

Advertiser prominence effects in search advertising*

Przemysław Jeziorski

University of California, Berkeley

Sridhar Moorthy

University of Toronto

May 27, 2015

*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—Matt 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, 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 advertising that appears at the top and right-hand sides of the search page when a user searches for something in an online search engine. Search ads feature two types of prominence: prominence of ad position and 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 Live Search, and measures of advertiser prominence from Alexa.com, we find that ad position and advertiser prominence are substitutes, not complements. Specifically, in searches for digital 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 in the top position. Qualitatively similar results are obtained for several other search strings. These findings demonstrate, first, that advertiser identity 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 Google. The ads—sponsored links, as they are also called—appear at the top and right-hand side of the “results page” in an ordered list (see Figure 1). This form of advertising is relatively new; it didn’t exist until 1998. However, already by 2003, it was the largest online ad medium, which it continues to be with nearly 40% of the market.¹

One of the distinguishing features of search ads is that the ads themselves are fairly uninteresting objects—certainly nowhere near as rich as TV and print ads. Other than identifying the advertiser and providing a clickable link to the advertiser’s website, they are relatively contentless.² To the extent they work, they do so by taking the consumer to the advertiser’s website, where the real selling takes place. The pricing of search ads reflects this: advertisers pay for clicks, not for exposure. In fact, a key term of art in search advertising is click-through-rate (CTR), the fraction of exposures that result in a click.

This paper is about the determinants of CTR. Obviously, ad position is implicated: we expect ads in higher positions to generate more clicks than ads in lower positions. But what is the marginal value of ad position—how much higher will CTR be if an advertiser moves up a position? Does marginal value of ad position depend on ad position? Does advertiser matter? If so, how does advertiser prominence interact with ad position?

The literature in economics and marketing has noted the importance of these questions, but empirical answers remain elusive. Edelman et al. (2007) and Varian’s (2007) theoretical models simply assume that the CTR decreases from top to bottom, independent of advertiser identity. More common is the assumption that CTR depends both on ad position and adver-

¹Market share data from IAB Internet Advertising Revenue Reports April 2003 and April 2014. Oremus (2013) notes that a search engine called GoTo.com (later Overture) was the first to offer paid-search ads. Growth accelerated in 2002 when Google launched its own pay-per-click, auction-based search-advertising product called AdWords Select.

²Not literally, of course—the ads do have words in them. Our point is that other than the identity of the advertiser, they generally do not offer much credible information. Perhaps for this reason, ad copy—whether it is quality oriented or price oriented—has been found to have no influence on the click-through rate of a search ad (Animesh et al. 2011).

