

Externalities in Keyword Auctions: an Empirical and Theoretical Assessment *

Renato Gomes
Northwestern University
Department of Economics
rdbgomes@gmail.com

Nicole Immorlica
Northwestern University
Department of EECS
nicimm@gmail.com

Evangelos Markakis
Athens University of
Economics and Business
Department of Informatics
markakis@gmail.com

ABSTRACT

It is widely believed that the value of acquiring a slot in a sponsored search list (that comes along with the organic links in a search engine's result page) highly depends on who else is shown in the other sponsored positions. To capture such externality effects, we consider a model of keyword advertising where bidders participate in a Generalized Second Price (GSP) auction and users perform ordered search (they browse from the top to the bottom of the sponsored list and make their clicking decisions slot by slot). Our contribution is twofold: first, we use impression and click data from Microsoft Live to estimate the ordered search model. With these estimates in hand, we are able to assess how the click-through rate of an ad is affected by the user's click history and by the other competing links. Further, we compare the clicking predictions of our ordered search model to those of the most widely used model of user behavior: the separable click-through rate model. Second, we study complete information Nash equilibria of the GSP under different scoring rules. First, we characterize the efficient and revenue-maximizing complete information Nash equilibrium (under any scoring rule) and show that such an equilibrium can be implemented with any set of advertisers if and only if a particular weighting rule that combines click-through rates and continuation probabilities is used. Interestingly, this is the same ranking rule derived in [11] for solving the efficient allocation problem. On the negative side, we show that there is no scoring rule that implements an efficient equilibrium with VCG payments (VCG equilibrium) for all profiles of valuations and search parameters. This result extends [8], who argue that the rank-by-revenue GSP does not possess a VCG equilibrium.

1. INTRODUCTION

Sponsored search advertising is a booming industry that accounts for a significant part of the revenue made by search engines. For queries with most commercial interest, Google, Yahoo! and MSN Live make available to advertisers up to three links above the or-

*Most of this work was done while all authors were at the Center for Math and Computer Science (CWI), Amsterdam. The project was partially funded by the Microsoft external research program "Beyond Search: Semantic Computing and Internet Economics".

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

ganic results (these are the *mainline slots*), up to eight links besides the organic results (*sidebar slots*) and, more recently, MSN Live even sells links below the organic results (*bottom slots*).

As such, an advertiser that bids for a sponsored position is seldom alone; and is usually joined by his fiercest competitors. Indeed, it is widely believed that the value of acquiring a sponsored slot highly depends on the identity and position of the other advertisers. Putting it differently, advertisers impose *externalities* on each other, which affect their click-through rates and might have consequences on their bidding behavior.

The literature on sponsored search auctions mostly assumes click-through rates are separable, i.e., the click through rate of a bidder is a product of two quantities, the first expressing the quality of the bidder and the second the quality of the slot she occupies. Such models cannot capture the externalities that one advertiser imposes on the others. To capture these externality effects, we depart from the separable model and study a model that integrates the users' search behavior and the advertiser's bidding behavior in the Generalized Second Price (GSP) auction run by search engines. We will assume that users perform *ordered search*, which means that (i) they browse the sponsored links from top to bottom and (ii) they take clicking decisions slot by slot. After reading each ad, users decide whether to click on it or not and, subsequently, decide whether to continue browsing the sponsored list or to simply skip it altogether (for a formal definition and motivation for this model, see Section 2). With this formulation, we are able to estimate *continuation probabilities* for each ad (which are simply the probabilities of continuing searching the sponsored list after clicking or not on some ad) and *conditional click-through rates* for each ad (which tell the probability of a click conditional on the user's previous clicking history). Continuation probabilities capture *position externalities*, that is, they capture the negative impact that top links impose on the click-through rates of bottom links (as users stop browsing either because their search needs were already fulfilled or because they got tired of previous bad matches). In turn, conditional click-through rates capture *information externalities*, as we can assess how the information collected by the user by clicking on one given link impacts the click-through rates of the other links he eventually reads.

On the auction side of the model, advertisers submit their bids taking click-through rates as implied by ordered search. As prescribed by the rules of the GSP, search engines then multiply each bid by a weight defined by a *scoring rule* (which solely depends on each advertiser's characteristics), producing a score for each advertiser. Advertisers are then ranked by their score; slots are assigned in decreasing order of scores and each advertiser pays per click the minimum bid necessary to keep his position.

We use the model described above to make both empirical and

theoretical contributions. On the empirical side, we used three months of impression and clicking data from Microsoft Live to estimate the ordered search model. In this version we report our findings from three selected search terms: *ipod*, *diet pill* and *avg antivirus*. We plan to substantiate further our findings in a longer version of this work. For each of the selected keywords, we selected the logs in which the most clicked advertisers occupied the mainline slots. Our main empirical findings can be summarized as follows: first, our dataset suggests that both position and information externalities are economically and statistically significant - and the returns to keyword advertising (in terms of clicks) strongly depend on the identity of the other advertisers. Secondly, our estimates suggest that users roughly divide in two groups: the first group has a low clicking probability and usually drops the sponsored list without going through all the mainline slots. In contrast, the second group of users clicks more often and tends to read most of the sponsored links (price research behavior). Thirdly, we made a non-parametric comparison of the clicking predictions of our ordered search model to those of the separable click-through rate model (the most popular model of user’s behavior). We found that the ordered search model performs as well, and often better, than the separable click-through rate model for all keywords considered. However, this was to some extent expected as the ordered search model has more parameters than the separable one and hence this makes it easier to fit the data.

On the theoretical side, we investigate the complete information equilibria of the GSP (with click-through rates implied by ordered search) under different scoring rules. First, we characterize the efficient and revenue-maximizing complete information Nash equilibrium (under any scoring rule). We then show that this equilibrium can be implemented under any valuation profiles and advertiser’s search parameters if and only if the search engine ranks bids using a particular weighting rule that combines click-through rates and continuation probabilities. Interestingly, this is the same ranking rule derived in [11] for solving the efficient allocation problem (in a non-strategic environment).

We season this positive result with an impossibility theorem: we show there is no scoring rule that implements an efficient equilibrium where advertisers pay their VCG payments for all valuations and search parameters (this is by far the most analyzed equilibrium of the separable click-through rate model). This extends an observation first made by [8], who argue that such an equilibrium does not exist in the rank by revenue GSP.

1.1 Comparisons with Related Work

Ad auctions became an active area of research in the past few years due to their important role as a revenue source for search engines. The initial theoretical literature [2, 6, 16] studied the equilibria of these auctions under a model of separable click-through-rates. In these models, there are slot-dependent and advertiser-dependent click-through-rates, and the click-through-rate of an advertiser in a given slot is simply the product of these two base click-through-rates. As such, these models can not account for the externalities that advertisers impose on each other.

The issue of externalities in ad auctions has recently attracted quite a bit of attention from the research community [3, 1, 11, 8, 7]. Initial studies were largely theoretical, and involved proposing models for user-search behavior that would explain externalities. Athey and Ellison [3] proposed one of the first such models. In their work, they assume that users search in a top-down manner and that clicking is costly. They then derive the resulting equilibria. Closely related are the cascade models of Aggarwal et al [1] and Kempe and Mahdian [11]. These models associate with each ad a click-

through-rate as well as a continuation probability representing the probability that a user continues the search after viewing the given ad. They then proceed to solve the winner determination problem in their models. Recently Giotis and Karlin [8] studied the equilibria of the cascade model in GSP auctions.¹

Our model of ordered search generalizes these previous models slightly by allowing click-through-rates and continuation probabilities to depend on the clicking history of the user. This enables us to model both position externalities as well as information externalities. Our empirical work shows that both effects are significant

Our contribution to the existing literature is two-fold: we use our model to document externalities empirically and then proceed to study equilibria of various auction types analytically (using the simplified cascade model to keep the problem tractable). This is the first paper to empirically document externalities in sponsored search. In a subsequent work, [9] estimated a model of unordered search in which users read all advertisements before choosing a subset of them to click on. We believe this is a valid and worthwhile exploring model of users’ behavior. We nevertheless think that ordered search is a more natural starting point. Indeed, it is hard to reconcile the assumption that users perform unordered search with advertisers’ competition to obtain the top positions (why to pay more to get a top slot if users read the whole list anyway?). Moreover, unlike [9], we allow click-through rates to depend on the click history of users (this captures users’ learning by browsing).

