Advertiser prominence effects in search advertising*

Przemysław Jeziorski
University of California, Berkeley

Sridhar Moorthy
University of Toronto

September 23, 2016

*We are grateful to Microsoft Research for providing us the data which forms the basis of this work. Many thanks also to the *Management Science* editorial board—Matthew Shum, the Associate Editor, and three anonymous referees—for expert advice on how to revise the paper. Avi Goldfarb, Tanjim Hossain, Kinshuk Jerath, Sridhar Narayanan, Matthew Osborne, Ilya Segal, Miguel Villas-Boas, Yi Zhu, and audiences at the 2014 Theory + Practice in Marketing Conference and the 2014 QME Conference offered several helpful comments and suggestions. This research was supported by Grant #s 435–2013–0704 and 864–2007–0306 from the Social Sciences and Humanities Research Council of Canada to the second author.

Abstract

Search advertising is the ordered list of advertisements that appears when a user searches for something in an online search engine. By construction, these ads differ in prominence: ads higher up the list are more prominent than ads lower down the list. However, search ads also differ in prominence in another way: prominence of advertiser. This paper examines how these two types of prominence interact in determining the click-through-rate of these ads. Using individual-level click-stream data from Microsoft's Live Search platform, and measures of advertiser prominence from Alexa.com, we find that ad position and advertiser prominence are substitutes. Specifically, in searches for camera brands, a retailer not in the Top-100 of Alexa rankings has a 30–50% higher click-through-rate (CTR) in position one than in position two, whereas a retailer in the Top-100 of Alexa rankings has only a 0–13% higher CTR for the same position improvement. Qualitatively similar results are obtained for several other search strings. These findings demonstrate, first, that advertiser brand matters even for search ads, and, second, the way it matters, is the opposite of what is usually assumed in the theoretical literature on search advertising.

1 Introduction

Paid search advertising is the advertising that occurs when a consumer searches for something in an online search engine such as Bing or Google. The search ads, also called sponsored links, appear at the top and right-hand side of the results page in an ordered list (as in Figure 1). Unlike traditional advertising, search advertising space is sold on a pay-for-performance basis: advertisers pay the search engine for clicks, not for exposure. Ad positions are usually auctioned off in a generalized second-price auction (Aggarwal et al. 2006, Edelman et al. 2007, Varian 2007). This generally leads to a higher cost-per-click (CPC) for higher ad positions. However, since higher ad positions also yield higher click-through-rates (CTR)—the so-called "position effect"(Brooks 2006, Animesh et al. 2011, Rutz et al. 2012, Narayanan and Kalyanam 2015)[1]—advertisers face a tradeoff in deciding whether to bid more in the quest for higher positions or bid less in order to increase margin per click. Ad platforms face a mechanism design problem: how to assign advertisers to advertising slots and how to price those slots in order to maximize revenue. Both of these decisions depend heavily on the nature of the position effect for different advertisers.

The theoretical literature on search advertising, focusing on the optimal mechanism design aspects of the generalized second-price auction (Edelman et al. 2007, Varian 2007), assumes that CTR decreases from top to bottom, independent of advertiser identity. More common is the assumption that CTR depends on ad position and advertiser identity, multiplicatively. That is, $\mathrm{CTR}_{jk} = \alpha_j \beta_k$, where α_j is an ad position-specific factor and β_k is an advertiser-specific factor (Aggarwal et al. 2006, Katona and Sarvary 2010, Pin and Key 2011, Nekipelov 2014). Empirical studies, however, suggest that CTR may have a more complicated relationship with ad position and advertiser identity. For instance, Narayanan and Kalyanam (2015) show that higher ad positions lead to smaller CTR-improvements for bigger advertisers, and for a given advertiser, when they are assigned higher "quality scores" by the search engine. Disaggregate analyses reveal even more nuanced patterns. Jeziorski

[1] What drives the position effect is the subject of a separate literature. Perhaps ads in higher positions are more likely to be noticed than ads in lower positions (Richardson et al. 2007, Buscher et al. 2010, Arkhangelsky et al. 2013, Ursu 2015), or perhaps consumers read meaning into ad position, viewing ads in higher positions to be better prospects than ads in lower positions (Animesh et al. 2010, Chen and He 2011, Ursu 2015).

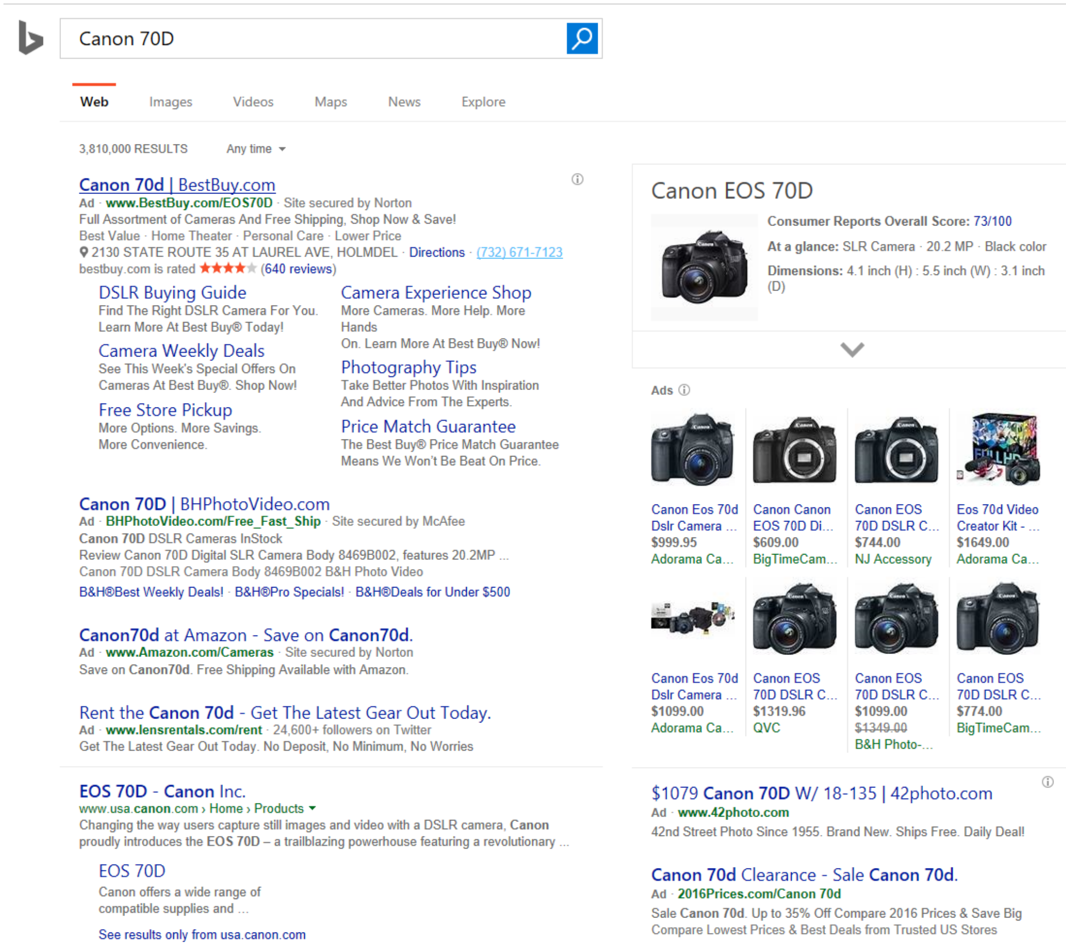

Figure 1: Search advertising

and Segal (2015) show that the CTR of particular advertisers in particular positions depends on who else is advertising and in what positions.

In this paper, we draw attention to advertiser prominence as another advertiser-specific factor that could challenge the assumptions of the theoretical models. As Figure 1 suggests, search ads differ not only in position but also in advertiser prominence. For example, a user looking at this display might recognize the retailers Amazon and Best Buy, but perhaps not the others.[2] With ads varying both in position and advertiser prominence, a natural question to ask is, what role does each type of prominence play in consumers' clicking decisions? How

[2]Traffic data from Alexa.com are consistent with this intuition. On a recent search at this website, Amazon had the 4th-highest U.S. traffic among all websites, followed by Best Buy at 49, B&H at 291, and 42nd Street Photo at 21,373.

will the two forms of ad prominence interact in determining CTR? Specifically, how will the marginal CTR of ad position change with advertiser prominence: who benefits more from moving up a slot, an advertiser with low prominence or an advertiser with high prominence? These questions are relevant to academics and practitioners alike, for they have implications for auction design and ad budgeting. More fundamentally, they speak to whether search advertising follows the same rules as traditional advertising or whether it is in a special category altogether.

We examine the interaction between ad position prominence and advertiser prominence by bringing together two data sets: (i) individual-level responses to search ads on Microsoft's Live Search platform, and (ii) contemporaneous measures of advertiser prominence from Alexa.com, a company that ranks advertisers by their website traffic. The click-stream data were provided to us by Microsoft Research as part of their *Beyond Search* initiative.[3] It is a curated sub-sample of search impressions from Live Search for the period August-November 2007. Two features of the data are noteworthy. First, ad positions show significant variation over short periods of time, even minutes. Much of this variation is due to the random perturbation of bids by Microsoft's ad-delivery platform while allocating advertisers to slots. Second, according to Microsoft Research, our sample does not contain individually targeted advertising, i.e., advertisers in the data were not adjusting their bids based on user characteristics such as sex, geographic location or demographics. These features of the data allow us, and others who have used this data before, e.g., Gomes et al. (2009) and Jeziorski and Segal (2015), to use within-advertiser position variation to identify advertiser-level position effects and correlate them with the measures of prominence from Alexa.