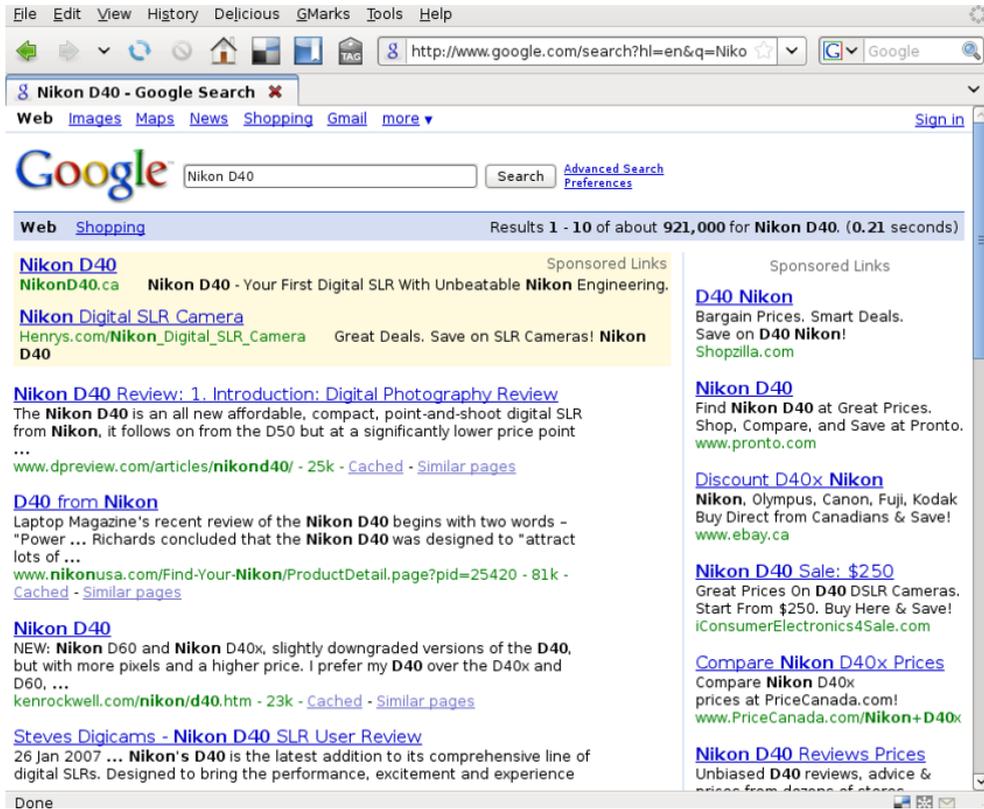


Figure 1: Search advertising

tiser prominence, multiplicatively. That is, $CTR_{ij} = \alpha_j \beta_i$, where α_j is an ad position factor and β_i is an advertiser prominence factor (Aggarwal et al. 2006, Katona and Sarvary 2010, Pin and Key 2011, Nekipelov 2014). While there is empirical support for higher ad positions eliciting higher CTR in the aggregate (Brooks 2006, Animesh et al. 2011, Rutz et al. 2012), disaggregate analysis shows considerable variation around this pattern. For example, Jerath et al. (2011) and Narayanan and Kalyanam (2014) demonstrate that particular advertisers in lower positions get higher CTRs than advertisers above them. Jeziorski and Segal (2014) find that the CTR for a particular advertiser in a particular ad position depends on who else is advertising and in what positions.

In this paper we combine individual-level clickstream data from Microsoft Live Search and advertiser prominence data from Alexa.com to study how ad position and advertiser prominence affect CTR. Consumers in our data are searching for particular brands of products, e.g., Nikon. Advertisers to such keywords are typically retailers.

We find that in nearly two-thirds of the impressions with clicks, consumers click non-sequentially; a click on the j th ad is generally *not* the j th click. Furthermore, in nearly a third of the impressions with two or more clicks, the consumer changes direction: she proceeds upward after proceeding downward. These observations suggest the importance of advertiser identity and of context—who the advertiser is and who else is advertising, both affect consumer behavior toward search ads. Our subsequent results pin down more precisely how these things matter. We find that the multiplicative assumption underlying much of the theoretical literature completely mischaracterizes the nature of the interaction between ad position and advertiser prominence. One of the implications of the multiplicative assumption is that the marginal CTR-value of ad position, $(\alpha_j - \alpha_{j+1})\beta_i$, is increasing in advertiser prominence. Instead, we find just the opposite: more prominent advertisers benefit less from improvements in ad position than less prominent advertisers. The multiplicative assumption also implies that *relative* marginal value of ad position, $(CTR_{ij} - CTR_{ij+1})/CTR_{ij+1}$, is independent of advertiser prominence. We find that advertiser prominence matters both in an absolute sense as well as in a relative sense. This last result implies, in particular, that the elasticity of CTR with respect to ad position is decreasing in advertiser prominence.

2 Data and institutional details

When a user submits a search query (“search string”) with commercial value to a search engine, a list of ads appears, what we call an “impression.” Sponsored links appear at the top and right-hand side of the so-called organic links—the 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). The search ads in an impression are determined by a generalized second-price auction (Edelman et al. 2007). Businesses 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 by proprietary “quality scores” and runs an auction, the outcome of which is an ordering of ads and a price-per-click for each advertising slot (Edelman et al. 2007, Varian 2007). Quality scores and bids are keyword- and advertiser/ad-specific. The search