An interesting counterpart of this work is [5], which provides experimental evidence of position bias in organic search. In this paper, the authors deliberately flip result positions of a major search engine to compare four different models of users’ browsing behavior. They conclude that a cascade model, in which users browse from top to bottom and drop the search after any click, offers better fit than the separable click-through rate model.

On the theoretical side, this work advances the equilibrium analysis of the GSP in the presence of externalities. The previous theoretical literature to study GSP equilibria [2, 6, 14, 16, 12] mostly focused on the separable click-through-rate model. The only exceptions are Athey and Ellison [3] and Giotis and Karlin [8], mentioned above. Our work differs from these papers as, instead of studying equilibria of particular weighing rules, we analyze what weighing rules can implement particular equilibria. As such, this work also pioneers the analysis of scoring rules in the presence of externalities. In the context of the separable model, [2, 12] and [13] compared various bid weighing procedures regarding their revenue properties.

2. THE ORDERED SEARCH MODEL

In our model, we are given a set $N = \{1, \dots, n\}$ of advertisers who must be placed on up to K slots. Throughout the remainder of this empirical section, we will investigate scenarios with $n = 3$ advertisers and $K = 3$ slots, returning to the general model in the theoretical section.

In order to study externalities in sponsored advertising, we develop a model of users’ behavior that assumes ordered search. The main elements of this model are, first, that users make their choices about clicking on sponsored links by analyzing one link at a time and, secondly, that they browse sponsored results from top to bottom. Our focus on such an ordered search model is motivated by

¹Cascade (or sequential search) models had been previously studied in the labor economics and industrial organization literatures: [4] and [15] are excellent surveys. Unlike the applications considered in these literatures, where firms choose prices and wages respectively, we assume that advertisers have no strategic options (qualities in sponsored search are exogenous).

various reasons. First, as the work of [5] demonstrates, position bias is present in organic search. In particular [5] compares a sequential search model with four other models (including the separable model) and concludes that sequential search provides the best fit to the click logs they have considered. Secondly, sequential search is further substantiated as a natural way to browse through a list of ads by the eye-tracking experiments of Joachims *et al.* [10], where it is observed that users search and click in a top down manner. Moreover, as the value per click of each advertiser tends to be correlated with its relevance, ordered search is a good heuristic for users (see [3]).

Given such a model, the users will typically not click on all the ads of the list, as it is costly both in terms of time and cognitive effort to go through a website and assimilate its content. For this reason, users only click on a link if it looks good enough to compensate for its browsing cost. Moreover, users typically change their willingness to incur this browsing cost as they collect new information through their search, and hence the decision about whether to continue reading ads naturally depends on the click history of the user. To formalize these ideas, we denote the click history of users as they browse through the sponsored links by $H = \{j : \text{link } j \text{ received a click}\}$.² Here we will focus on two types of externalities:

Information Externalities. An ad imposes information externalities on others by providing a user who has clicked on his link with information regarding the search – e.g., prices or product reviews. This, in turn, affects the user’s willingness to click on all links displayed below in the sponsored search list. To make these points formally, let’s denote the expected quality of slot j by u_j . In order to save on browsing costs, a searcher with click history H clicks on link j only if its perceived quality exceeds some optimal threshold, which we denote by T_H . We set $H = \{\emptyset\}$ if no links were previously clicked (no extra information gathered through search), $H = \{j\}$ if only link j was clicked and $H = \{j, k\}$ if links j and k were clicked in this order. We let the clicking threshold T_H on the ad’s perceived quality depend on the information gathered by the searcher in his previous clicks, but assume that T_H is not affected by the precise order of clicks. That is, we impose that $T_{\{j,k\}} = T_{\{k,j\}}$.

In addition, we summarize any user specific bias towards a link by the random term ε^{ij} . Hence, a user with click history H that reaches the slot occupied by advertiser j clicks on it if and only if

$$u_j - \varepsilon^{ij} \geq T_H.$$

We assume that the idiosyncratic preference parameters ε^{ij} are independently and identically distributed across bidders and advertisers, with a cumulative distribution function F . Thus, the probability that a searcher i with click history H gives a click on link j is given by:

$$F_j(H) \equiv \text{Prob} \{\varepsilon^i \leq u_j - T_H\} = F(u_j - T_H).$$

By virtue of browsing from the top, users have no previous clicks when they analyze the first slot. Hence, if advertiser j occupies the first position, his chance of getting a click, which we call *click-through rate*, is $F_j \equiv F_j(\{\emptyset\})$.

The difference between advertiser j ’s baseline click-through rate, F_j , and his conditional click-through rate, $F_j(H)$, $H \neq \emptyset$, indicates the impact of information externalities.

Position Externalities. An ad additionally imposes externalities on other ads by virtue of its position in the ordered search list. This

²Note we abstract away order information; i.e., we assume a user’s behavior depends on past clicks, but not on the order in which the clicks were made.

can happen in one of two manners: first, the user may tire of the search if the ads he has read appear to be poorly related to the search term; second, the user may leave the search if an ad he has read and clicked on has satisfied his search need. We capture the first effect with a parameter, λ_j , that indicates the probability a user keeps browsing the sponsored links after reading ad j and *choosing not to click on it*. We capture the second effect with a parameter, γ_j , that indicates the probability a user keeps browsing the sponsored links after *clicking link j* . The parameters λ_j and γ_j are referred to as the *continuation probabilities* of ad j and jointly capture its position externalities imposed on the ads that follow.

Note that, unlike many models in the literature, in our model the position externalities may depend on both the advertiser and clicking behavior of the user.

We model the user behavior for a given sponsored list using the above parameters as follows. She reads the first ad A_1 in the list and clicks on it with probability F_{A_1} . Conditional on clicking on A_1 , she reads the second ad A_2 with probability γ_{A_1} and clicks on it with probability $F_{A_2}(\{A_1\})$. Conditional on not clicking on A_1 , she reads the second ad with probability λ_{A_1} and clicks on it with probability F_{A_2} . Thus, the probability she clicks on ad A_1 is simply F_{A_1} while the probability she clicks on ad A_2 is

$$(1 - F_{A_1})\lambda_{A_1}F_{A_2} + F_{A_1}\gamma_{A_1}F_{A_2}(\{A_1\}).$$

This behavior extends to multiple advertisers in the natural way.

2.1 Identification

We will now describe the structure of our data and discuss the identification of the model described above. In this version we have included in our data set all the impression logs in which users searched for *ipods*, *diet pills* or *antivirus*. We will incorporate more keywords in a future version of this work. Within each of these queries, we selected the impressions in which the three most clicked advertisers are displayed in the mainline slots (which are the three top positions in a search results page).