Our main analysis focuses on the ads served to users who searched for particular brands of cameras (e.g., Nikon). Since product brand searches usually reveal a transactional intent for the brand in question, advertisers to these searches tend to be retailers. Retailers also happen to be the largest single category of search advertisers;[4] hence our results are applicable to a large portion of the search-advertising market.

[3] `https://blogs.msdn.microsoft.com/msdnat/2007/09/11/request-for-proposal-beyond-search-semantic-computing-and-internet-economics`.

[4] Interactive Advertising Bureau, 2015 Full-Year Results; see `http://www.iab.com/wp-content/uploads/2016/04/IAB_Internet_Advertising_Revenue_FY_2015.pdf`.

To preview our results, first, we find evidence of non-sequential clicking behavior. In nearly two-thirds of our impressions with clicks, a click on the jth ad is generally not the jth click. Furthermore, in nearly a third of the impressions with two or more clicks, the consumer changes direction, i.e., proceeds upward after proceeding downward. These observations suggest that advertiser identity might play a role in people's clicking behavior. We find that ad position prominence and advertiser prominence are substitutes in consumers' clicking behavior; the marginal-CTR value of position is greater for less prominent advertisers than for more prominent advertisers. For instance, in the Nikon keyword, CTR increases by approximately 2 percentage points when an advertiser *not* in the Top-100 of Alexa rankings goes from the second position to the first position, whereas for a Top-100 Alexa advertiser, the same position improvement yields a statistically negligible CTR-effect. Results for the other camera brands, Canon and Olympus, are similar, as are the results for non-camera brands such as Nike and Puma, as well as for non-brand keywords like "textbook."

Our results contradict the independence assumption of Edelman et al. (2007) and Varian (2007). They also contradict the multiplicative assumption of Aggarwal et al. (2006), Katona and Sarvary (2010), Pin and Key (2011), and Nekipelov (2014). For this model says that, to the extent advertiser prominence matters for CTR, those effects must be embedded in β_k, which, as a result, must be (i) positive, and (ii) larger for more prominent advertisers. However, then, $\text{CTR}_{jk} - \text{CTR}_{j+1,k}$ must increase in advertiser prominence, whereas we find that it is decreasing. In particular, our results show that CTR-elasticity of ad position is decreasing in advertiser prominence, which has implications for how advertisers should bid and how the position auction itself should be designed.

The observation that less prominent advertisers benefit more from ad position echoes findings of a similar nature reported in the literature. For example, Jerath et al. (2011) find that particular advertisers in lower positions may get higher CTRs than advertisers above them. Narayanan and Kalyanam (2015) observe that the two smallest advertisers in their 3-firm data set have stronger position effects than the third. They also find that advertisers experience larger position effects in keywords where they receive lower quality scores from the search engine than in keywords where they receive higher quality scores. In contrast, (i) our analysis is keyword by keyword, (ii) our advertising data includes all advertisers who advertise to a keyword, and (iii) our advertiser prominence measure is Alexa rank, which is not only publicly observable, but also exhibits wide variation in our data.

Our results also complement Blake et al.'s (2015) finding, on the basis of field experiments at eBay, that "the effectiveness of SEM [search engine marketing] is small for a well-known company like eBay ... because users would have found the advertisers' sites anyway." Specifically, they find that for keywords containing the eBay brand, search advertising did not generate incremental clicks and sales. For keywords not containing eBay's brand, also, search advertising did not lead to incremental sales, but in this case there was a statistically significant increase in the number of new registered users and in sales to infrequent buyers (but not to frequent buyers). This suggests that advertiser prominence and effectiveness of search advertising are related, because the eBay brand is likely to be more salient among frequent users. Blake et al. (2015) focus is on the effectiveness of search advertising as a whole for a prominent advertiser like eBay. By contrast, our focus is within search advertising; we examine multiple advertisers to determine if consumers respond differently to different ad positions depending on the advertiser's overall prominence. So, whereas Blake et al. (2015) end up concluding that search ads and organic links are substitutes, we end up concluding that search ad position and advertiser prominence are substitutes.

Goldman and Rao (2014) have concurrently obtained results similar to ours using past experiments at Bing as the basis for identifying variation in advertiser position. They, too, document heterogeneity in position effects across different advertisers. In particular, they find that advertisers not featured in the keyword benefit more from higher positions, than the advertiser featured in the keyword. For keywords in which no advertiser is featured, less well-known (per Alexa rank) and "higher quality"[5] advertisers benefit more from higher position. In obtaining these results, Goldman and Rao (2014) pool over many keywords, whereas we treat each keyword separately, which allows the interaction between prominence and position effect to be keyword specific. Furthermore, our branded keywords are product-brand keywords, not advertiser-brand keywords.

[5]By "higher quality" Goldman and Rao mean advertisers whose websites have lower "bounce" rates.

2 Data and institutional details

When a user submits a search query (a "search string") with commercial value to a search engine, a list of ads appears, what we call an "impression."[6] Sponsored links appear at the top and right-hand side of the so-called organic links—the unpaid links the search engine itself produces based on its proprietary algorithms. Each search ad is a brief paragraph of text—perhaps two or three lines—of which the most notable part is the advertiser's (clickable) web address (see Figure 1). Advertisers bid for advertising slots associated with particular keyword(s) by submitting the price per click they are willing to pay. The search engine weights the bids with proprietary quality scores and runs a generalized second-price auction (Edelman et al. 2007, Varian 2007), the outcome of which is an ordering of ads and a price-per-click for each advertising slot. The search engine is compensated only if a consumer clicks on a sponsored ad.[7]

Our search ads data come from Live Search, Microsoft's search platform from 2006 to 2009 (before it morphed into Bing). In 2007, the time of our data, Live Search had about 10% of all online search queries; in comparison, Google had 51% and Yahoo had 22%.[8] In paid-search advertising, Microsoft's market share was even smaller. Until May 2006, Live Search's ad auctions were being conducted by Yahoo; in May 2006, Microsoft switched to its own ad delivery platform, adCenter.

In 2008, as part of its *Beyond Search* initiative, Microsoft Research made available to a limited set of academics a curated subsample of 20 million search impressions. The impressions were sampled from the ones that appeared in a roughly three-month period: 10 August-1 November 2007. The sampling scheme involved selecting an impression at random from the log and then including all the other impressions displayed to the same user during the same session. As already noted, the sampled impressions did not contain any user-targeted advertising.[9] Impressions that were part of longer user sessions had a propor-

[6]Note that our definition of an impression is different from Goldman and Rao's (2014). They define impressions as "ad appearances"; we define it as "ad list appearances."

[7]See Pin and Key (2011) and Nekipelov (2014) for more institutional details about search advertising.

[8]Dow Jones News Service, 9 May 2007.

[9]This is a statement about our sample, not about Microsoft's ad-targeting capabilities in general. According to `https://www.clickz.com/better-targeting-with-adcenter/47836/` and

tionally higher probability of being in the data set than impressions from shorter sessions. However, since the vast majority of the sessions contained only one impression, associated with a particular user's search, we believe sample selection caused by including longer user sessions is negligible. The average length of a session was about ten minutes. Search impressions and user activity are well documented in the data. For each impression we have the search string that originated the impression, the list of sponsored links, order of sponsored links, identity of advertisers, and time stamps of all clicks on sponsored links in a session.

This data set is fairly unique for the search advertising literature. Other than Gomes et al. (2009) and Jeziorski and Segal (2015) who use the same data source as us, and Goldman and Rao (2014), we are not aware of any empirical study that uses impression-level data, with *all* advertisers in every impression accounted for. By contrast, Ghose and Yang (2009), Yang and Ghose (2010), Agarwal et al. (2011), and Animesh et al. (2011) have data from a single advertiser only, and Narayanan and Kalyanam (2015) have data from four advertisers (that eventually merged into one firm), which gets reduced to three when they conduct between-advertiser analyses. Clearly, in order to examine the role of advertiser prominence in consumers' clicking decisions, we need a data set with multiple advertisers varying in prominence. And to examine the role of ad position, we need rich variation in ad positions, not only across advertisers, but also within advertisers, across impressions. Our impression-level data has this variation as we will discuss later. By contrast, Ghose and Yang (2009) and Yang and Ghose (2010) use weekly average data, and Animesh et al. (2011) and Narayanan and Kalyanam (2015) use daily average data, aggregating over many impressions in the process.

For our main analysis, we focus on three camera-brand keywords: "Canon," "Nikon," and "Olympus." We chose these brands because they (a) generate a large number of impressions in our data set, and (b) are inherently specific enough to elicit the same set of advertisers in a large proportion of their impressions—for any manifestation of the keyword

https://www.seroundtable.com/archives/002904.html, in 2007 Microsoft could target the limited subset of users who voluntarily shared their personal data with Microsoft while setting up accounts on various Microsoft properties such as MSN, Hotmail, etc. We thank an anonymous referee for these and other references listed in the online appendix.

(e.g., "Nikon," "Nikon camera," "Nikon D40," etc.) and any match specification ("broad match" or "exact match").[10] Such keywords tend to be well-known product brands that participate in a fairly narrow set of related product categories; advertisers advertising to them tend to be retailers.[11] For example, Canon, Nikon and Olympus are makers of photographic equipment—cameras, lenses, and photographic accessories. Retailers who sell any one of these photographic categories, tend to sell all of them.[12]

Besides our main analysis on the three camera brands, we examine a set of brands that participate in a broader set of categories—Nike, Adidas, Puma, Sony, Yamaha, and Maytag—and two non-branded keywords, "textbooks" and "e-books," that are really product categories (of which the second is broader than the first). Our final sample consists of $28,153$ search impressions for the three camera brands and $252,138$ search impressions for all the other keywords combined.