engine is compensated only if a consumer clicks on a sponsored ad.³

Our search ads data come from Microsoft’s search engine, *Live Search*, the precursor to today’s Bing. In 2008, as part of its Beyond Search initiative, Microsoft made available to a limited set of academics a data set containing 20 million search impressions chosen randomly from the ones that appeared in a roughly three-month period, 10 August-1 November 2007.⁴ 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. Impressions that were part of longer user sessions thus had a proportionally 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, both in impressions and in users, is essentially random. 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 keyword that generated it, list of sponsored links, order of sponsored links, identity of advertisers, and the time stamps of all clicks on sponsored links in a session. Thus, this data set is fairly unique. Other than Jeziorski and Segal (2014), who use the same data source as us, we are not aware of any empirical study that uses *impression*-level data, with *all* advertisers in an impression accounted for. By contrast, Ghose and Yang (2009) and Yang and Ghose (2010) use weekly average data, and Animesh et al. (2011) and Narayanan and Kalyanam (2014) use daily average data, aggregating over many impressions in the process. Furthermore, the first three papers are based on data from a single advertiser, whereas the fourth is based on data from four advertisers that eventually merged into the same firm.

We restrict our attention to keywords that have a large number of impressions and are specific enough to involve the same set of advertisers in a large proportion of the impressions. Such keywords tend to be well-known brands that participate in a fairly narrow set of related product categories, e.g., Nikon. Advertisers advertising to such keywords tend to be retailers.⁵

³See Pin and Key (2011) and Nekipelov (2014) for more institutional details about search advertising.

⁴According to techcrunch.com, Live Search had 9.1% of the U.S. online search market in May 2008 (versus 61.6% for the market leader, Google). This translated to about 900 million search queries per month.

⁵Retailers are the largest single category of search advertisers, according to the Interactive Advertising

Our main analysis concerns three digital camera brands: Canon, Nikon, and Olympus. In addition, 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 digital camera brands and 252,138 search impressions for all the other keywords combined.⁶

We supplement the Microsoft data with data from Alexa.com. Alexa ranks websites by their traffic using a “global traffic panel, which is a sample of millions of Internet users using one of over 25,000 different browser extensions,” and, directly, for “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 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 retailer during our sample period. Using the daily data we computed a 3-month average Alexa rank for each retailer and used that as a measure of advertiser prominence in our analysis.⁷

3 Analysis

Tables 1-3 contain descriptive statistics from our data set.

Bureau, accounting for almost 20% of advertising spending in 2012.

⁶Our analysis is keyword-specific, which essentially means that we are controlling for “unobserved search string effects” by running separate regressions for each keyword. For branded keywords, to the extent the search string involved words other than the brand and the advertisers targeted different ad positions based on these unobserved words, our results may be non-robust. This consideration dictates our choice of the digital camera brands as the primary object of our analysis. The analysis of the broader keywords should be regarded as complementary, because it provides evidence for the external validity of our findings beyond the digital camera market.

⁷The variation in retailers’ Alexa ranks within the 3-month period of our sample is negligible, suggesting that these ranks are relatively stable characteristics of advertisers. For other studies using market share like measures of advertiser prominence, see Goldberg and Hartwick (1990), Brynjolfsson and Smith (2000), Pham and Johar (2001), and Agarwal et al. (2012).

	Digital cameras	Broader 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 below slots	1.07%	0.74%

Table 1: Descriptive statistics

Table 1 contains basic facts for search impressions corresponding to digital camera keywords and the broader keywords. 13% of digital camera impressions have at least one click on a search ad and 2.2% of impressions have two or more clicks on search ads.⁸ 63% of the impressions with at least one click are impressions with non-sequential clicks.⁹ In other words, most ad impressions 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

⁸The first of these numbers is higher than the typical CTR percentages 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 2014).

⁹As 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.

	Keyword		
	Nikon	Canon	Olympus
Advertiser 1	5.9%	14.0%	3.0%
Advertiser 2	10.7%	2.6%	4.9%
Advertiser 3	11.5%	2.4%	4.4%
Advertiser 4	3.1%	3.4%	4.1%
Other advertisers	2.6%	2.4%	2.3%

Table 2: Click-through-rates of the top-4 most clicked advertisers for each camera brand. Note that the advertisers are not necessarily the same across keywords (cf. footnote 10).

decrease is rather steep with fourth and below slots receiving one-fifth of the CTR of the top slot. This suggests either that high CTR is a natural property of high ad positions or that the top ad positions attract more prominent advertisers who attract more clicks.