A few observations are in order: first, the appeal of an advertiser (which translates into her click-through rate) fundamentally depends on the search objective of the user (for example, the purchase of an ipod). As a consequence, a model of user behavior makes most sense when restricted to a single search objective. Accordingly, our data set only contains clicking logs of advertisers displayed along the results of some specific search objective. The same search objective may be expressed by different queries, though: *ipod*, *buy ipod*, *ipod purchase*, *apple ipod*, *cheap ipod* are all different ways of expressing the same search objective. Hence, we include in our data set all click logs with our target keyword in the search query.

Second, we only analyzed, within each query, the three most effective advertisers. The reason for this is two-fold. First, in order to identify the model described above we need some variation on clicking histories that only the most popular advertisers jointly exhibit (we will formalize this observation in Lemma 1 below). There was an insufficient number of observations with variation for the four most popular advertisers. Second, the number of parameters in our model grows exponentially with the number of advertisers, and so it becomes more difficult, experimentally, to track and maintain the parameters.

Third, our ordered model makes the implicit assumption that all users read the first sponsored link before dropping the sponsored search results. In fact, as the eye tracking experiments in [10] attest, these are the most visible links to the users. Even if this were not the case, our model would still capture negative externalities, as baseline click-through rates would be underestimated, while the

estimates for conditional click-through rates would still be consistent.

Let the three most effective sponsored search advertisers for a given keyword be denoted by j , k and l . In our database, when a user submits a query, he sees a sponsored list displaying ads of at least two advertisers among j , k and l in some specific order. The user may click none, one or more of these links, also in some specific order. We denote such an event by a pair of tuples, each of three elements from the set $\{j, k, l, \emptyset\}$. The first tuple denotes the advertisers that were displayed and the second tuple denotes the advertisers that were clicked. If a slot is left empty, or an advertiser is not clicked, we use the symbol \emptyset to mark this in the tuple. If, for example, only advertisers j and k were displayed (in the first two slots) and only k was clicked, we denote this event by $\{j, k, \emptyset; k, \emptyset, \emptyset\}$.

We will now derive the distribution of observables in our model. As an example, consider the event $\{j, k, l; j, \emptyset, \emptyset\}$. Such an observation is consistent with three search paths: first, the user may have clicked on j and then decided to stop searching. Second, the user may have clicked on j , continued searching, felt that slot k was not appealing and then decided to stop searching. Finally, the user may have clicked on j , continued searching, felt that slot k was not worth-clicking, still decided to keep searching and finally considered l unappealing as well. As such, the probability of the observable $\{j, k, l; j, \emptyset, \emptyset\}$ is:

$$\begin{aligned} \text{Prob}(\{j, k, l; j, \emptyset, \emptyset\}) &= F_j(1 - \gamma_j) + \\ &F_j\gamma_j(1 - F_k(\{j\}))(1 - \lambda_k) + \\ &F_j\gamma_j(1 - F_k(\{j\}))\lambda_k(1 - F_l(\{j\})). \end{aligned}$$

One can analogously compute that:

$$\begin{aligned} \text{Prob}(\{j, k, \emptyset; \emptyset, \emptyset, \emptyset\}) &= (1 - F_j)(1 - \lambda_j) + \\ &(1 - F_j)\lambda_j(1 - F_k), \\ \text{Prob}(\{j, k, \emptyset; j, \emptyset, \emptyset\}) &= F_j(1 - \gamma_j) + F_j\gamma_j(1 - F_k(\{j\})), \\ \text{Prob}(\{j, k, \emptyset; k, \emptyset, \emptyset\}) &= (1 - F_j)\gamma_j F_k, \\ \text{Prob}(\{j, k, \emptyset; j, k, \emptyset\}) &= F_j\gamma_j F_k(\{j\}), \\ \text{Prob}(\{j, k, l; k, \emptyset, \emptyset\}) &= (1 - F_j)\lambda_j F_k(1 - \gamma_k) + \\ &(1 - F_j)\lambda_j F_k\gamma_k(1 - F_l(\{k\})), \\ \text{Prob}(\{j, k, l; l, \emptyset, \emptyset\}) &= (1 - F_j)\lambda_j(1 - F_k)\lambda_k F_l, \\ \text{Prob}(\{j, k, l; j, k, \emptyset\}) &= F_j\gamma_j F_k(\{j\})(1 - \gamma_k) + \\ &F_j\gamma_j F_k(\{j\})\gamma_k(1 - F_l(\{j, k\})), \\ \text{Prob}(\{j, k, l; j, j, \emptyset\}) &= F_j\gamma_j(1 - F_k(\{j\}))\lambda_k F_l(\{j\}), \\ \text{Prob}(\{j, k, l; k, l, \emptyset\}) &= (1 - F_j)\lambda_j F_k\gamma_k F_l(\{k\}), \\ \text{Prob}(\{j, k, l; j, k, l\}) &= F_j\gamma_j F_k(\{j\})\gamma_k F_l(\{j, k\}). \end{aligned}$$

The above equations fully describe the distribution of observables of our model when the the mainline slots display at least 2 of the 3 advertisers we consider for each keyword. The lemma below proves that our ordered search model is fully identified, that is, it shows that different vectors of parameters are never observationally equivalent. To simplify exposition, let's denote for all j :

$$\mathbf{F}_j \equiv (F_j, F_j(\{k\}), F_j(\{l\}), F_j(\{k, l\})), \quad \mathbf{F} \equiv (\mathbf{F}_j, \mathbf{F}_k, \mathbf{F}_l);$$

and:

$$\theta_j \equiv (\gamma_j, \lambda_j), \quad \theta \equiv (\theta_j, \theta_k, \theta_l).$$

LEMMA 1. *The ordered search model with Sponsored Search lists of size 2 and 3 is identified, that is, for any two vector of parameters (\mathbf{F}, θ) and (\mathbf{F}', θ') , if the above equations take the same values, then $(\mathbf{F}, \theta) = (\mathbf{F}', \theta')$.*

keyword	advertisers	# of obs.
ipod	(A): store.apple.com	8,398
	(B): cellphoneshop.net	
	(C): nextag.com	
diet pill	(A): pricesexposed.net	4,652
	(B): dietpillvalueguide.com	
	(C): certiphene.com	
avg antivirus	(A): Avg-Hq.com	1,336
	(B): avg-for-free.com	
	(C): free-avg-download.com	

Table 1: Keywords and Advertisers

2.2 Data Description

Our data consists of impression and clicking records associated to queries that contained the keywords *ipods*, *diet pills* and *avg antivirus* in Microsoft's Live Search. We chose these keywords because, first, a user that searches for any of them has a well defined objective and, second, because they are highly advertised. Within each of these keywords, we selected the three most popular advertisers (in number of clicks) and considered all impressions in which at least two of these advertisers are displayed.³

For the keyword *ipod*, the Apple Store (www.store.apple.com) is the most important advertiser, followed by the online retailer of electronics Cell Phone Shop (www.cellphoneshop.net) and by the price research website Nextag (www.nextag.com). All the 8398 *ipod* observations in our sample refer to impressions that happened between August 1st and November 1st of 2007.

The most popular advertisers for *diet pills* are, first, the meta-search website Price Exposed (pricesexposed.net), followed by the diet pills retailer *dietpillvalueguide.com* and then by *certiphene.com* (which only sells the diet pill certiphene). All 4,652 impressions considered happened between August 1st and October 1st of 2007.

For *avg antivirus*, the most popular advertiser is the official AVG website, followed by the unofficial distributors of the AVG antivirus *avg-for-free.com* and *free-avg-download.com*. The 1,336 observations range from September 1st to November 1st of 2007. The sample provided by Microsoft AdWords displays impressions associated to different keywords with varying intensities through time. This is why ranges differ across the selected keywords; and we have no reason to expect such differences might affect the estimates of our model.