We supplement the Microsoft data with data from Alexa.com. Alexa is a company that ranks websites by their traffic using a global traffic panel, which, according to its website, "is a sample of millions of Internet users using one of over 25,000 different browser extensions," and "sites that have chosen to install the Alexa script on their site and certify their metrics" (`www.alexa.com/about`). Its rank "is a measure of how a website is doing relative to all other

[10]These terms refer to how the advertiser would like its keyword to be interpreted by the search engine when generating ads. For example, an advertiser buying the "Nikon D40" keyword under an "exact match" specification seeks to advertise only when that specific keyword is searched; by contrast, when this advertiser purchases "Nikon D40" under a "broad match" specification, then the advertiser's intent is to advertise to searches for "Nikon D40 camera," "Nikon camera," "Nikon flash," "D40 lens," etc.

[11]Our definition of branded keywords is thus quite different from Blake et al.'s (2015), who use the term to refer to keywords where the advertiser brand itself is featured, such as "toasters Macy's" or simply "Macy's." A user searching for "Macy's"—without additional specification—could be searching for a wide variety of products, from toasters to mattresses. While Macy's would probably be interested in advertising to this user regardless, a retailer such as Craig's Beds in New York City would probably prefer to advertise to "mattresses Macy's" but not to "Macy's."

[12]Digital camera brand keywords were chosen to facilitate our focus on the relationship between advertiser prominence and position, as opposed to the relationship between organic and sponsored links. While the advertisers to these keywords are retailers, the top organic results pertain to manufacturers or camera-review sites, such as dpreview.com. The relationship between organic and sponsored links is analyzed by Yang and Ghose (2010), Blake et al. (2015), and Simonov et al. (2015).

sites on the web over the past 3 months. The rank is calculated using a combination of the estimated average daily unique visitors to the site and the estimated number of pageviews on the site over the past 3 months." We used Alexa's API to download daily rankings of each of our advertisers during our sample period. The variance in daily rankings is negligible; we therefore use the 3-month average rank of each advertiser as a measure of its prominence, with lower ranks signifying greater prominence.[13]

3 Basic features of the data

Table 1 presents the key statistics of search impressions for the camera brand keywords and other keywords. 13% of camera brand impressions have at least one click on a search ad and 2.2% of those impressions have two or more clicks on search ads.[14] 63% of the impressions with at least one click are impressions with non-sequential clicks.[15] In other words, most ad impressions for camera brands do not involve clicking from top to bottom. Furthermore, in 28% of the impressions with at least two clicks, users click on a higher ad after clicking on a lower ad. This pattern contradicts the so-called cascade models of clicking behavior (Craswell et al. 2008). The last four rows of Table 1 show decreasing CTR as we go down the ad list. The decrease is rather steep, with fourth and lower slots receiving less than one-fifth of the CTR of the top slot. "Other keywords" behave more or less similarly to the camera brands, the most notable difference being that the proportion of non-sequential clicks is significantly smaller.

Corresponding descriptive statistics about the advertisers advertising to camera brand searches are presented in Table 2. (Advertiser statistics for other keywords are omitted for

[13]For other studies using Alexa ranks and similar measures of advertiser prominence see Goldberg and Hartwick (1990), Brynjolfsson and Smith (2000), Pham and Johar (2001), Animesh et al. (2010), Goldman and Rao (2014), and Agarwal et al. (2015).

[14]The first of these numbers is higher than the typical CTRs reported in the literature because it is the CTR for an *impression*, not for a particular advertiser. Our advertiser-specific CTRs, in Table 3, are in line with what has been reported before (Animesh et al. 2011, Narayanan and Kalyanam 2015).

[15]As noted earlier, non-sequential clicks are those for which click position does not equal ad position. For example, a consumer whose first click is on the second ad has made a non-sequential click, as has a consumer whose second click is on the first ad. Non-sequential clicks always involve "jumps" over some ads.

	Camera brands	Other keywords
Impressions	28,153	252,138
Impressions with no clicks	87.04%	91.73%
Impressions with 1 click	10.75%	7.33%
Impressions with 2 or more clicks	2.69%	1.28%
Impressions with non-sequential clicks (out of impressions with at least 1 click)	62.54%	45.29%
Impressions with out-of-order clicks (out of impressions with at least 2 clicks)	28.25%	27.49%
CTR of the top slot	5.65%	4.91%
CTR of the second slot	3.86%	2.76%
CTR of the third slot	2.69%	1.97%
CTR of fourth and lower slots	1.07%	0.74%

Table 1: Descriptive statistics

brevity.) In this table we show the top-4 advertisers with the most clicks for each brand; the rest are pooled into a catch-all "other advertisers" group. The CTR column shows the CTRs of these advertisers, sorted by total number of clicks. The sorting is brand-specific. Thus, the advertiser with the most clicks for Canon has a CTR of 14%, but this same advertiser, when advertising to Nikon, is the third-most clicked advertiser for that keyword, and gets a CTR of 11.5% there; in turn, the most-clicked advertiser for Nikon has a CTR of 5.9% there, but doesn't figure in the top-4 for any of the other keywords. Even though the top-4 advertisers are not necessarily the same across camera brands, they often overlap, because retailers selling any one of these camera brands is likely to be selling all of them. Since our analysis does not pool impressions across camera brands, we treat each brand's advertisers as distinct. Note that the CTRs in Table 2 are not monotonic. This is because some advertisers figure in more impressions than others and hence generate more clicks even with lower CTRs.

The Alexa rank column in Table 2 shows the average Alexa world-ranks of the four most-clicked advertisers and "other advertisers." In general, the 4-most-clicked advertisers have

	Keyword	CTR	Alexa rank
1st-most clicked advertiser	Nikon	5.9%	$14,650$
2nd-most clicked advertiser	Nikon	10.7%	515
3rd-most clicked advertiser	Nikon	11.5%	$66,379$
4th-most clicked advertiser	Nikon	3.1%	24
Other advertisers	Nikon	2.6%	$314,537$
% of "other advertisers" ranked in Alexa Top-100	Nikon	-	7%
1st-most clicked advertiser	Canon	14.0%	$66,379$
2nd-most clicked advertiser	Canon	2.6%	24
3rd-most clicked advertiser	Canon	2.4%	98
4th-most clicked advertiser	Canon	3.4%	540
Other advertisers	Canon	2.4%	$255,582$
% of "other advertisers" ranked in Alexa Top-100	Canon	-	6%
1st-most clicked advertiser	Olympus	3.0%	24
2nd-most clicked advertiser	Olympus	4.9%	1117
3rd-most clicked advertiser	Olympus	4.4%	$16,455$
4th-most clicked advertiser	Olympus	4.1%	540
Other advertisers	Olympus	2.3%	$253,532$
% of "other advertisers" ranked in Alexa Top-100	Olympus	-	9%

Table presents CTRs and Alexa ranks of the top 4 most-clicked advertisers for each keyword; as the Alexa ranks indicate, these advertisers are not necessarily the same across keywords. Higher Alexa ranks correspond to less prominent advertisers.

Table 2: Click-through-rates and Alexa (world) ranks of advertisers

lower Alexa ranks, i.e., are more prominent, than "other advertisers." However, amongst the 4-most-clicked advertisers are several non-prominent advertisers, such as the most-clicked advertiser for Canon. Additionally, note that 6–9 percent of the low-clicked "other advertisers" are ranked in the Top-100 at Alexa, which places them amongst the most recognizable advertisers on the Internet.

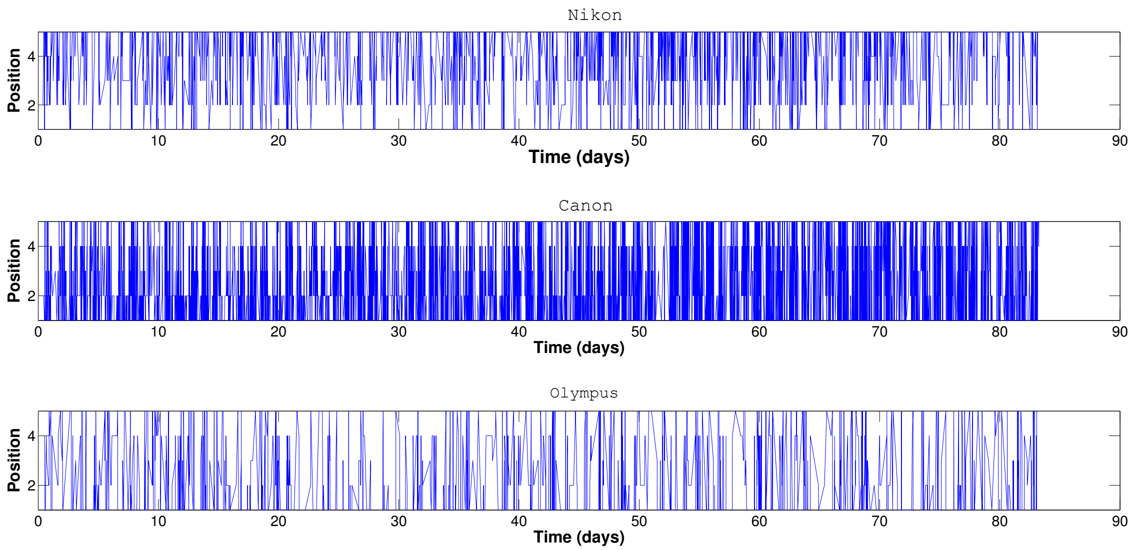

Figure 2: Position variation of the most clicked advertiser.

In stark contrast to the stability in advertisers' Alexa ranks, is the pronounced instability in their ad positions. Figure 2 shows the ad position of the most-clicked advertiser for each camera brand, clicked-impression to clicked-impression, over the three months of our data. These graphs demonstrate that there is significant variation in the position of the most clicked ad over very short periods of time—even minutes. The coefficient of variation of ad position is greater than 40% for the average advertiser with at least ten impressions. Others working with Microsoft data have reported similar levels of ad position variation (Pin and Key 2011, Athey et al. 2014). For example, the latter note that "for the search phrases we consider, the most commonly observed advertisements have a standard deviation of their position number ranging from about one third of a position, to about 2 positions."