Corresponding descriptive statistics for advertisers in digital camera keywords are presented in Tables 2-3 (we omit the statistics for broader keywords for brevity). In these tables we show the top-4 advertisers (ranking is separate for each keyword) with the most clicks; the rest are pooled into a catch-all “other advertisers” category.¹⁰ Table 2 shows the CTRs of the chosen advertisers sorted by total number of clicks. Note that the CTRs are not monotonic. This is because some advertisers have more impressions than others and hence generate more clicks even with a lower CTR. Table 3 presents the average Alexa ranks of the four most-clicked advertisers and “other advertisers.” We note that in general the most-clicked advertisers have higher Alexa ranks than “other advertisers.” However, amongst the most-clicked advertisers, we do observe advertisers with low Alexa ranks, such as advertiser 3 for Nikon (who is also advertiser 1 for Canon). Additionally, note that 7–9 percent of the low-CTR “other advertisers” have a Top-100 Alexa ranking, which places them amongst the most recognizable advertisers on the Internet.

In order to disentangle the effect of ad position and the effect of advertiser prominence

¹⁰As might be expected, the top-4 advertisers are often, but not always, the same across camera brands—after all, these are camera retailers with similar brand assortments. Since we do not pool impressions across camera brands, however, we treat the advertisers for each camera brand as distinct.

	Keyword		
	Nikon	Canon	Olympus
Advertiser 1	14,650	66,379	24
Advertiser 2	515	24	1117
Advertiser 3	66,379	98	16,455
Advertiser 4	24	540	540
Other advertisers	314,537	255,582	253,532
% of “other advertisers” ranked in Alexa Top-100	7	6	9

Table 3: Average Alexa ranks of retailers by search string (higher ranks correspond to less prominent advertisers). Note that the advertisers are not necessarily the same across search strings (cf. footnote 10).

on CTR, we estimate a linear probability model for each keyword. The dependent variable is a click, and explanatory variables include advertiser fixed effects, ad position fixed effects, and interactions between ad position and advertiser prominence. The latter is represented in three different ways: by a dummy variable identifying Top-100 Alexa-ranked advertisers (Alexa rank $\leq 100 = 1$, Alexa rank $> 100 = 0$), a dummy variable identifying Top-500 Alexa-ranked advertisers, and as a continuous variable, as log of the reciprocal of the Alexa rank.

In Tables 4-6 we present the results of various regressions of this form for the keywords, Nikon, Canon and Olympus, respectively. The top of each table confirms the presence of strong ad position effects even after controlling for advertiser via fixed 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 Olympus, and beyond position 4 for Canon.

		Nikon						
Ad position	I	II	III	IV	V	VI	VII	
1	0.054** (0.003)	0.053** (0.003)	0.051** (0.003)	0.056** (0.003)	0.055** (0.004)	0.051** (0.003)	0.051** (0.003)	
2	0.032** (0.003)	0.031** (0.003)	0.031** (0.003)	0.032** (0.003)	0.030** (0.003)	0.032** (0.003)	0.032** (0.003)	
3	0.017** (0.003)	0.016** (0.003)	0.017** (0.003)	0.017** (0.003)	0.016** (0.003)	0.017** (0.003)	0.017** (0.003)	
4	-0.001 (0.003)	-0.002 (0.003)	-0.002 (0.003)	-0.001 (0.003)	-0.001 (0.003)	-0.002 (0.003)	-0.002 (0.003)	
5	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	0.003 (0.003)	
1×Top Alexa	-0.023** (0.008)	-0.020** (0.009)	-	-0.025** (0.007)	-0.023** (0.007)	-	-	
2×Top Alexa	-	0.010 (0.008)	-	-	0.010 (0.007)	-	-	
3×Top Alexa	-	0.013 (0.008)	-	-	0.006 (0.006)	-	-	
1-3×Top Alexa	-	-	0.002 (0.006)	-	-	-0.001 (0.005)	-	
$1 \times a^{-1}$	-	-	-	-	-	-	-0.0048 (0.0134)	
N	39016	39016	39016	39016	39016	39016	39004	
R^2	0.030	0.030	0.029	0.030	0.030	0.029	0.029	

Table 4: Results from a linear probability model that predicts clicks as a function of ad position for the Nikon keyword. The model includes interactions between ad position and advertiser prominence (Alexa rank) and controls for advertiser heterogeneity via fixed effects (not reported). Specifically, specifications I, II and III, contain interactions between, respectively, the highest ad position, ad positions 1-3 individually, and ad positions 1-3 collectively, and a dummy that captures a Top-100 Alexa-ranked advertiser. Columns IV, V and VI contain the same analysis but with Top-500 Alexa dummy instead of Top-100 Alexa dummy. Column VII contains a specification in which position 1 is interacted with the logarithm of the inverse Alexa score.