All keywords possess a leading advertiser that occupies the first position in most of the observations. For *ipod*, the Apple Store occupies the first slot in roughly 77% of the cases, while the Cell Phone Shop appears in 22% of the observations. The situation is reversed when we look at the second slot: the Cell Phone Shop is there in almost 70% of the observations, while the Apple Store and Nextag appear respectively in 17% and 13% of the cases. As table 2 below makes clear, advertising for *diet pills* or *avg antivirus* display a similar pattern.

For all the keywords considered, approximately one out of four impressions got at least one click (25.26% for *ipods*, 24.24% for *diet pills* and 35.55% for *avg antivirus*). As one should expect, click-through rates are decreasing for most of the queries: among the clicks associated to *diet pill*, 56.73% occurred in the first slot, 34.04% in the second and 9.21% in the third. For *ipod*, the concen-

³Regarding the impressions that contain only two of the three selected advertisers in the mainline slots, we only kept those logs which display our selected advertisers in the first two positions. By doing this, we can disregard the advertisers on slot 3 and below without biasing our estimates.

slot	ipod	diet pill
first	(A): 6,460 (76.92%)	(A): 1,912 (41.10%)
	(B): 1,864 (22.20%)	(B): 908 (19.52%)
	(C): 74 (0.88%)	(C): 1,832 (39.38%)
second	(A): 1,438 (17.12%)	(A): 1,848 (39.72%)
	(B): 5,826 (69.37%)	(B): 1,988 (42.73%)
	(C): 1,134 (13.50%)	(C): 816 (17.54%)
third	(A): 26 (0.31%)	(A): 472 (10.15%)
	(B): 22 (0.26%)	(B): 692 (14.88%)
	(C): 950 (11.31%)	(C): 668 (14.36%)
	(other): 7,400 (88.12%)	(other): 2,820 (60.62%)
antivirus		
first	(A): 1,233 (92.29%)	
	(B): 71 (5.31%)	
	(C): 32 (2.40%)	
second	(A): 88 (6.59%)	
	(B): 674 (50.45%)	
	(C): 574 (42.96%)	
third	(A): 9 (0.67%)	
	(B): 21 (1.57%)	
	(C): 355 (26.57%)	
	(other): 951 (71.18%)	

Table 2: Distribution of Advertisers per Slot

slot	ipod	diet pill	antivirus
first	1,572 (74.08%)	640 (56.73%)	205 (43.15%)
second	524 (24.69%)	384 (34.04%)	259 (54.52%)
third	30 (1.41%)	104 (9.21%)	11 (2.31%)
total	2,122 (100%)	1,128 (100%)	475 (100%)

Table 3: Distribution of Clicks per Slot

tration of clicks in the first slot is even higher, as one can see from table 3. The keyword *avg antivirus* is an interesting exception, as most of the clicks happened in the second slot (54.5%).

2.3 Estimation Results

At this stage, it is not possible to tell whether a high click-through rate in the first slot is simply due to users' behavior or is the effect of very high quality advertisers. In the same vein, what explains the very low click-through rate in the third slot for *ipod*? Is it because advertisers are bad matches for the users' search or is it the result of search externalities imposed by the links in the first two slots?

In order to evaluate externalities, we must estimate the parameters of our model. We do this with the well-established *maximum likelihood method*, which selects values for the parameters that maximize the probability of the sample. First we must derive an expression, called the *log-likelihood*, for the (log of) probability of the sample given the parameters of the model.⁴ Our log-likelihood function is:

$$\log L = \sum_n \log [\text{Prob}(\{j_n, k_n, l_n; c_n^1, c_n^2, c_n^3\})],$$

where the probability of observations $\{j_n, k_n, l_n; c_n^1, c_n^2, c_n^3\}$ is given by the equations in Section 2.1.

Next we estimate the parameters to be those that maximize the log-likelihood. Before discussing our estimation results, we need

⁴It is common to use the log of the probability as opposed to the probability itself to simplify the algebra. As log is a monotone function, maximizing the log-likelihood corresponds to maximizing the likelihood.

to make one important observation. The conditional click-through rate of some advertiser j , $F_j(\{k\})$, is the probability that a random user clicks on ad j given that this user clicked on advertiser k 's link and kept searching until he read j 's link. Note that $F_j(\{k\})$ abstracts from position externalities, as this is the probability that a user that *read* the ad gives a click on it. We have three reasons to think that conditional click-through rates should differ from baseline click-through rates. First, link k may offer low prices for ipods, in which case, even if the user keeps browsing the sponsored list after clicking on k (an event of probability γ_k), he will be less likely to click on j or on any other link. This is the *negative externality* effect, which pushes, let's say $F_j(\{k\})$, to be less than F_j . Second, link k may increase the users' willingness to click on j , which may happen if, for example, link k is a meta-search website. In this case, $F_j(\{k\})$ is greater than F_j , which corresponds to a *positive externality* effect.

These first two reasons for F_j to depart from $F_j(\{k\})$ relate to information externalities. There is a third reason, though, not related to externalities but to the structure of our data, that may explain why $F_j \neq F_j(\{k\})$: the group of users that make at least one click may be fundamentally different from the total pool of users that perform searches on Microsoft Live. As such, the conditional click-through rate $F_j(\{k\})$ reflects the probability of j getting a click among a quite selected group of users. It is natural to think that these users click more often on sponsored links than a common user; and this should push $F_j(\{k\})$ to be higher than F_j . We call this the *selection* effect.

As a consequence, we can safely interpret estimates such that $F_j > F_j(\{k\})$ as evidence that advertiser k imposes a negative externality on advertiser j . Nevertheless, if $F_j < F_j(\{k\})$, as we don't observe any users' characteristics, we can't tell apart positive externalities from purely selection effect. We need to keep this in mind in order to interpret the estimation results.

One can directly test whether the selection effect is driving our estimates by looking at the continuation probabilities λ_j and γ_j . Clearly, absent any selection effect and granted j is not a meta-search website, λ_j , the probability that a user keeps browsing after *not* clicking on j , is expected to be higher than γ_j , the probability that a user keeps browsing after clicking on j . The reason for this is that users may only fulfill their search needs if they do click on j , in which case they are not expected to return to the results page. As a consequence, having λ_j significantly lower than γ_j is strong evidence in favor of the selection effect, as the subgroup of users that indeed make clicks is much more likely to patronize sponsored search.

We are now able to discuss our estimation results, which are displayed at Table 3. We find that for the three search terms we investigated, selection effects were ubiquitous. Nonetheless, we observed significant negative externalities in two of them (*ipod* and *avg antivirus*). For the third keyword (*diet pills*), we observed that conditional click-through-rates were higher than the base-line click-through-rates, although it is not possible to determine whether to attribute this to the selection effect or to positive externalities. In the following subsections, we discuss the results for each keyword in detail.

2.3.1 ipod Results

For this keyword, the lead advertiser (the Apple Store) has a very high click-through rate: 21%. Its competitors, the Cell Phone Shop and Nextag, have 8.7% and 10.4%, respectively. These estimates can be interpreted as the probability that the first slot gets a click when it is occupied by one of these three advertisers. The difference between the Apple Store click-through rate and that of

its competitors is significant at the 1% level. As such, in the *ipod* case, the lead advertiser (who occupies the top position in 76% of the observations – see Table 2) is also the most effective in attracting clicks.