4 Results

A. Empirical model. The variation in ad positions within advertisers, and the cross-sectional variation in Alexa ranks across advertisers, provides the basis of our empirical strategy for disentangling the effects of ad position and advertiser prominence.

We estimate a series of linear probability models for each keyword in which the dependent variable is a click, and the explanatory variables include advertiser fixed effects, ad position fixed effects, and interactions between ad position and advertiser prominence. That is, we estimate several variants of the following model:

$$\Pr(\text{Click}_{ijk}) = \alpha_j + \beta_k + \gamma \times j \times \text{Alexa}_k + \epsilon_{ijk}, \tag{1}$$

where i indexes impressions, j denotes ad positions, and k identifies advertisers. Since we run these regressions for each keyword, parameters α_j, β_k and γ are allowed to vary across keywords in an unrestricted way.

Advertiser prominence is proxied by Alexa ranks, in three different ways: (i) as a dummy variable identifying Top-100 Alexa-ranked advertisers (Alexa rank $\leq 100 = 1$, Alexa rank $> 100 = 0$), (ii) a dummy variable identifying Top-500 Alexa-ranked advertisers[16] (Alexa rank $\leq 500 = 1$, Alexa rank $> 500 = 0$), and (iii) as the continuous variable, $1/(\text{Alexa rank})$.

Equation (1) is not to be interpreted as a structural equation, but rather as a reduced-form specification. Its purpose is to assess how position effects vary across advertisers with different levels of prominence; it is not meant to provide a specification test of the multiplicative model (as in Jeziorski and Segal 2015). The main advantage of the reduced-form approach is that it allows us to control better for the unobserved heterogeneity of advertisers by using more fixed effects without making specific structural assumptions.

B. Identification. Identification in our model rests on the assumption that within-advertiser variation in ad positions in the data is exogenous. In turn, it relies on two assumptions about our data set: (i) that no user targeting is present, i.e., advertisers' positions (within each

[16] 18% of the advertisers for the digital camera keywords are Top-100 Alexa ranked and 10% of the advertisers are Top-500 Alexa ranked, but not Top-100 Alexa ranked. These percentages are similar across keywords. Moving from the Top-100 to Top-500 Alexa ranking adds a few prominent brick-and-mortar retailers.

keyword) do not vary based on user characteristics, and (ii) that short-run changes in advertiser characteristics, unobserved by the econometrician but observed by the user, are not synchronized with changes in advertiser positions.

As noted earlier, according to representations made by Microsoft Research, the former assumption is true. Lacking data on user characteristics, we cannot test this assumption further.[17] As for the second assumption, we discuss why it might be true below; later, in Section 5, we test the robustness of our results with respect to this assumption. Notwithstanding all this, if either of these assumptions is not satisfied, our results may no longer be interpreted as causal.

What drives the within-advertiser variation in ad position in our data? We believe in large measure this is due to the way Microsoft's adCenter administers the auction. As Gomes et al. (2009), who use the same data as us, state:

> ... a significant source of variation is due to the allocation procedure itself. Microsoft AdCenter applies a randomization procedure that perturbs submitted bids and (non-deterministically) changes the slot allocations. This makes the variation exogenous.

When advertisers submit bids for a keyword, they intend to obtain a certain position. However, the position that they actually obtain is a noisy realization of their intention. In other words,

$$j_{ki} = j'_{ki} + \epsilon'_{ik}, \tag{2}$$

where j_{ki} is advertiser k's actual ad position in impression i, j'_{ki} is advertiser k's intended ad position in impression i, and ϵ'_{ik} is a mean-zero random variable, independent of j'_{ki}, representing the noise introduced by the auction mechanism.[18]

[17]We conducted a test for geo-targeting by examining location-based keywords, such as "Chicago" and "Seattle," under the presumption that location-based keywords have less heterogeneity in the location of the user. We report the results in the online appendix. As noted there, our interaction effects continue to appear, but the estimates are noisier due to lower CTR in these keywords. We note that besides news outlets, employment agencies, and local retailers, some of the advertisers in these keywords are tourist agencies who might be interested in geo-targeting.

[18]Intended ad position, one might reasonably assume, is a function of advertiser characteristics—both short-term and long-term—as it is in all forms of advertising. In TV and print advertising as well, different

Microsoft-induced noise is not the only reason advertisers can't realize their intentions. Basic parameters of the position auction are themselves uncertain. No advertiser knows, for instance, how many others will be bidding in a particular auction, nor the identities of those advertisers. Moreover, they do not know what quality scores the search engine will assign them in the current auction, nor their competitors' quality scores, both of which play an important role in determining ad position. Pin and Key (2011, p. 70) suggest that these unobserved-by-the-advertiser variations drive the variations in ad positions within advertisers:

> ... most ads enter many auctions with the same bid, facing multiple opponents whose identities differ from one auction to the other, and with highly irregular weights attributed to the bids.[19]'

Similarly, Athey and Nekipelov (2010) state:

> In practice quality scores do vary from query to query, queries arrive more quickly than advertisers can change their bids, and advertisers cannot perfectly predict changes in quality scores [...] Although bids can be changed in real time, the system that runs the real-time auction is updated only periodically based on the state at the time of the update, so that if bids are adjusted in rapid succession, some values of the bids might never be applied.

In 2007, the time of our data, changes to the submitted bids required a manual intervention by the advertiser. Advertisers chose bids over a series of auctions, not auction by auction. Thus, synchronizing changes in ad position to very short-term changes in advertiser characteristics was difficult. We elaborate on this more in Section 5.

---

advertisers seek different ad placements, and the same advertiser may seek different ad placements at different times. However, the big difference between traditional advertising and search advertising is that in the former intended ad position is generally equal to actual ad position: TV and print advertisers purchase specific ad placements directly, based on posted or negotiated prices.

[19] In their study the typical advertiser's quality score has a coefficient of variation of 40%—almost exactly the coefficient of variation we observe for position.

C. Camera keywords. Tables 3-5 show the results of various regressions of type (1) for the camera keywords.

The top panel of each table confirms the presence of strong ad position effects. Being in the top ad position increases CTR by an average of 4- to 6-percentage points depending on the camera brand (relative to ad positions 6 and lower). The position effect is not statistically significant beyond position 3 for Nikon, and beyond position 4 for Olympus.

	Top-100 Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.053** (0.004)	0.051** (0.004)	0.055** (0.004)	0.051** (0.004)	0.051** (0.004)
Pos. 2	0.031** (0.003)	0.031** (0.003)	0.030** (0.003)	0.032** (0.003)	0.032** (0.003)
Pos. 3	0.016** (0.003)	0.017** (0.003)	0.016** (0.003)	0.017** (0.003)	0.017** (0.003)
Pos. 4	-0.002 (0.002)	-0.002 (0.002)	-0.001 (0.002)	-0.002 (0.002)	-0.002 (0.002)
Pos. 5	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)
Pos. 1×Top Alexa	-0.020* (0.010)	-	-0.023** (0.008)	-	-
Pos. 2×Top Alexa	0.010 (0.010)	-	0.010 (0.008)	-	-
Pos. 3×Top Alexa	0.013 (0.008)	-	0.006 (0.006)	-	-
Pos. 1-3×Top Alexa	-	0.002 (0.006)	-	-0.001 (0.004)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0050 (0.0115)
N	39016	39016	39016	39016	39004
R^2	0.029	0.029	0.030	0.029	0.029

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$.

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 3: Linear probability model predicting clicks for the Nikon keyword

The bottom panels of these tables show the interaction between advertiser prominence and

	Top-100 Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.041** (0.002)	0.039** (0.002)	0.041** (0.002)	0.040** (0.002)	0.039** (0.002)
Pos. 2	0.028** (0.002)	0.027** (0.002)	0.028** (0.002)	0.028** (0.002)	0.027** (0.002)
Pos. 3	0.018** (0.002)	0.019** (0.002)	0.017** (0.002)	0.019** (0.002)	0.019** (0.002)
Pos. 4	0.005** (0.001)	0.005** (0.001)	0.005** (0.001)	0.005** (0.001)	0.005** (0.001)
Pos. 5	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)	0.002* (0.001)
Pos. 1×Top Alexa	-0.010** (0.004)	-	-0.008** (0.004)	-	-
Pos. 2×Top Alexa	-0.005 (0.004)	-	-0.005 (0.003)	-	-
Pos. 3×Top Alexa	0.001 (0.004)	-	0.002 (0.003)	-	-
Pos. 1-3×Top Alexa	-	-0.005* (0.003)	-	-0.003 (0.002)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0030 (0.0062)
N	105427	105427	105427	105427	101229
R^2	0.022	0.022	0.022	0.022	0.022

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$.

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 4: Linear probability model predicting clicks for the Canon keyword.

ad position. We estimate five specifications and we frequently find a statistically significant impact of advertiser prominence on the size of the position effect. The most robust effect is the decrease of the effect of the top slot (the difference between ad position 1 and ad position 2) for advertisers in Top-100 and in Top-500 of the Alexa classification, relative to less prominent advertisers. Specifically, for the Nikon keyword, being in the top of the Alexa classification completely nullifies the marginal effect of the top slot. In other words, retailers with a Top-100 and Top-500 Alexa ranking obtain the same CTRs in ad positions 1 and

	Top-100 Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.044** (0.005)	0.045** (0.005)	0.048** (0.005)	0.047** (0.005)	0.045** (0.005)
Pos. 2	0.025** (0.004)	0.025** (0.004)	0.026** (0.005)	0.027** (0.004)	0.025** (0.004)
Pos. 3	0.015** (0.004)	0.014** (0.004)	0.015** (0.004)	0.016** (0.004)	0.014** (0.003)
Pos. 4	0.005* (0.003)	0.005* (0.003)	0.005* (0.003)	0.005* (0.003)	0.006** (0.003)
Pos. 5	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.004 (0.003)
Pos. 1×Top Alexa	0.008 (0.015)	-	-0.014* (0.009)	-	-
Pos. 2×Top Alexa	-0.006 (0.010)	-	-0.006 (0.008)	-	-
Pos. 3×Top Alexa	-0.014** (0.007)	-	-0.006 (0.006)	-	-
Pos. 1-3×Top Alexa	-	-0.005 (0.006)	-	-0.009* (0.005)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0226** (0.0100)
N	18608	18608	18608	18608	18177
R^2	0.018	0.018	0.018	0.018	0.018

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$.