Ad position	Canon						
	I	II	III	IV	V	VI	VII
1	0.040** (0.002)	0.040** (0.002)	0.039** (0.002)	0.041** (0.002)	0.041** (0.002)	0.039** (0.002)	0.039** (0.002)
2	0.027** (0.002)	0.027** (0.002)	0.027** (0.002)	0.027** (0.002)	0.028** (0.002)	0.028** (0.002)	0.027** (0.002)
3	0.018** (0.002)	0.018** (0.002)	0.019** (0.002)	0.018** (0.002)	0.017** (0.002)	0.019** (0.002)	0.019** (0.002)
4	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005** (0.002)
5	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)
1×Top Alexa	-0.008** (0.003)	-0.009** (0.004)	-	-0.007** (0.003)	-0.008** (0.003)	-	-
2×Top Alexa	-	-0.005 (0.004)	-	-	-0.005 (0.003)	-	-
3×Top Alexa	-	0.001 (0.004)	-	-	0.002 (0.003)	-	-
1-3×Top Alexa	-	-	-0.005* (0.003)	-	-	-0.003 (0.002)	-
$1 \times a^{-1}$	-	-	-	-	-	-	-0.0029 (0.0075)
N	105427	105427	105427	105427	105427	105427	101229
R^2	0.022	0.022	0.022	0.022	0.022	0.022	0.022

Table 5: Results from a linear probability model that predicts clicks as a function of ad position for the ‘Canon keyword. The model includes interactions between ad position and advertiser prominence (Alexa rank) and controls for advertiser heterogeneity via fixed effects (not reported). Specifically, specifications I, II and III, contain interactions between, respectively, the highest ad position, ad positions 1-3 individually, and ad positions 1-3 collectively, and a dummy that captures a Top-100 Alexa-ranked advertiser. Columns IV, V and VI contain the same analysis but with Top-500 Alexa dummy instead of Top-100 Alexa dummy. Column VII contains a specification in which position 1 is interacted with a logarithm of the inverse Alexa score.

Olympus							
Ad position	I	II	III	IV	V	VI	VII
1	0.043** (0.004)	0.044** (0.004)	0.045** (0.004)	0.048** (0.004)	0.048** (0.005)	0.047** (0.004)	0.045** (0.004)
2	0.024** (0.004)	0.025** (0.004)	0.025** (0.004)	0.024** (0.004)	0.026** (0.004)	0.027** (0.004)	0.025** (0.004)
3	0.013** (0.004)	0.015** (0.004)	0.014** (0.004)	0.013** (0.004)	0.015** (0.004)	0.016** (0.004)	0.014** (0.004)
4	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)	0.006 (0.004)
5	0.003 (0.004)	0.003 (0.004)	0.003 (0.004)	0.003 (0.004)	0.003 (0.004)	0.003 (0.004)	0.004 (0.004)
1×Top Alexa	0.012 (0.010)	0.008 (0.010)	-	-0.012 (0.007)	-0.014* (0.008)	-	-
2×Top Alexa	-	-0.006 (0.009)	-	-	-0.006 (0.007)	-	-
3×Top Alexa	-	-0.014 (0.009)	-	-	-0.006 (0.007)	-	-
1-3×Top Alexa	-	-	-0.005 (0.006)	-	-	-0.009* (0.005)	-
$1 \times a^{-1}$	-	-	-	-	-	-	-0.0226 (0.0159)
N	18608	18608	18608	18608	18608	18608	18177
R^2	0.018	0.018	0.018	0.018	0.018	0.018	0.018

Table 6: Results from a linear probability model that predicts clicks as a function of ad position for the Olympus keyword. The model includes interactions between ad position and advertiser prominence (Alexa rank) and controls for advertiser heterogeneity via fixed effects (not reported). Specifically, specifications I, II and III, contain interactions between, respectively, the highest ad position, ad positions 1-3 individually, and ad positions 1-3 collectively, and a dummy that captures a Top-100 Alexa-ranked advertiser. Columns IV, V and VI contain the same analysis but with Top-500 Alexa dummy instead of Top-100 Alexa dummy. Column VII contains a specification in which position 1 is interacted with a logarithm of the inverse Alexa score.