Our estimates detect that Apple Store imposes a negative externality on the Cell Phone Shop (as $F_B = 0.08 > 0.04 = F_B(\{A\})$, and the difference is significant at 5%) and on Nextag (as $F_C = 0.10 > 0.04 = F_C(\{A\})$, and the difference is significant at 5%). This means that the information provided by the Apple Store website reduced by half the appeal to a random user of the links to the Cell Phone Shop or the Nextag. The lack of observations in which users click on Nextag and then click on Apple Store or the Cell Phone Shop prevents us from being able to estimate $\gamma_C, F_A(\{C\})$, $F_A(\{B, C\})$, $F_B(\{C\})$ and $F_B(\{A, C\})$.

The selection effect indeed seems to play a role in our estimates. Looking at the *ipod* results, one can see that γ 's are higher than λ 's for at least two advertisers: for the Apple Store, $\gamma_A = 0.94 > 0.76 = \lambda_A$ (although the difference is not significant) and for the Cell Phone Shop, $\gamma_B = 1 > 0.62 = \lambda_B$ (significant at 15%). This suggests that users that make one click in a sponsored link are more likely to keep browsing the sponsored list. As the results presented above point out, though, the selection effect wasn't strong enough to shadow the negative externalities that the Apple Store imposes on its competitors.

2.3.2 diet pill Results

Alike the *ipod* case, the leading advertiser for *diet pill* is also the most effective in terms of attracting users: the click-through rate of *pricesexposed.net*, roughly 21%, is significantly (at 1% level) higher than that of its competitors (15% for *dietpillvalueguide.com* and 5% for *certiphene.com*).

We didn't find evidence of negative information externalities among *diet pill* advertisers. For *pricesexposed.net*, the click-through rate jumps from roughly 21% to 31% if *certiphene.com* was previously clicked; and the difference is significant at 10%. The same happens with *dietpillvalueguide.com*: its click-through rate goes from 15% to either 66% (in case *certiphene.com* got a click) or to 33% (in case *certiphene.com* and *pricesexposed.net* had clicks); and both differences are significant at 5%.

Interestingly, the click-through rate of *certiphene.com* jumps from 5% to 8% (difference significant at 5%) if *dietpillvalueguide.com* was previously clicked by the user. Since *dietpillvalueguide.com* is a website specialized in comparing diet products, one can think that positive reviews of the Certiphene pills might explain this difference (positive externality).

As discussed above, we cannot rule out that the selection effect explains this difference, though. Indeed, our estimates imply that users are more likely to keep browsing the sponsored links if they clicked on *certiphene.com*: $\gamma_C = 1 > 0.57 = \lambda_C$ (significant at 5%).

2.3.3 avg antivirus Results

Unlike the previous keywords, the leading advertiser for *avg antivirus* is not the one with highest CTR. In fact, *Avg-Hq.com* has the lowest CTR (15%), while *avg-for-free.com* and *free-avg-download.com* have a 20% and 21% CTRs, respectively (higher than *Avg-Hq.com*'s CTR at a 15% confidence level).

Our estimates detect that *avg-for-free.com* imposes a negative externality on *Avg-Hq.com*, as $F_A = 0.15 > 0 = F_A(\{B\})$ (significant at 5%). As in the *ipod* case, the lack of observations in which users click on *free-avg-download.com* and then click on *Avg-Hq.com* or *avg-for-free.com* makes it impossible to estimate $\gamma_C, F_A(\{C\})$, $F_A(\{B, C\})$, $F_B(\{C\})$ and $F_B(\{A, C\})$.

keyword	ipod	diet pill	antivirus
F_A	0.210 (0.005)	0.210 (0.008)	0.151 (0.010)
$F_A(\{B\})$	0.250 (0.038)	0.232 (0.032)	0.00 (0.074)
$F_A(\{C\})$	—	0.317 (0.065)	—
$F_A(\{B, C\})$	—	0.664 (0.075)	—
F_B	0.087 (0.006)	0.150 (0.009)	0.206 (0.038)
$F_B(\{A\})$	0.030 (0.022)	0.146 (0.034)	0.364 (0.050)
$F_B(\{C\})$	—	0.663 (0.080)	—
$F_B(\{A, C\})$	—	0.334 (0.083)	—
F_C	0.104 (0.012)	0.051 (0.004)	0.215 (0.042)
$F_C(\{A\})$	0.040 (0.032)	0.052 (0.017)	0.242 (0.042)
$F_C(\{B\})$	0.095 (0.032)	0.088 (0.029)	0.121 (0.889)
$F_C(\{A, B\})$	0.327 (0.190)	0.664 (0.089)	0.125 (0.699)
λ_A	0.676 (0.056)	0.760 (0.064)	1.0 (0.217)
λ_B	0.627 (0.042)	0.673 (0.057)	0.183 (0.049)
λ_C	1.00 (0.057)	0.579 (0.037)	0.424 (0.201)
γ_A	1.00 (0.777)	0.940 (0.195)	1.00 (0.231)
γ_B	1.00 (0.820)	1.00 (0.743)	0.686 (0.902)
γ_C	—	1.00 (0.892)	—

Table 4: Estimates of the Ordered Search Model

keyword	ipod	diet pill	antivirus
f_A	0.216 (0.005)	0.205 (0.008)	0.144 (0.009)
f_B	0.085 (0.004)	0.164 (0.009)	0.256 (0.036)
f_C	0.107 (0.011)	0.057 (0.004)	0.253 (0.038)
s^1	1.00 —	1.00 —	1.00 —
s^2	0.676 (0.036)	0.671 (0.037)	0.961 (0.136)
s^3	0.400 (0.072)	0.699 (0.056)	0.144 (0.043)

Table 5: Estimates of the Separable CTR Model

2.3.4 Addressing Endogeneity

One might question the consistency of our estimates by arguing that the variation on slot allocations may be endogeneous, that is, advertisers may change their bids (to alter their positions) as a response to different groups of users (that browse the web in different time periods). We believe on the contrary that a significant source of variation is due to the allocation procedure itself. Microsoft Ad-Center applies a randomization procedure that perturbs submitted bids and (non-deterministically) changes the slot allocations. This makes the variation exogenous.

2.4 Model Validation

It's commonly assumed, by both the search engines and the literature on keyword auctions, that the users' clicking behavior follows a separable click-through rate model. According to this simple model, the probability that advertiser j gets a click when she occupies slot k is given by:

$$\text{Prob}(\text{advertiser } j \text{ gets a click on slot } k) = s^k \cdot f_j,$$

where s^k is the slot-specific click-through rate and f_j is the advertiser-specific click-through rate. It then follows that, according to this model, the probability of any clicking log $\{j, k, l; c^1, c^2, c^3\}$ is given by:

$$\text{Prob}(\{j, k, l; c^1, c^2, c^3\}) = (s^1 f_j)^{I_j} (1 - s^1 f_j)^{1-I_j} (s^2 f_k)^{I_k} \cdot (1 - s^2 f_k)^{1-I_k} (s^3 f_l)^{I_l} (1 - s^3 f_l)^{1-I_l},$$

where I_j is an indicator function that equals 1 if and only if link j was clicked.

We applied the separable click-through rate model to the same data base used to estimate the ordered search model. After normalizing $s^1 = 1$ (which is necessary for identification), we obtained the results shown in Table 5. Note that, as one would expect, the estimates of advertiser-specific click-through rates in the separable model, f_A , f_B and f_C , are roughly equal to the baseline click-through rates of the ordered search model.

This subsection is devoted to comparing the ordered search model (to simplify, ordered model) from this paper with the widely used separable click-through rate model (separable model from now on). Although having less parameters, the separable model cannot be written as a constrained version of the ordered model. This is because the separable model specifies a slot-specific term s^k into click-through rates that does not depend on who occupies the previous slots. As such, unlike the ordered model, the separable model disregards externalities between advertisers, but, on the other hand, allows for slot-specific effects not contemplated by the ordered model.