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 5: Linear probability model predicting clicks for the Olympus keyword.

2, whereas, retailers not in the Top-100 of Alexa rankings obtain nearly 50% higher CTR in the top position versus position 2.[20] The impact of advertiser prominence is similar for the Canon keyword. Here, not Top-100 Alexa retailers enjoy a 1.3-percentage point CTR increase (approximately 33% increase) when moving from position 2 to position 1, whereas,

[20]This does not mean that top Alexa-ranked retailers get lower CTRs than non-top Alexa-ranked retailers. As noted earlier, the model includes advertiser fixed effects. Top Alexa-ranked retailers tend to have larger fixed effects than non-top Alexa-ranked retailers.

Top-100 Alexa retailers enjoy only 0.3-percentage point CTR increase (approximately 8% increase) for the same position improvement. Lastly, for the Olympus keyword, sample sizes are significantly smaller, so it is harder to find statistically significant advertiser prominence effects. Still, even here, being a Top-500 Alexa-ranked advertiser nearly nullifies the marginal effect of the top ad position, and also has a significant negative effect on the CTRs of the top-3 slots, collectively. In particular, Top-500 Alexa retailers suffer a 1-percentage point lower CTR-gap between the top-three slots and lower slots, which is equivalent to nearly nullifying the positive impact of slot 3 and decreasing the positive impacts of slot 1 and 2 by 20% and 33%, respectively.

D. Non-camera keywords. The analysis of non-camera keywords is summarized in Table 6. For each keyword, we report the interaction between top ad position and advertiser prominence from two regressions, differing only in the measure of advertiser prominence used. In other words, these regressions replicate column I and column III in Tables 3–5. As can be seen, the interaction effects noted above for camera keywords are essentially replicated for non-camera keywords. Despite the greater breath of these keywords, in five of the eight keywords, high-Alexa-ranked advertisers do not benefit from the top ad position as much as lower-Alexa-ranked advertisers. Moreover, the magnitude of the interaction coefficients is very similar to what we observed with cameras.

Interaction between the Alexa dummies and ad position is statistically insignificant for Yamaha, Maytag and "textbooks." This could be because there is, in fact, no interaction between advertiser prominence and ad position for these keywords, but we suspect that the negative result is related to some peculiarities of these keywords. Yamaha is arguably the broadest keyword in our sample because it contains retailers who sell stereo equipment as well as retailers who sell motorcycles. Advertisers who are prominent in one category may not be prominent in another. Consumer behavior will, be guided by what category they had in mind when they searched, which we do not observe generally. This contributes noise, and perhaps even bias, in our estimates. On the other hand, Maytag has the smallest number of clicks in our data; we suspect that the Maytag interaction coefficients would be significant if more data were available. Lastly, "textbooks" is an unusual keyword because it elicits advertising

Nike	Position 1 × Top-100 Alexa	-0.039** (0.007)
	Position 1 × Top-500 Alexa	-0.042** (0.005)
Adidas	Position 1 × Top-100 Alexa	-0.011 (0.015)
	Position 1 × Top-500 Alexa	-0.021* (0.012)
Puma	Position 1 × Top-100 Alexa	-0.029* (0.016)
	Position 1 × Top-500 Alexa	-0.034** (0.012)
Sony	Position 1 × Top-100 Alexa	-0.022** (0.003)
	Position 1 × Top-500 Alexa	-0.015** (0.003)
Yamaha	Position 1 × Top-100 Alexa	0.004 (0.005)
	Position 1 × Top-500 Alexa	0.002 (0.004)
Maytag	Position 1 × Top-100 Alexa	-0.013 (0.012)
	Position 1 × Top-500 Alexa	-0.008 (0.012)
Textbooks	Position 1 × Top-100 Alexa	0.008 (0.011)
	Position 1 × Top-500 Alexa	0.002 (0.010)
Ebook	Position 1 × Top-100 Alexa	-0.016** (0.001)
	Position 1 × Top-500 Alexa	-0.015** (0.001)

Each cell is a separate regression replicating columns I and III of Tables 3–5. We report only relevant interaction terms.

Table 6: Ad position-advertiser prominence interaction effects for non-camera keywords

from a retailer with an extremely suggestive .com name. This advertiser is essentially a monopolist, harvesting nearly 50% of all clicks. It has a strong position effect despite being not very prominent on Alexa—which is consistent with our other findings. However, due to the large concentration of clicks at this retailer, it is hard to statistically identify the position effects of other retailers, which results in the insignificant Alexa-ad position interaction.

In short, Tables 3–6 suggest that across a broad cross-section of search keywords there is a negative interaction between ad position prominence and advertiser prominence. More prominent advertisers seem to gain fewer clicks from higher search ad positions than less prominent advertisers. Before we take these findings as conclusive, however, we should evaluate the robustness of our estimates. This we do next.

5 Robustness

The principal concern about our results is that ad positions are not exogenous but rather that they are correlated with unobserved advertiser and user characteristics. In that case, the data-generating process for clicks may not be as stated in (1), but rather

$$\Pr(\text{Click}_{ikj_k}) = \alpha_{j_{ki}} + \beta_k + \gamma \times j_{ki} \times \text{Alexa}_k + \epsilon_{ikj_{ki}}, \tag{3}$$

where j_{ki} is the position assigned to advertiser k in impression i by the search platform. Then, to the extent assigned ad positions are correlated with unobserved advertiser and user characteristics affecting the probability of a click, our position and interaction effect estimates might be biased due to endogeneity.

As already noted, according to the information provided to us, all users searching a particular keyword at a given time faced the same distribution of advertisers across slots. Our regressions take into account observed advertiser characteristics via Alexa ranks and unobserved advertiser characteristics via fixed effects. However, advertiser characteristics may be changing over time. For instance, a retailer's inventory in Nikon cameras might be different at different times, which prompts the retailer to be more or less aggressive in its pricing policy with respect to these cameras, which, in turn, prompts it to be more or less aggressive in seeking higher ad positions at different times. This concern, however, depends on two presumptions: (i) that advertisers can fine-tune their positions to reflect changes in their characteristics,[21] and (ii) that consumers observe private signals about those changes

[21] A correlation between ad position and underlying retail prices can develop even without the retailer seeking higher ad positions under advantageous circumstances if the search engine does the work for him—for instance, by boosting quality scores whenever it detects lower prices at the retailer's landing page. However, for this mechanism, as well, the subsequent argument applies.

(not observed by the econometrician). Importantly, if consumers learn about changes in underlying advertiser characteristics only through ad position, endogeneity is not a concern. As noted in the Introduction, we do not take a position in this paper about how ad position affects CTR. It could be because consumers pay more attention to better-positioned ads; or it could be because consumers read meaning into position and interpret better-positioned ads as better clicking prospects.[22] Both are legitimate ad position effects for our purposes. In 2005, according to Fallows (2005), 62% of Internet users were not aware of the distinction between organic and paid-search results—only 1 in 6 searchers could "consistently distinguish between paid and unpaid results." Thus, it is likely that most consumers simply reacted to the overall prominence of the advertisers and their positions when making their clicking decisions.

User advertiser characteristics may also be changing over time. Gomes et al. (2009), who use the same data, express this concern in the following way:

> One might question the consistency of our estimates by arguing that the variation on slot allocations may be endogenous, 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).

However, this concern relies on the same presumptions that we discussed above.

Finally, as noted in Section 4B, the process underlying the determination of ad positions itself mitigates endogeneity. In the remainder of this section, we provide evidence that our estimates are driven by the random short-term variation in ad positions caused by the auction process itself and not the relatively longer-term variation in intended ad positions that responds to changes in unobserved advertiser and user characteristics. As evidence, we start with Figure 2. As noted already, this picture shows that there are no systematic patterns in advertiser positions for any of the camera keywords, Figure 3 shows another variant of this analysis.[23] Here we show the extent of position variation for search strings that do not

[22] Jeziorski and Segal (2015) discuss these two mechanisms in more detail and argue that each contributes to the position effect equally.

[23] These position variation graphs also suggest that neither throttling nor bid-scaling was going on to any substantial extent in our data. Throttling is when the search engine limits the participation of an advertiser in various auctions (to which it would be otherwise qualified) in order to stay within the advertiser's budget.