The bottoms of these tables show the interaction between advertiser prominence and ad position. We estimate three specifications with these interactions and in each case find a statistically significant impact of advertiser prominence on the size of the position effect. Namely, for the Nikon keyword, being in the top-100 of the Alexa classification nullifies the marginal effect of the top slot (the difference between ad position 1 and ad position 2). In other words, retailers that enjoy a top-100 Alexa ranking obtain the same CTRs in ad positions 1 and 2, whereas, not top-100 retailers enjoy nearly 50% higher CTR in the top position.¹¹ The impact of advertiser prominence is similar for the Canon keyword. Here, not top-100 Alexa retailers enjoy 1.3-percentage point increase (approximately 33% increase) in CTR when moving from second position to the top position, whereas, top-100 Alexa retailers enjoy only 0.5-percentage point increase (approximately 13% increase). Lastly, for the Olympus keyword, sample sizes are significantly smaller, so it is harder to find statistically significant advertiser prominence effects. Still, even here, we find that being a top-500 Alexa-ranked advertiser nearly nullifies the effect of top ad position. There advertiser prominence also has a significant aggregate effect on the top-3 slots. In particular, top-500 Alexa retailers enjoy a 1-percentage point lower gap between the top-three slots and lower slots, which is equivalent to nullifying the positive impact of slot 3 and decreasing the positive impacts of slot 1 and 2 by 20% and 33%, respectively.

The analysis of broader keywords is summarized in Table 7. Each row is a separate regression replicating either column I or column IV in Tables 4, 5 and 6 depending on whether advertiser prominence is captured as a Top-100 ranking in Alexa or as a Top-500 ranking in Alexa. We report only the interaction coefficient summarizing how the top ad position interacts with advertiser prominence. As can be seen, the interaction effects are robust. Despite the greater breath of these keywords, in 5 of the 8 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 saw earlier for digital cameras.

¹¹This doesn't mean, of course, that top Alexa-ranked retailers get lower CTRs than lower Alexa-ranked retailers. As noted earlier, the model includes advertiser fixed effects. Top Alexa-ranked retailers tend to have larger or not statistically different fixed effects than low Alexa-ranked retailers.

Nike	Position 1 × Top 100	-0.039** (0.006)
	Position 1 × Top 500	-0.042** (0.005)
Adidas	Position 1 × Top 100	-0.011 (0.011)
	Position 1 × Top 500	-0.021** (0.010)
Puma	Position 1 × Top 100	-0.029** (0.013)
	Position 1 × Top 500	-0.034** (0.010)
Sony	Position 1 × Top 100	-0.022** (0.003)
	Position 1 × Top 500	-0.015** (0.002)
Yamaha	Position 1 × Top 100	0.004 (0.003)
	Position 1 × Top 500	0.002 (0.003)
Maytag	Position 1 × Top 100	-0.013 (0.010)
	Position 1 × Top 500	-0.008 (0.009)
Textbook	Position 1 × Top 100	0.008 (0.009)
	Position 1 × Top 500	0.002 (0.008)
Ebook	Position 1 × Top 100	-0.016** (0.001)
	Position 1 × Top 500	-0.015** (0.001)

Table 7: Interaction effects between the top ad position and advertiser prominence (dummy variables indicating Top-100 or Top-500 Alexa ranks) from a linear probability model predicting clicks as a function of ad position for each of the keywords listed in column 1. The model includes ad position and advertiser fixed effects (not reported). Each cell is a separate regression replicating columns I and IV of Tables 4, 5 and 6.