To compare these models, we contrasted the probabilities of clicking logs implied by each model to raw moments computed from our data base. The ordered model has more parameters than the separable model, and so it should come as no surprise that it provides a better fit for the data. This is a drawback from purely non-parametrical comparisons; and in future work we shall address these concerns by applying non-nested hypotheses testing techniques (as in [17]).

Table 6 below present our results. To save space, the event 'user clicks on advertiser J occupying the first slot' is denoted simply by J and the event 'user clicks on advertiser J occupying the second slot given that he also clicked on advertiser K occupying the first slot' is denoted by $J|K$. Finally, the event 'user clicks on advertiser J occupying the second slot given that he did not click on advertiser K occupying the first slot' is denoted by $J|\sim K$. The column labeled 'realized CTR' gives the empirical distribution of

these events, while the columns 'ordered' and 'separable' present the probabilities predicted by the ordered and the separable models, respectively.

Let's analyze table 6, whose left side is constructed based on clicking logs associated to the keyword *ipod*. Among the events considered, the ordered model predictions are at least as good as the separable models predictions (as measured by their deviations to the empirical distribution) in 14 out of 15 events (as indicated by \star). More importantly, the only event in which the separable model performs better than the ordered model ($B|\sim C$) may be found in less than 1% of the sample (as C occupies the first slot in only 0.88% of the observations - see table 2).

The same pattern is observed for the keywords *diet pills* and *avg antivirus*. As table 6 shows, the ordered model provides better predictions for 14 out of the 15 events considered. Further, the cases in which the separable model performs better ($C|\sim B$ and $B|\sim C$, respectively) occur in less than 2% of the observations for both keywords.

The 15 events considered span all possible clicking histories for the first 2 slots. As a consequence, by applying Bayes rule to the predictions displayed on table 6 and 7, one can derive the implied probabilities for any clicking log involving only slots 1 and 2. Thus, from the discussion above, we can safely conclude that the predictions of the ordered model outperform those of the separable model for all keywords considered.

3. EQUILIBRIUM ANALYSIS

We'll now analyze how advertisers bid given that users do ordered search. We return to a model with a set of N advertisers denoted by A_j , $j \in \{1, \dots, n\}$ and K . Each advertiser A_j has a value of v_{A_j} per click. Search engines use the following generalization of the second-price auction to sell sponsored links: first, each advertiser A_j submits a single bid b_{A_j} representing his willingness to pay per click. Then each advertiser's bid is multiplied by a weight w_{A_j} that solely depends on his characteristics, producing a score $s_{A_j} = w_{A_j} \cdot b_{A_j}$. Next, advertisers are ranked in decreasing order of their scores and the j^{th} highest ranked advertiser gets the j^{th} highest slot. When an advertiser receives a click, he is charged a price equal to the smallest bid he could have submitted that would have allowed him to maintain his position in the sponsored list. Labeling advertisers such that A_i denotes the advertiser ranked in the i^{th} slot, we see that advertiser A_j pays p_{A_j} where:

$$p_{A_j} \cdot w_{A_j} = b_{A_{j+1}} \cdot w_{j+1} \quad \text{which gives} \quad p_{A_j} = \frac{b_{A_{j+1}} \cdot w_{A_{j+1}}}{w_{A_j}}.$$

The total payment of advertiser A_j is then $p_{A_j} \cdot q^j$, where q^j is the total number of clicks of slot j .⁵ To simplify the analysis, we'll take the ordered search model of the previous section and assume that baseline and conditional click-through rates are the same for each advertiser, that is, we assume that $F_{A_j} = F_{A_j}(H)$ for any click history H . Although our empirical exercise suggests that baseline and conditional click-through rates indeed differ, this assumption is necessary to bring tractability to our theoretical model of bidding. Further, as we will argue later, our main theoretical conclusions remain valid under the more general ordered search model of the previous section.

With this assumption in hand, the total number of clicks of the

⁵Note q^j is actually a function of the entire assignment of advertisers to slots preceding j ; we denote it by q^j to simplify notation.

Prob.	<i>ipod</i>			<i>diet pills</i>			<i>avg antivirus</i>		
	realized CTR	ordered	separable	realized CTR	ordered	separable	realized CTR	ordered	separable
A	0.21	0.21 [★]	0.22	0.23	0.21 [★]	0.21	0.15	0.15 [★]	0.14
B	0.09	0.09 [★]	0.09	0.15	0.15 [★]	0.16	0.17	0.21 [★]	0.26
C	0.05	0.10 [★]	0.11	0.03	0.05 [★]	0.06	0.22	0.22 [★]	0.25
$A B$	0.26	0.25 [★]	0.15	0.22	0.23 [★]	0.14	0.00	0.00 [★]	0.14
$A \sim B$	0.13	0.13 [★]	0.15	0.10	0.14 [★]	0.14	0.04	0.03 [★]	0.14
$A C$	0.00	0.00 [★]	0.15	0.36	0.32 [★]	0.14	0.00	0.00 [★]	0.14
$A \sim C$	0.33	0.21 [★]	0.15	0.11	0.12 [★]	0.14	0.06	0.06 [★]	0.14
$B A$	0.04	0.03 [★]	0.06	0.21	0.14 [★]	0.11	0.04	0.03 [★]	0.06
$B \sim A$	0.06	0.06 [★]	0.06	0.10	0.11 [★]	0.11	0.06	0.06 [★]	0.06
$B C$	0.00	0.00 [★]	0.06	0.60	0.66 [★]	0.11	0.00	0.00 [★]	0.06
$B \sim C$	0.05	0.09	0.06 [★]	0.10	0.09 [★]	0.11	0.05	0.09	0.06 [★]
$C A$	0.05	0.04 [★]	0.07	0.07	0.05 [★]	0.04	0.05	0.04 [★]	0.07
$C \sim A$	0.07	0.07 [★]	0.07	0.06	0.04 [★]	0.04	0.07	0.07 [★]	0.07
$C B$	0.19	0.09 [★]	0.07	0.11	0.09 [★]	0.04	0.19	0.09 [★]	0.07
$C \sim B$	0.08	0.07 [★]	0.07	0.07	0.03	0.04 [★]	0.08	0.07 [★]	0.07

Table 6: Model Validation

j^{th} slot is given by:

$$q^j = F_{A_j} \cdot \prod_{k=1}^{j-1} c_{A_k}, \quad \text{where } c_{A_k} = F_{A_k} \gamma_{A_k} + (1 - F_{A_k}) \lambda_{A_k}.$$

Each term c_{A_k} accounts for the fraction of users that continue browsing the sponsored list after coming across advertiser A_k . As such, the total number of clicks of slot j is the product of advertiser A_j 's click-through rate (F_{A_j}) and the total number of users that reach that position ($\prod_{k=1}^{j-1} c_{A_k}$). Advertiser A_j 's payoff is then $(v_{A_j} - p_{A_j})q^j$.

We are interested in analyzing the complete information Nash equilibria and resulting efficiency of various scoring rules. A *complete information Nash equilibrium* is a vector of bids such that no advertiser can unilaterally change his bid and improve his payoff. The *efficiency* of an equilibrium is simply the sum of all advertisers' value per click times total number of clicks. The *optimum social welfare* is the assignment of advertisers to slots with maximum efficiency. Given our labeling scheme in which the j^{th} slot is occupied by advertiser $A_j \in N$, the optimum social welfare can be written as:

$$W(N) = \max_{A_1, \dots, A_N \in N} \sum_{j=1}^N q^j v_{A_j} \quad (1)$$

3.1 Can Scoring Rules Help?

Search engines have often changed their auction rules for keyword advertising in order to increase revenue. Yahoo! first dropped a generalized first-price auction and adopted the rank-by-bid GSP in early 1997. Ten years later, and with a much wider base of advertisers, Yahoo! opted for a less drastic change and simply altered its scoring rule from rank-by-bid to rank-by revenue (in which case $w_{A_j} = F_{A_j}$). Microsoft's Live Search followed the same path and also in 2007 moved from the rank-by-bid to the rank-by-revenue GSP. Recently, Google also changed its scoring rule, although its precise functional form was not made public.