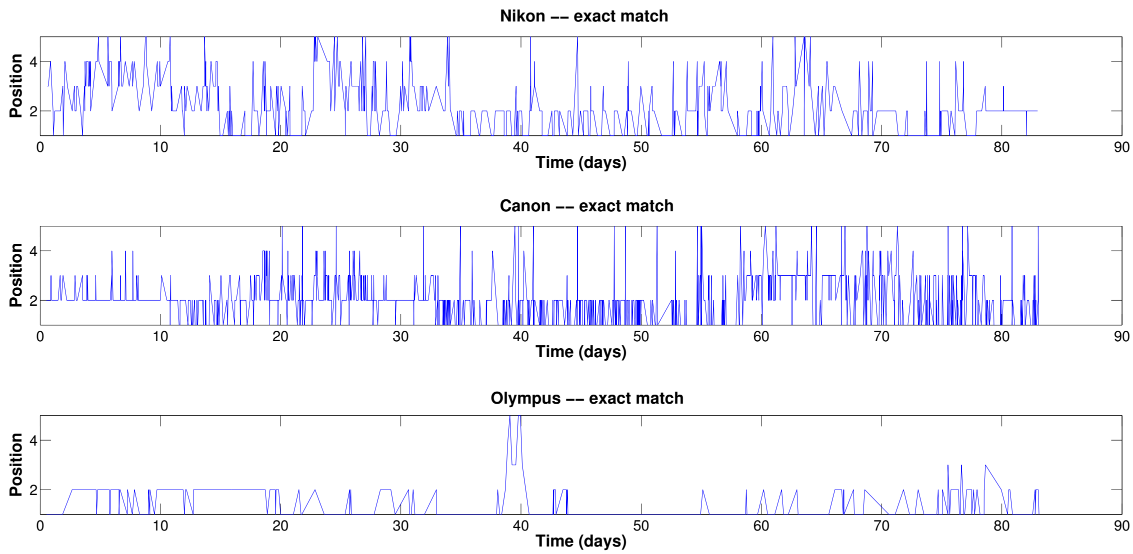

Figure 3: Position variation of the most clicked advertiser – exactly matched keywords.

contain extra words beyond "Nikon," "Canon," or "Olympus." Again, the pattern looks random, albeit coarser than before, reflecting the smaller number of data points. Finally, Table 7 shows a variance decomposition of ad position for the camera brands, by month, week, and day. We find that there is more variance in ad positions in the short-term than in the long-term.

We push this analysis in two directions. First, looking at the variation in ad positions over

Bid-scaling is when the search engine decreases the advertiser's bids automatically to stay within budget. The latter requires the advertiser's consent, and according to Karande et al. (2013), not many advertisers exercise this option. Under bid-scaling, assuming a first-come, first-served regime, one would expect to see trends developing over time where the budget-constrained advertiser gradually loses position. The fact that ad positions retain their essentially random character for the most-clicked advertiser over the entire 3-month period suggests either that bid-scaling was not being employed by this advertiser, or, that the most-clicked advertiser was not close to hitting its budget constraint. Another consequence of automated budget constraints is exit. Such exit could introduce bias if it is correlated with systematic selection of consumers; e.g., if exit only happens later in the day, and if consumers browsing at that time have different preferences. However, there are limits to this concern in our data, because we observe little exit; most of our variation is due to position flips between advertisers. The first-difference model discussed later in this section is a good test for such endogeneity.

	Keyword		
	Nikon	Canon	Olympus
Month	3.5%	1.1%	0.7%
Week	4.9%	2.4%	3.5%
Day	14.5%	7.1%	16.0%

The numbers correspond to the percentage of the within-advertiser variance of the position explained by cross-month, cross-week and cross-day variation.

Table 7: Variance decomposition of within-advertiser variation in positions.

very short periods of time, such as 5s—during which unobserved advertiser characteristics are unlikely to be changing—we run the linear probability model on first-differenced pairs of impressions very close in time. If our position effect estimates in Tables 3–6 are being estimated off of the exogenous short-term variation in ad positions, then this first-difference analysis should produce similar estimates. Second, to control for variation in unobserved advertiser characteristics over longer time intervals—say, two weeks—we re-estimate our linear-probability model using advertiser-time fixed effects. As we will show, this analysis, too, produces estimates close to our original ones.

A. First-differencing pairs of impressions close in time. For each advertiser we choose impression pairs in which the advertiser is displayed in different positions within a short period of time—ranging from 5 seconds to up to 12 hours—one of which produced a click while the other did not. We construct observations by first-differencing the impressions in each such pair, effectively "differencing out" any unobserved advertiser characteristics that might affect clicking behavior (because they are likely to stay constant over the short time interval). Then we run the following linear-probability regression:

$$c_{k,i} - c_{k,i-1} = \beta(\iota_{k,i} - \iota_{k,i-1}) + \epsilon_{k,i} - \epsilon_{k,i-1}, \tag{4}$$

where $c_{k,i}$ (resp. $c_{k,i-1}$) is the indicator function for a click on advertiser k in impression i (resp. $i-1$, a short time earlier), and $\iota_{k,i}$, e.g., $(0,0,\ldots,1,0,\ldots,0)$, and $\iota_{k,i-1}$, e.g., $(0,0,\ldots,0,0,\ldots,1)$, are vectors indicating the different ad positions of advertiser k in impressions i and $i-1$, respectively.[24] Under the realistic assumption that unobserved advertiser characteristics are constant within the short time interval between impressions $i-1$ and i in which the advertiser occupied different positions, the error terms cancel out, and the ad position effect estimates, $\hat{\beta}$, from this regression are robust to position endogeneity.

Table 8 presents these estimates; for comparison, we include fixed-effect estimates (from a specification without Alexa interaction effects).[25] The economic differences between the first-differenced estimates and the the fixed-effect estimates are small (except for Canon 5s, which we discuss below). For example, for the Nikon keyword, among impressions as close as 5s apart, the top ad position effect is .055 which is close to the fixed-effect estimate of .050 (and even closer to the column 1 estimate in Table 3). Even impressions up to 1h or 12h apart, which deliver more precise estimates because of the larger number of observations, are still close to the fixed-effect estimates.[26]

Only in the case of Canon, in the 5s window, are the first-difference estimates substantially different from the fixed-effect estimates. In this particular case, despite the large number of observations, there are only 26 relevant first-difference pairs containing one clicked and one not-clicked ad, i.e., there are only 26 observations with $c_{k,i} - c_{k,i-1} \neq 0$. In this case, even mild clustering of standard errors, would widen the 5% confidence intervals of 5s window estimates enough to contain the fixed-effects estimators.

We also performed a first-difference analysis of impressions where the search string exactly

[24]Note that the number of observations in these regressions is frequently larger than what we had for the fixed-effects regressions in Tables 3-5. This is because we pick all pairs of advertisers that switch positions, which combinatorially increases the size of the sample. For this reason, one may rightfully worry about clustering of standard errors; to alleviate this concern we use unclustered (underestimated) standard errors, under which we are more prone to detect differences between the first-differenced estimates and our main specification.

[25]Excluding Alexa interaction effects provides a conservative test of ad position endogeneity because it does not rely on Alexa ranks to absorb heterogeneity in long-term unobserved advertiser characteristics.

[26]Deriving a robust statistical test between the two models is difficult because the first-differenced model is estimated on a selected subsample. Presuming that the selection is not too large, and ignoring the difference in sample sizes between models, we performed a specification test suggested by Clogg et al. (1995). We cannot reject the null (at the 5% level for Nikon and Olympus and at the 1% level for Canon) that the fixed-effects model generates unbiased estimates of the top position effect, if the 5s model is the true data-generating process.

Nikon

	5s	1h	12h	Fixed effects
Position 1	0.055 (0.019)	0.049 (0.002)	0.058 (0.001)	0.050 (0.004)
Position 2	0.049 (0.014)	0.032 (0.002)	0.032 (0.001)	0.031 (0.005)
Position 3	0.009 (0.013)	0.017 (0.002)	0.014 (0.001)	0.016 (0.005)
N	480	43,301	593,877	39,016

Canon

	5s	1h	12h	Fixed effects
Position 1	0.014 (0.010)	0.033 (0.001)	0.038 (0.000)	0.036 (0.002)
Position 2	0.015 (0.009)	0.025 (0.001)	0.024 (0.000)	0.024 (0.003)
Position 3	-0.007 (0.008)	0.017 (0.001)	0.017 (0.000)	0.016 (0.003)
N	1,118	173,204	2,734,901	105,425

Olympus

	5s	1h	12h	Fixed effects
Position 1	0.046 (0.058)	0.042 (0.004)	0.036 (0.001)	0.042 (0.005)
Position 2	0.002 (0.046)	0.023 (0.004)	0.021 (0.001)	0.021 (0.005)
Position 3	0.043 (0.039)	0.009 (0.004)	0.008 (0.001)	0.010 (0.005)
N	138	9,028	99,708	18,608

We include the fixed effects regressions without Alexa interactions for comparison.

Table 8: Position effects estimated using first-differenced impressions for various time windows.

matched the brand, i.e., there were no extra words beyond the brand searched. The number of impressions is substantially smaller now, forcing us to use longer time windows. For the 3-day window, the effect of position 1 for "Nikon" is estimated to be 0.059 (0.014), whose 5% confidence intervals overlap with those of the fixed-effect estimates. We are unable to perform this check for other keywords without substantially increasing the time window—which would negate the premise of the test.

B. Advertiser-time fixed effects. To control for endogeneity over longer time windows, we estimate a specification with advertiser-time fixed effects, recognizing time in 2-week windows. Any variation in unobserved advertiser characteristics over two-week periods is now being absorbed by advertiser-biweekly fixed effects.

As Tables 9-11 show, these estimates are virtually identical to what we saw before in Tables 3-5. In other words, accounting for the passage of time in two-week intervals, within advertisers, does not affect position-effect estimates. Most important, the negative interaction between ad position and advertiser prominence remains—if anything, it is marginally strengthened.[27] We implement a formal test for nested models proposed by Clogg et al. (1995). The null hypothesis is that there is no estimation bias when estimating a model with advertiser fixed effects, if the model with advertiser-time fixed effects is the true data generating process. We cannot reject that, at the 5% level, the top position effects are estimated without bias in 13 out of 15 reported specifications (we marginally reject the null in the Specification II of Canon keyword and in the Specification V of Nikon keyword). We also cannot reject that the main interaction effects are estimated without bias in 13 out of 15 specifications (we marginally reject the null in the Specifications I and V of the Olympus keyword).