We find insignificant interaction between Alexa and position for Yamaha, Maytag and “textbooks”. We cannot rule out that these interaction terms are indeed absent, however,

we suspect that the negative result is related to some peculiarities of these keywords. For example, Yamaha is arguably the broadest keyword in our sample because it contains retailers of such diverse products as music equipment and motorcycles. These different sub-categories may have different position effects, which contributes noise and may even lead to bias in our estimates. Keyword Maytag has the smallest number of clicks in our data; we suspect that Maytag coefficients would be significant if more data were available. Lastly, “textbooks” is an unusual keyword because it elicits advertising from a retailer with an extremely suggestive .com name. This advertiser is essentially a monopolist and harvests nearly 50% of all clicks. This retailer has a low Alexa rank and not coincidentally enjoys a strong position effect, which is consistent with our main hypothesis. However, due to 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-position interaction term.

Collectively, it is clear that our results are inconsistent with the multiplicative model of $CTR_{ij} = \alpha_j \beta_i$, popular in the theoretical literature (e.g., Aggarwal et al. 2006, Katona and Sarvary 2010, Pin and Key 2011). They suggest, in fact, a model such as $CTR_{ij} = \alpha_j / \beta_i$.

4 Are ad positions exogenous?

A potential concern with our estimates is that ad positions are not exactly exogenous—that ad positions might be correlated with unobserved advertiser characteristics (which also affect CTR). However, there are limits to how concerning this concern should be: after all our regressions recognize observed advertiser heterogeneity via Alexa ranks and unobserved advertiser heterogeneity via fixed effects. Still, there might be a residual concern that advertiser fixed effects do not completely control for unobserved advertiser characteristics. The argument would be that unobserved advertiser characteristics are not exactly fixed—they might have a dynamic component. For example, a retailer’s inventory position in Nikon cameras might change over time, which prompts this retailer to be more or less aggressive in its pricing policy, which, in turn, prompts it to be more or less aggressive in seeking higher ad positions. But this counter-argument itself has its limits: it depends on the assumption that consumers recognize a correlation between ad position and underlying offers when making

Nikon				
	Swap within 5s	Swap within 1h	Swap within a day	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	43301	593877	39016
Canon				
	Swap within 5s	Swap within 1h	Swap within a day	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	105427	173204	2734901	105425
Olympus				
	Swap within 5s	Swap within 1h	Swap within a day	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	9028	99708	18608

Table 8: Regression discontinuity estimates of position effects in comparison with our fixed-effect estimator. We construct observation pairs for which the position of the same domain changes within a short interval. Then we first-difference the time-domain effects; that is, we regress the difference between the indicator functions for clicks on the difference between the indicator functions for respective positions. The approach is equivalent to using thousands of time-domain fixed effects.

clicking decisions. For example, if consumers simply reacted to the prominence of ad positions, without probing what those positions might be saying about the underlying “quality” of advertisers’ offers, then endogeneity of ad positions would still not be a concern.

Narayanan and Kalyanam (2014) discuss several mechanisms that can result in position endogeneity and mention, among them: (i) strategic interactions caused by sophisticated bidding and bid throttling software, and (ii) promotions conducted by the advertisers that may be associated with changes in their bidding strategy. However, according to anecdotal evidence, the behavior of advertisers in 2007 was quite different from contemporary advertiser behavior, specifically: (i) advertisers changed their bids quite infrequently, (ii) the targeting of ads (and bid) on user characteristics was not possible, and (iii) automated bidding software was not available. Also, 2007 was the first year of operation of the Microsoft Live marketplace. For this reason, the quality scores of advertisers are noisy, which results in a considerable amount of exogenous ad placement variation.

Fortunately, our copious data set allows us to test for the exogeneity of ad positions using a regression discontinuity design. This analysis is presented in Table 8. For each advertiser we construct observation pairs in which the advertiser is displayed in different positions within a short period of time ranging from 5 seconds to 12 hours. We construct observations by first-differencing the unobserved ad quality for each pair, and we reestimate the position effects using linear probability model. Formally, we perform the regression

$$c_{a,t} - c_{a,t-1} = (\iota_{a,t} - \iota_{a,t-1})\beta + \epsilon_{a,t} - \epsilon_{a,t-1},$$

where $c_{a,t}$ is the indicator function for a click on domain a in impression t , and $\iota_{a,t} = (0, \dots, 1, \dots)$ is a vector indicating position of ad a .¹² Under the mild assumption that the ad quality is constant within the short interval $[t - 1, t]$, the ad quality cancels out and the estimates of position effects, $\hat{\beta}$, are robust to position endogeneity. We compare the regression-discontinuity ad position estimates with the simple fixed-effect estimates without