Search engines are very reluctant to make bold changes in their auction rules for mostly two reasons: first, advertisers are hardly willing to learn a completely new auction format; and may switch to a competitor if that happens. Second, it is believed that much may be achieved in terms of revenue and efficiency by simply ex-

ploring different scoring rules within the GSP format. In this subsection, we make this claim formally by studying how the choice of a scoring rule affects the set of complete information Nash equilibrium of the GSP.

We will focus on a very interesting, but so far neglected, equilibrium of the GSP: the one that maximizes the search engine's revenue among all pure strategy Nash equilibria. The next lemma derives the bid profile that maximizes revenue for the search engine:

LEMMA 2. *Consider the GSP with scoring rule w_{A_j} , selling K slots to $N > K$ advertisers. Let advertisers A_1, \dots, A_N be the efficient assignees of slots 1 to N and assume advertisers submit bids according to:*

$$b_{A_j} = (1 - c_{A_j}) \frac{w_{A_{j-1}}}{w_{A_j}} v_{A_{j-1}} + c_{A_j} \frac{w_{A_{j+1}}}{w_{A_j}} b_{A_{j+1}}$$

$$\text{for } j \in \{2, \dots, K\}, b_{A_{K+1}} = \frac{w_{A_K} v_{A_K}}{w_{A_{K+1}}}, b_{A_1} > b_{A_2} \quad (2)$$

$$\text{and } b_{A_j} < b_{A_{K+1}} \text{ for } j > K + 1. \quad (3)$$

If this bid profile constitutes a Nash equilibrium, then it maximizes the search engine's revenue among all pure strategy complete information Nash equilibria. We call it the greedy bid profile.

Proof. Consider the efficient allocation, that is, let advertisers A_1, \dots, A_N receive slots 1 to N in this order. The Nash equilibrium candidate that extracts most rents from advertisers has clearly two properties: first, the last advertiser to obtain a slot (who is A_K) enjoys a zero payoff. This implies that his payment per click, $\frac{w_{A_{K+1}}}{w_{A_K}} b_{A_{K+1}}$, has to be equal to his value per click, v_{A_K} , what gives equation (3). Second, all advertisers above A_K should be indifferent between following equilibrium strategies and undercutting the advertiser immediately below them. To see why this is true, imagine some advertiser A_j strictly prefers slot j to slot $j + 1$ (given this bid profile). In this case, if advertiser A_{j+1} slightly increases his bid, A_j has to pay more but still finds all deviations unprofitable (and the search engine's total revenue is higher). As a consequence, $b_{A_{j+1}}$ has to satisfy:

$$\left(\prod_{k=1}^{j-1} c_{A_k} \right) F_{A_j} \cdot \left(v_{A_j} - \frac{w_{A_{j+1}}}{w_{A_j}} b_{A_{j+1}} \right)$$

$$= \left(\prod_{k=1}^{j-1} c_{A_k} \right) c_{A_{j+1}} F_{A_j} \cdot \left(v_{A_j} - \frac{w_{A_{j+2}} b_{A_{j+2}}}{w_{A_j}} \right),$$

what, after switching indexes, gives equation (2). ■

As the next proposition shows, such a bid profile is an equilibrium for all $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$ if and only if weights are given by:

$$w_{A_j} = \frac{F_{A_j}}{1 - c_{A_j}} = \frac{F_{A_j}}{1 - (F_{A_j} \gamma_{A_j} + (1 - F_{A_j}) \lambda_{A_j})}.$$

Although at first awkward, the scoring rule above is a quite natural one. Indeed, as first proved by [11], advertiser j comes on top of advertiser k in the efficient allocation if and only if $v_{A_j} \cdot w_{A_j} \geq v_k \cdot w_{A_k}$.

PROPOSITION 1. *Consider the GSP with scoring rule w_{A_j} , selling K slots to $N > K$ advertisers. The greedy bid profile constitutes a complete information Nash equilibrium for all valuations and search parameters $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$ if and only if $w_{A_j} = \frac{F_{A_j}}{1 - c_{A_j}}$ (up to a multiplicative constant). In this case, the equilibrium allocation is efficient and the search engines's revenue is maximal.*

Proof. Let advertisers bid according to (2) and (3). By construction, no advertiser wants to undercut someone else's bid and get a slot below his own. Further, no bidders want to deviate upwards. To see why, let's first assume (for later confirmation) that $w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot b_{A_{j+1}}$ for all j . To get a contradiction, imagine some advertiser A_{j+1} strictly prefers slot j to slot $j + 1$ (under this bid profile). In this case:

$$\begin{aligned} & \left(\prod_{k=1}^{j-1} c_{A_k} \right) F_{A_{j+1}} \cdot \left(v_{A_{j+1}} - \frac{w_{A_j} b_{A_j}}{w_{A_{j+1}}} \right) \\ & > \left(\prod_{k=1}^{j-1} c_{A_k} \right) c_{A_j} F_{A_{j+1}} \cdot \left(v_{A_{j+1}} - \frac{w_{A_{j+2}} b_{A_{j+2}}}{w_{A_{j+1}}} \right), \end{aligned}$$

what simplifies to:

$$w_{A_{j+1}} v_{A_{j+1}} (1 - c_{A_j}) > w_{A_j} b_{A_j} - c_{A_j} w_{A_{j+2}} b_{A_{j+2}}. \quad (4)$$

Since $w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot b_{A_{j+1}}$ for all j , we have that:

$$w_{A_j} b_{A_j} - c_{A_j} w_{A_{j+2}} b_{A_{j+2}} \geq w_{A_j} b_{A_j} - c_{A_j} w_{A_{j+1}} b_{A_{j+1}}. \quad (5)$$

By the choice of the scoring rule, and the fact the allocation is efficient, we know that $w_{A_j} \cdot v_{A_j} \geq w_{A_{j+1}} \cdot v_{A_{j+1}}$. Thus:

$$w_{A_j} v_{A_j} (1 - c_{A_j}) \geq w_{A_{j+1}} v_{A_{j+1}} (1 - c_{A_j}). \quad (6)$$

Plugging (5) and (6) into (4), we obtain that:

$$w_{A_j} v_{A_j} (1 - c_{A_j}) > w_{A_j} b_{A_j} - c_{A_j} w_{A_{j+1}} b_{A_{j+1}}.$$

Using the definition of b_{A_j} from equation (2), the inequality above becomes:

$$w_{A_j} v_{A_j} > w_{A_{j-1}} v_{A_{j-1}},$$

contradicting efficiency. We conclude that if $w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot b_{A_{j+1}}$ holds for all j , then the greedy bid profile is a Nash equilibrium.

