen, R. L. Zhan. 2007. Ranked items auctions and online advertisement. *Production Oper. Management* **16**(4) 510–522.
- Goldfarb, A., C. Tucker. 2007. Search engine advertising: Pricing ads to context. NET Institute Working Paper 07-23, <http://ssrn.com/abstract=1021451>.
- Katona, Z., M. Sarvary. 2008. Network formation and the structure of the commercial World Wide Web. *Marketing Sci.* **27**(5) 764–778.
- Lutz, R. J. 1985. Affective and cognitive antecedents of attitude towards the ad: A conceptual framework. L. F. Alwitt, A. A. Mitchell, eds. *Psychological Processes and Advertising Effects: Theory, Research, and Applications*. Lawrence Erlbaum Associates, Hillsdale, NJ, 45–64.
- Maskin, E. S., J. G. Riley. 2000. Asymmetric auctions. *Rev. Econom. Stud.* **67**(3) 413–438.
- Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1) 58–73.
- Ofek, E., M. Sarvary. 2003. R&D, marketing and the success of next generation products. *Marketing Sci.* **22**(3) 355–370.
- Rutz, O. J., R. E. Bucklin. 2007a. A model of individual keyword performance in paid search advertising. Working paper, <http://ssrn.com/abstract=1024766>.
- Rutz, O. J., R. E. Bucklin. 2007b. From generic to branded: A model of spillover dynamics in paid search advertising. Working paper, <http://ssrn.com/abstract=1024765>.
- Schlosser, A. E., S. Shavitt, A. Kanfer. 1999. Survey of Internet users' attitude towards Internet advertising. *J. Interactive Marketing* **13**(3) 34–54.
- Varian, H. R. 2007. Position auctions. *Internat. J. Indust. Organ.* **25**(6) 1163–1178.
- Villas-Boas, J. M. 1993. Predicting advertising pulsing policies in an oligopoly: A model and empirical test. *Marketing Sci.* **12**(1) 88–102.
- Wilbur, K. C., Y. Zhu. 2009. Click fraud. *Marketing Sci.* **28**(2) 293–308.
- Yao, S., C. F. Mela. 2008. A dynamic model of sponsored search advertising. NET Institute Working Paper 08-16, <http://ssrn.com/abstract=1285775>.
- Zeithammer, R. 2006. Forward-looking bidding in online auctions. *J. Marketing Res.* **43**(3) 462–476.
- Zeithammer, R., P. Liu. 2008. When is auctioning preferred to posting a fixed selling price? Working paper, http://www.anderson.ucla.edu/faculty/robert.zeithammer/auctioned_goods_nov06.pdf.