¹²Note that the number of observations for the discontinuity regressions is frequently larger than for fixed effects regressions. It is caused by the fact that we pick all pairs of domains that switch positions, which combinatorially increases the size of the data. For this reason, one may rightfully worry about clustering of standard errors; nevertheless, we use unclustered (underestimated) standard errors, which provide the most conservative test for position endogeneity.

the interaction term with Alexa scores. The simple fixed-effects specification provides a conservative test of ad position endogeneity because it does not rely on Alexa scores to absorb some unobserved advertiser quality. As Table 8 shows, regression discontinuity estimates are not significantly different from the simple fixed-effect estimates and from our main estimates of position effects reported in Tables 4-6. For example, for the Nikon keyword, even among impressions as far apart as 5s, the top position effect is .055 which is close to the simple fixed-effect estimate of .050 (and even closer to the estimates in Table 4). Needless to say, impressions within 1h and 12h, which deliver more precise estimates, are still statistically indistinguishable from the simple fixed-effect estimates. Only in the case of Canon in the 5s window are the estimates substantially different from the simple fixed-effect estimates. Despite the large number of observations, the statistical precision of these estimates is low. This suggests that, in this particular case, there is not much variation in advertisers' positions. Consequently, even mild clustering of standard errors renders the difference between the 5s window and the fixed-effects estimates insignificant.

5 Discussion and conclusion

Among all the advertising media, no medium is as single-mindedly focused on ad placement 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 ad placement is not entirely misguided. Compared to TV and print advertising, search ads have much less to work with. Their room is limited and the content sparse: website address, some hyperbole, and not much else. If an advertiser deigns to make a factual claim, other relevant facts must necessarily be missing—and we are back in the realm of hyperbole. On the emotional front, the format is nowhere near as engaging as TV and print ads, despite recent improvements in this regard. 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 to search ads,

and that it matters in ways not previously recognized in the literature. Using individual-level click-through data of people responding to search ads on the Microsoft 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 digital 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. In comparison, a retailer in the top-100 of Alexa rankings has only a 0–13% higher CTR in the top position. In other words, ad position and advertiser prominence are substitutes as far as click-through rates are concerned, not complements, as the theoretical literature has commonly assumed. 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 corroborate the “position paradox” discussed by Jerath et al. (2011) and are consistent with the effects of “bigness” noted by Narayanan and Kalyanam (2014). Most importantly, they provide a way for advertisers with different brand endowments to evaluate the value of different ad positions.

Taken together these 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 a surprise that ad placement and advertiser prominence matter for search advertising—they also matter in TV and print advertising (see, e.g., Goldberg and Hartwick 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. Where search advertising likely differs from traditional media is in the relative importance of ad content versus ad position. Traditional media, because they allow much bigger sway to advertising content—because of bandwidth and because of engageability— have a greater capacity to build brands from scratch, and to overcome any brand deficiencies the advertiser comes endowed with. Search ads, because they don’t build brands by themselves, must leverage their brand endowments more.

References

- AGARWAL, A., K. HOSANAGAR, AND M. D. SMITH (2012): “Do Organic Results Help or Hurt Sponsored Search Performance?” Working Paper, University of Pennsylvania.
- 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., 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.
- BROOKS (2006): “The Atlas rank report: How search engine rank impacts traffic,” Tech. rep., Atlas Institute.
- BRYNJOLFSSON, E. AND M. D. SMITH (2000): “Frictionless Commerce? A Comparison of Internet and Conventional Retailers,” *Management Science*, 46, 563–585.
- 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.
- 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.
- 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 (2014): “What Makes them Click: Empirical Analysis of Consumer Demand for Search Advertising,” *American Economic Journal: Microeconomics*, forthcoming.
- 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 (2014): “Position Effects in Search Advertising: A Regression Discontinuity Approach,” Tech. rep., Stanford University.
- NEKIPELOV, D. (2014): “Eliciting Preferences of Sponsored Search Advertisers: Implications for Mechanism Design,” Tech. rep., University of Virginia.
- OREMUS, W. (2013): “Google’s Big Break,” *Slate*.
- 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.
- 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.
- 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.