C. Exact keywords. In our main analysis we pooled across all search phrases that contain the brand keyword in question. However, an objection might be raised that advertisers could be targeting different ad positions depending on the exact phrase searched, which could introduce endogeneity into our estimates. In order to assess to what extent this might be a problem, we repeat our camera-keyword analysis on the smaller set of impressions generated purely from "exact brand name searches"—these searches, by definition, are exactly for "Nikon," "Canon," and "Olympus." Naturally, the smaller number of observations reduces the precision of our estimates—the reason we pooled in the first place. We are able to obtain statistically precise estimates of the Alexa interaction term only for "Canon"—the most

[27] One caveat is that in the new specification it is hard to estimate advertiser-time fixed effects for some advertisers with little clicks within 2-week windows. However, we expect this to have a negligible effect on our results because (a) the advertisers whose fixed effects cannot be estimated generally have low prominence in Alexa, and (b) there are very few of them.

Dep. Var.	Top-100 Alexa dummy Click	Click	Top-500 Alexa dummy Click	Click	Reciprocal of Alexa Click
Pos. 1	0.053** (0.004)	0.051** (0.004)	0.056** (0.004)	0.052** (0.004)	0.052** (0.004)
Pos. 2	0.031** (0.003)	0.032** (0.003)	0.031** (0.003)	0.033** (0.003)	0.032** (0.003)
Pos. 3	0.017** (0.003)	0.018** (0.003)	0.017** (0.003)	0.019** (0.003)	0.018** (0.003)
Pos. 4	0.000 (0.002)	0.000 (0.002)	0.000 (0.002)	0.000 (0.002)	0.000 (0.002)
Pos. 5	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)
Pos. 1×Top Alexa	-0.018* (0.010)	-	-0.023** (0.008)	-	-
Pos. 2×Top Alexa	0.010 (0.010)	-	0.008 (0.007)	-	-
Pos. 3×Top Alexa	0.013 (0.008)	-	0.005 (0.006)	-	-
Pos. 1-3×Top Alexa	-	0.003 (0.006)	-	-0.002 (0.004)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0057 (0.0088)
N	39016	39016	39016	39016	39004
R^2	0.037	0.037	0.037	0.037	0.037

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 9: Linear probability model predicting clicks for the Nikon keyword with time fixed effects.

popular of the three camera brands. For "Nikon," the analysis is inconclusive because of wide confidence intervals,[28] and for "Olympus" it is simply infeasible because of multicollinearity.

The results for "Canon" are presented in Table 12. We would like to stress two regularities. First, despite the lower power, we find evidence of negative interaction between

[28] Even though impressions matching "Nikon" exactly do not deliver precise estimates about the interaction between ad position and Alexa rank, we are still able to estimate the main position effects precisely. We find that these estimates are quantitatively close to those reported in Table 3.

	Top 100-Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.041** (0.002)	0.040** (0.002)	0.042** (0.002)	0.040** (0.002)	0.040** (0.002)
Pos. 2	0.028** (0.002)	0.028** (0.002)	0.029** (0.002)	0.028** (0.002)	0.028** (0.002)
Pos. 3	0.018** (0.002)	0.019** (0.002)	0.018** (0.002)	0.020** (0.002)	0.019** (0.002)
Pos. 4	0.006** (0.001)	0.006** (0.001)	0.006** (0.001)	0.006** (0.001)	0.006** (0.001)
Pos. 5	0.002** (0.001)	0.002** (0.001)	0.002** (0.001)	0.002** (0.001)	0.003** (0.001)
Pos. 1×Top Alexa	-0.010** (0.004)	-	-0.009** (0.004)	-	-
Pos. 2×Top Alexa	-0.006* (0.004)	-	-0.006* (0.003)	-	-
Pos. 3×Top Alexa	0.000 (0.004)	-	0.002 (0.003)	-	-
Pos. 1-3×Top Alexa	-	-0.006** (0.002)	-	-0.004* (0.002)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0067 (0.0061)
N	105427	105427	105427	105427	101229
R^2	0.023	0.023	0.023	0.023	0.023

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 10: Linear probability model predicting clicks for the Canon keyword with time fixed effects.

Alexa prominence and ad position. Second, the main position effects recovered using this subsample are close to those recovered in the pooled analysis, as reported in Table 4. These results suggest that possible endogeneity coming from the auxiliary words beyond the brand name does not pollute our main conclusions.[29]

[29]We repeated this "exact keyword" analysis on some of the non-branded, broader keywords used by Jeziorski and Segal (2015), namely "games," "white pages," and "weather." These results are reported in the On-line Appendix. Again, the interaction between ad position and Alexa prominence is negative. We acknowledge that these results are not as clean as the digital camera results, since the above keywords may

	Top-100 Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.043** (0.005)	0.044** (0.005)	0.048** (0.005)	0.047** (0.005)	0.044** (0.005)
Pos. 2	0.026** (0.004)	0.025** (0.004)	0.027** (0.005)	0.028** (0.004)	0.025** (0.004)
Pos. 3	0.016** (0.004)	0.014** (0.004)	0.016** (0.004)	0.017** (0.004)	0.014** (0.003)
Pos. 4	0.005* (0.003)	0.005* (0.003)	0.005* (0.003)	0.005* (0.003)	0.005* (0.003)
Pos. 5	0.004 (0.003)	0.004 (0.003)	0.004 (0.003)	0.004 (0.003)	0.004 (0.003)
Pos. 1×Top Alexa	0.008 (0.014)	-	-0.016* (0.009)	-	-
Pos. 2×Top Alexa	-0.008 (0.009)	-	-0.009 (0.007)	-	-
Pos. 3×Top Alexa	-0.014** (0.007)	-	-0.008 (0.006)	-	-
Pos. 1-3×Top Alexa	-	-0.006 (0.006)	-	-0.011** (0.004)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0151 (0.0106)
N	18608	18608	18608	18608	18177
R^2	0.027	0.027	0.027	0.027	0.025

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 11: Linear probability model predicting clicks for the Olympus keyword with time fixed effects.

In short, we believe our estimates of ad position effects, and the interaction between ad position prominence and advertiser prominence, are unbiased estimates of the underlying parameters, driven by the short-term, exogenous variation in ad positions that each advertiser experiences in our data. In particular, the finding that advertiser prominence and ad position prominence interact negatively in determining CTR is robust.

have some overlap between organic and sponsored search.

	Top-100 Alexa dummy		Top-500 Alexa dummy		Reciprocal of Alexa
Dep. Var.	Click	Click	Click	Click	Click
Pos. 1	0.033** (0.006)	0.031** (0.005)	0.035** (0.006)	0.033** (0.005)	0.034** (0.005)
Pos. 2	0.013** (0.006)	0.009** (0.005)	0.014** (0.006)	0.011** (0.005)	0.011** (0.005)
Pos. 3	-0.001 (0.004)	0.001 (0.003)	-0.001 (0.004)	0.002 (0.004)	0.001 (0.003)
Pos. 4	-0.002 (0.002)	-0.002 (0.002)	-0.002 (0.002)	-0.001 (0.002)	-0.002 (0.002)
Pos. 5	-0.000 (0.002)	-0.000 (0.002)	-0.000 (0.002)	0.000 (0.002)	0.001 (0.002)
Pos. 1×Top Alexa	-0.009 (0.009)	-	-0.015* (0.009)	-	-
Pos. 2×Top Alexa	-0.012* (0.007)	-	-0.017** (0.007)	-	-
Pos. 3×Top Alexa	0.002 (0.005)	-	-0.001 (0.005)	-	-
Pos. 1-3×Top Alexa	-	-0.005 (0.003)	-	-0.009** (0.003)	-
Pos. 1×(Alexa rank)$^{-1}$	-	-	-	-	-0.0239** (0.0076)
N	12592	12592	12592	12592	11361
R^2	0.016	0.016	0.017	0.017	0.017

Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$.

Model includes advertiser fixed effects, main effects for ad position (Pos.), and its interaction with advertiser prominence (Top Alexa), represented as Top-100 Alexa rank in columns I-II, as Top-500 Alexa rank in columns III-IV, and as inverse Alexa rank in column V.

Table 12: Linear probability model predicting clicks for the exact match of Canon keyword.

6 Discussion and conclusion

Among all advertising media, no medium is as single-mindedly focused on ad position as online search advertising. The auction mechanism underlying search advertising exists solely to place advertisements into slots and to price those slots. This concern for placement is not entirely misguided. Compared to TV and print advertising, search ads have much less to work with. They have limited room, and their content is sparse. If an advertiser deigns to make a factual claim, other relevant facts must necessarily be missing. On the emotional front, the format is nowhere near as engaging as TV and print ads. The bottom line is that search ads are not meant to brand-build, or to close a sale, but instead to persuade consumers to click on an ad so that the advertiser's website itself can perform the task of brand-building or closing the sale. In this context, being positioned high on a list of search ads understandably matters.

The contribution of this paper is to show that advertiser brand also matters. And it matters in ways not previously recognized in the literature. Using individual-level click-through data of consumers responding to search ads on Microsoft's Live Search platform, and web-traffic data from Alexa.com, we have shown that ad position matters more for less prominent retailers than for more prominent retailers. Specifically, in searches for camera brands, a retailer not in the Top-100 of Alexa rankings has a 30–50% higher click-through-rate (CTR) in position one than in position two, whereas a retailer in the Top-100 of Alexa rankings has only a 0–13% higher CTR for the same position improvement. In other words, ad position and advertiser prominence are substitutes as far as click-through is concerned, not complements, as assumed in the multiplicative models of Aggarwal et al. (2006), Katona and Sarvary (2010), Pin and Key (2011), and Nekipelov (2014). This explains why consumers do not always click from top to bottom, and why particular advertisers in lower positions generate more clicks than other advertisers in higher positions. Thus, our results are consistent with the "position paradox" discussed by Jerath et al. (2011), and the effects of advertiser "bigness" and "quality scores" noted by Narayanan and Kalyanam (2015). Most importantly, they provide a way for advertisers with different brand endowments to evaluate the marginal value of different ad positions appropriately, and to bid accordingly.