It only remains to be shown that the bids described by (2) and (3) are indeed such that $w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot b_{A_{j+1}}$ for all A_j . The proof is by induction. First, it is a matter of algebra to see that $w_{A_{j+1}} \cdot b_{A_{j+1}} \geq w_{A_{j+2}} \cdot b_{A_{j+2}}$ implies $w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot$

$b_{A_{j+1}}$. Second, this induction hypothesis is true for advertiser A_K , as $w_{A_K} \cdot b_{A_K} \geq w_{A_{K+1}} \cdot b_{A_{K+1}}$ if and only if:

$$\begin{aligned} (1 - c_{A_K}) w_{A_{K-1}} v_{A_{K-1}} + w_{A_{K+1}} c_{A_K} b_{A_{K+1}} & \geq w_{A_{K+1}} \cdot b_{A_{K+1}} \\ \Leftrightarrow (1 - c_{A_K}) w_{A_{K-1}} v_{A_{K-1}} & \geq (1 - c_{A_K}) w_{A_K} v_{A_K} \\ \Leftrightarrow w_{A_{K-1}} v_{A_{K-1}} & \geq w_{A_K} v_{A_K}. \end{aligned}$$

By [11], the last inequality is true for all $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$ if and only if $w_{A_K} = \frac{F_{A_K}}{1 - c_{A_K}}$ (up to a multiplicative constant). ■

Our next proposition brings a pessimistic message about what scoring can achieve in the GSP. It shows that there is no scoring rule for which an efficient equilibrium where each advertiser pays his Vickrey-Clark-Groves payments exists for all profiles of valuations and search parameters. This extends a result by [8], who shows that the GSP equipped with the "rank-by-revenue" scoring function ($w_{A_K} = F_{A_K}$) does not possess an efficient equilibrium that implements VCG payments.

Recall the Vickrey-Clark-Groves (VCG from now on) payments charge each advertiser the welfare difference imposed on the others:

$$p_{A_j}^V = W(N - \{A_j\}) - (W(N) - q^j v_{A_j})$$

where $W(N)$ is the welfare as defined by equation (1).

PROPOSITION 2. *Consider the GSP selling K slots to $N > K$ advertisers. There is no scoring rule w_{A_j} which depends solely on advertiser A_j 's search parameters $(F_{A_j}, \gamma_{A_j}, \lambda_{A_j})$ that implements an efficient equilibrium with VCG payments for all valuations and search parameters $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$.*

Proof. By the payment rule of the GSP, bids that implement VCG payments have to be such that:

$$q^{j-1} \cdot \frac{w_{A_j} b_{A_j}}{w_{A_{j-1}}} = p_{A_{j-1}}^V.$$

This implies that advertisers have to bid according to:

$$b_{A_j} = \frac{w_{A_{j-1}}}{w_{A_j}} \cdot \frac{p_{A_{j-1}}^V}{q^{j-1}} \quad \text{for } j \in \{2, \dots, K+1\}, \quad (7)$$

$$b_{A_1} > b_{A_2} \quad \text{and} \quad b_{A_j} < b_{A_{K+1}} \quad \text{for } j > K+1. \quad (8)$$

With these bids in hand, we have to pick a scoring rule w_{A_j} such that the order of scores corresponds to the efficient ranking of advertisers, that is,

$$w_{A_j} \cdot b_{A_j} \geq w_{A_{j+1}} \cdot b_{A_{j+1}} \quad (9)$$

if and only if A_j is assigned a slot above A_{j+1} in the efficient allocation.

The idea of the proof is to show that there is no scoring rule depending only on $F_{A_j}, \gamma_{A_j}, \lambda_{A_j}$ that preserves the inequality (9) for all profiles $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$.

To see why this is true, let's plug the bids (7) in the inequality (9) to obtain that any scoring rule that implements a VCG equilibrium has to satisfy:

$$\frac{w_{A_{j-1}}}{F_{A_{j-1}}} p_{A_{j-1}}^V \geq \frac{w_{A_j}}{c_{A_{j-1}} F_{A_j}} p_{A_j}^V.$$

To make the argument, let's take a profile of primitives

$$\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$$

such that the efficient assignees of slots A_j and A_{j-1} have identical click-through rates and continuation probabilities, that is, let $(F_{A_j}, \gamma_{A_j}, \lambda_{A_j}) = (F_{A_{j-1}}, \gamma_{A_{j-1}}, \lambda_{A_{j-1}})$. Since the scoring

rule is anonymous and can only depend on this quantities, it follows that $w_{A_j} = w_{A_{j-1}}$. Condition (10) then becomes:

$$p_{A_{j-1}}^V \geq \frac{1}{c_{A_{j-1}}} p_{A_j}^V.$$

The VCG payments $p_{A_{j-1}}^V$ and $p_{A_j}^V$ are clearly bounded by the total welfare obtained by advertisers other than A_j and A_{j-1} . As a consequence, one can pick continuation probabilities $\gamma_{A_{j-1}}, \lambda_{A_{j-1}}$ small enough (and consequently $c_{A_{j-1}}$ small enough) to violate the inequality above. The strategy of this proof generalizes and closely follows the argument given by [8] for why the rank-by-revenue GSP may not possess a VCG equilibrium. ■

4. CONCLUSION

This work contributes in two fronts: in the empirical side, we document information and position externalities among sponsored search advertisers. Our results bring suggestive evidence that part of the population of users perform price research through the sponsored list (as a user that clicks in a sponsored link is more likely to keep browsing the sponsored list than users that don't make clicks at all). Finally, our empirical model of search behavior is shown to have more predictive power than the widely popular separable click-through rate model.

On the theoretical side, we take click-through rates as produced by ordered search and study the GSP equilibrium properties under different scoring rules. We derive the unique scoring rule that implements the revenue-maximizing complete information Nash equilibrium of the GSP (under any scoring rule). We temper this positive claim with a strong inexistence result: extending the analysis of [8], we show that no scoring rule implements an efficient equilibrium with VCG payments for all profiles of valuations and search parameters.

These results fundamentally rely on the assumption that users browse from the top to the bottom of the sponsored list and take clicking decisions link by link (what we call ordered search). It would be interesting to extend the analysis (both empirical and theoretical) to allow users to (optimally) perform other search procedures.

5. REFERENCES

- [1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pal. Sponsored search auctions with markovian users. In *Workshop on Internet and Network Economics (WINE)*, 2008.
- [2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proc. the Seventh ACM Conference on Electronic Commerce*, Ann Arbor, Michigan, USA, 2006.
- [3] S. Athey and G. Ellison. Position auctions with consumer search. working paper, 2008.
- [4] M. Baye, J. Morgan, and P. Scholten". Information, search and price dispersion. In T. Hendershott, editor, *Handbooks in Information Systems*, volume 1 of *Handbooks in Information Systems*. Elsevier, 2006.
- [5] N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In *WSDM'08*, pages 87–94, 2008.
- [6] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth in keywords. *American Economic Review*, 97(1):242–259, 2007.
- [7] A. Ghosh and M. Mahdian. Externalities in online advertising. In *International World Wide Web Conference (WWW)*, 2008.
- [8] I. Giotis and A. R. Karlin. On the equilibria and efficiency of the gsp mechanism in keyword auctions with externalities. In *WINE'08*, 2008.
- [9] P. Jeziorskiy and I. Segal. What makes them click: Empirical analysis of consumer demand for search advertising. working paper, 2009.
- [10] T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately predicting clickthrough data as implicit feedback. In *SIGIR*, pages 154–161, 2005.
- [11] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *Workshop on Internet and Network Economics (WINE)*, Shanghai, China, 2008.
- [12] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proc. the Seventh ACM Conference on Electronic Commerce*, Ann Arbor, Michigan, USA, 2006.
- [13] S. Lahaie and D. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proc. the Eighth ACM Conference on Electronic Commerce*, San Diego, California, USA, 2007.
- [14] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. 28: Sponsored search auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [15] D. T. Mortensen. Job search and labor market analysis. In O. Ashenfelter and R. Layard, editors, *Handbook of Labor Economics*.
- [16] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.
- [17] Q. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57(2):307–333, 1989.