From a theoretical point of view, the attractiveness of the generalized second-price auction (GSP) underlying most search engines is predicated on the result that at least one of its "locally envy-free" equilibria provides the same payoffs to advertisers as the dominant strategy equilibrium in the Vickrey-Clark-Groves (VCG) mechanism (Edelman et al. 2007). However, Edelman et al. (2007) derive this result under the assumption "that all advertisers are identical along dimensions other than per-click value." As they note (in their footnote 6):

> The analysis would have to change considerably if there were specific advertiser-position effects. The magnitude of these advertiser-position effects is ultimately an empirical question, and we do not have the kind of data that would allow us to answer it; however, judging from the fact that the two major search engines effectively ignore it in their mechanisms (Yahoo! ignores CTRs altogether; Google computes an advertiser's estimated CTR conditional on the advertiser attaining the first position), we believe it to be small.

We know now that specific advertiser-position effects are present, and, indeed, are nontrivial. With this, even the notion of efficiency changes. To illustrate, consider two advertisers A and B with values per click $v_A > v_B$, competing for two advertising slots. Under the Edelman et al. (2007) assumptions, efficiency calls for allocating slot 1 to A and slot 2 to B, i.e., the allocation AB. However, if B's CTR in slot 1 greater than A's CTR in slot 1, while their CTRs in slot 2 are about the same, it may turn out that $\mathrm{CTR}_{\mathrm{B\ in\ BA}} \times v_B + \mathrm{CTR}_{\mathrm{A\ in\ BA}} \times v_A > \mathrm{CTR}_{\mathrm{B\ in\ AB}} \times v_B + \mathrm{CTR}_{\mathrm{A\ in\ AB}} \times v_A$, making the allocation BA more efficient. In other words, our empirical results have implications for the efficient allocation itself. Little is known about the optimality of generalized second-price auctions when advertiser prominence and ad position prominence interact as substitutes. The "quality score"-weighted auctions conducted by search engines, however, seem to treat advertiser prominence and ad position prominence as complements.[30] As Hsieh et al. (2015) note, they are therefore likely to reward prominent advertisers with discounts while their less prominent brethren are implicitly penalized, hardly a prescription for "envy-freeness" when the latter actually benefit more from being placed higher. Since envy-freeness is a common equilibrium refinement in the theory of GSP auctions, the violation of this property may be consequential for the results in the literature. In particular, the relative performance of GSP auctions vis-á-vis the VCG mechanism, in empirically relevant settings, becomes an open question.

Taken together, our results point to the similarities and differences between search advertising and traditional advertising. It is well understood that all ads must gain attention to be effective. So it is not surprising that ad placement and advertiser prominence matter in search advertising—they also matter in TV and print advertising (see, e.g., Goldberg and Hartwick

[30] "Quality scores," we conjecture, reflect to some degree advertiser prominence.

1990, Pieters and Wedel 2004). What this study shows is that the interaction effects between the two are also similar across media: just as a front-page ad or a TV commercial in the Super Bowl is more valuable to an up-and-coming-brand than to an established brand, a higher ad position is more valuable to a less prominent search advertiser than to a more prominent search advertiser. The reason may be as simple as saturation effects—an already-prominent advertiser can only increase its prominence so much via ad position. Alternatively, it may be a case of advertiser prominence overwhelming position prominence when it is large. Where search advertising likely differs from traditional advertising is in the relative importance of ad content versus ad position. Traditional media, because they allow much bigger sway to advertising content—because the have greater bandwidth and because they are inherently more engaging—have a greater capacity to build brands from scratch, and to overcome any brand deficiencies the advertiser comes endowed with. Search ads, because they do not build brands by themselves, must leverage their existing brand endowments more.

References

Agarwal, A., K. Hosanagar, and M. D. Smith (2011): "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets," *Journal of Marketing Research*, 48, 1057–1073.

——— (2015): "Do organic results help or hurt sponsored search performance?" *Information Systems Research*, 26, 695–713.

Aggarwal, G., A. Goel, and R. Motwani (2006): "Truthful auctions for pricing search keywords," in *Proceedings of the 7th ACM conference on Electronic commerce*, ACM, 1–7.

Animesh, A., V. Ramachandran, and S. Viswanathan (2010): "Research Note: Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation," *Information Systems Research*, 21, 190–201.

Animesh, A., S. Viswanathan, and R. Agarwal (2011): "Competing "Creatively" in Sponsored Search Markets: The Effect of Rank, Differentiation Strategy, and Competition on Performance," *Information Systems Research*, 22, 153–169.

Arkhangelsky, D., S. Izmalkov, and D. Khakimova (2013): "On evaluation of ctrs of different positions in sponsored search auctions," in *14th ACM Conference on Electronic Commerce, poster*.

Athey, S. and D. Nekipelov (2010): "A structural model of sponsored search advertising auctions," in *Sixth ad auctions workshop*.

Athey, S. C., A. Schwaighofer, and D. Nekipelov (2014): "Tool for analysis of advertising auctions," US Patent 8,650,084.

Blake, T., C. Nosko, and S. Tadelis (2015): "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment," *Econometrica*, 83, 155–174.

Brooks (2006): "The Atlas rank report: How search engine rank impacts traffic," Tech. rep., Atlas Institute.

Bryniolfsson, E. and M. D. Smith (2000): "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 46, 563–585.

Buscher, G., S. T. Dumais, and E. Cutrell (2010): "The Good, the Bad, and the Random: An Eye-tracking Study of Ad Quality in Web Search," in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM, SIGIR '10, 42–49.

Chen, Y. and C. He (2011): "Paid Placement: Advertising and Search on the Internet*," *The Economic Journal*, 121, F309–F328.

Clogg, C. C., E. Petkova, and A. Haritou (1995): "Statistical Methods for Comparing Regression Coefficients Between Models," *American Journal of Sociology*, 100, 1261–1293.

Craswell, N., O. Zoeter, M. Taylor, and B. Ramsey (2008): "An experimental comparison of click position-bias models," in *Proceedings of the 2008 International Conference on Web Search and Data Mining*, ACM, 87–94.

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

Fallows, D. (2005): "Search Engine Users," *Pew Research Center: Internet, Science & Tech*.

Ghose, A. and S. Yang (2009): "An empirical analysis of search engine advertising: Sponsored search in electronic markets," *Management Science*, 55, 1605–1622.

Goldberg, M. E. and J. Hartwick (1990): "The effects of advertiser reputation and extremity of advertising claim on advertising effectiveness," *Journal of Consumer Research*, 172–179.

Goldman, M. and J. M. Rao (2014): "Experiments as Instruments: Heterogeneous Position Effects in Sponsored Search Auctions," Working paper, University of California, San Diego.

Gomes, R., N. Immorlica, and E. Markakis (2009): "Externalities in keyword auctions: An empirical and theoretical assessment," in *International Workshop on Internet and Network Economics*, Springer, 172–183.

Hsieh, Y.-W., M. Shum, and S. Yang (2015): "To Score or Not to Score? Estimates of a Sponsored Search Auctions Model," Working Paper, USC.

Jerath, K., L. Ma, Y.-H. Park, and K. Srinivasan (2011): "A "position paradox" in sponsored search auctions," *Marketing Science*, 30, 612–627.

Jeziorski, P. and I. Segal (2015): "What makes them click: Empirical analysis of consumer demand for search advertising," *American Economic Journal: Microeconomics*, 7, 24–53.

Karande, C., A. Mehta, and R. Srikant (2013): "Optimizing Budget Constrained Spend in Search Advertising," in *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, ACM, WSDM '13, 697–706.

Katona, Z. and M. Sarvary (2010): "The race for sponsored links: Bidding patterns for search advertising," *Marketing Science*, 29, 199–215.

Narayanan, S. and K. Kalyanam (2015): "Position Effects in Search Advertising and Their Moderators: A Regression Discontinuity Approach," *Marketing Science*, 34, 388–407.

Nekipelov, D. (2014): "Eliciting preferences of sponsored search advertisers: implications for mechanism design," *ACM SIGecom Exchanges*, 13, 72–76.

Pham, M. T. and G. V. Johar (2001): "Market prominence biases in sponsor identification: processes and consequentiality," *Psychology and Marketing*, 18, 123–143.

Pieters, R. and M. Wedel (2004): "Attention capture and transfer in advertising: Brand, pictorial, and text-size effects," *Journal of Marketing*, 68, 36–50.

Pin, F. and P. Key (2011): "Stochastic Variability in Sponsored Search Auctions: Observations and Models," in *Proceedings of the 12th ACM Conference on Electronic Commerce*, ACM, EC '11, 61–70.

Richardson, M., E. Dominowska, and R. Ragno (2007): "Predicting Clicks: Estimating the Click-through Rate for New Ads," in *Proceedings of the 16th International Conference on World Wide Web*, New York, NY, USA: ACM, WWW '07, 521–530.

Rutz, O. J., R. E. Bucklin, and G. P. Sonnier (2012): "A latent instrumental variables approach to modeling keyword conversion in paid search advertising," *Journal of Marketing Research*, 49, 306–319.

Simonov, A., C. Nosko, and J. M. Rao (2015): "Competition and Crowd-Out for Brand Keywords in Sponsored Search," *Available at SSRN 2668265*.

Ursu, R. M. (2015): "The Power of Rankings: Quantifying the Effects of Rankings on Online Consumer Search and Choice," *Available at SSRN 2729325*.

Varian, H. R. (2007): "Position auctions," *International Journal of Industrial Organization*, 25, 1163–1178.

Yang, S. and A. Ghose (2010): "Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?" *Marketing Science*, 29, 602–623.