

Competition and Crowd-Out for Brand Keywords in Sponsored Search

Andrey Simonov,^a Chris Nosko,^b Justin M. Rao^c

^a Columbia University, New York, New York 10027; ^b Amazon, Inc., Seattle, Washington 98109 ^c HomeAway, Inc., Austin, Texas 78703

Contact: asimonov@gsb.columbia.edu, <http://orcid.org/0000-0002-3238-9829> (AS); cnosko@gmail.com (CN); justinmrao@outlook.com (JMR)

Received: September 30, 2015

Revised: September 3, 2016; June 16, 2017

Accepted: June 22, 2017

Published Online in Articles in Advance:
January 29, 2018

<https://doi.org/10.1287/mksc.2017.1065>

Copyright: © 2018 INFORMS

Abstract. On search keywords with trademarked terms, the brand owner (“focal brand”) and other relevant firms compete for consumers. For the focal brand, paid clicks have a direct substitute in the organic links below the paid ad(s). The proximity of this substitute depends on whether competing firms are aggressively bidding to siphon off traffic. We study the returns to focal brands and competitors using large-scale experiments on Bing with data from thousands of brands. When no competitors are present, we find a positive, statistically significant impact of brand ads of 1%–4%, with larger brands having a smaller causal effect. In this case, the effective “cost per incremental click” is significantly higher than what focal brands typically pay on other keywords. When the focal brand ad is present, competitors in paid positions 2–4 can “steal” 1%–5% of the focal brand’s clicks and raise its costs by shifting traffic to the paid link. Finally, for a set of brands that face competition on their brand search but choose not to advertise, competitors “steal” 18%–42% of clicks, suggesting a strong causal effect of position. Under such position effects, we find the return on investment on defensive advertising to be strongly positive.

History: Avi Goldfarb served as the senior editor and Juanjuan Zhang served as associate editor for this article.

Supplemental Material: Data and the online appendix are available at <https://doi.org/10.1287/mksc.2017.1065>.

Keywords: advertising • paid search • field experiments • causal inference • electronic commerce

1. Introduction

In sponsored search advertising, queries that include a company’s trademarked name form a substantial portion of advertising expenditure. At first glance this practice seems sensible: Consumers learn about products through a variety of channels and seek them out online using a search engine. Entering a specific brand signals product awareness, and click-through rates (CTR) tend to be higher than on similar non-brand keywords. Furthermore, since bidding on another firm’s trademarked term is legal, competing firms can step in and siphon off traffic, and focal brands may want to aggressively bid to minimize “traffic stealing” (as in the theoretical model of Sayedi et al. 2014). On the other hand, there is evidence that advertising on brand keywords might be ineffective. Because the focal brand’s website is the most relevant to the query, it almost always occupies the first “organic” result, which is shown just below the paid link(s), creating the possibility that the paid link crowds out free clicks. A recent paper by Blake, Nosko and Tadelis (2015, herein “BNT”) finds almost complete crowd-out for a single, well known brand, eBay. Using a controlled experiment, they document that when eBay stopped bidding on its own keywords, 99.5% of traffic was retained via the organic link.

In this paper, we aim to deepen our understanding of brand search in light of these considerations. We do so using a large-scale field experiment run on Bing. At the time of the experiment, the maximum allowable number of ads placed above the organic results on Bing and Google was four. In the experiment, this cap was exogenously reduced to 0, 1, 2 or 3, which allows us to study user behavior in the absence of paid search advertising all together and with an exogenous cap placed on the number of ads. We focus on the 2,500 most searched brands, which provides a sample with rich variance in measures of brand capital.

Our analysis reveals that brand search advertising can be effective and that effectiveness hinges on two key factors, i.e., (1) the prominence of the focal brand, and (2) whether competitors are bidding on the focal brand’s trademark. In the absence of competing ads, less prominent brands tend to gain more incremental traffic from advertising on their keywords. In the presence of competitors, the focal brand’s advertising “defends” the top paid position on the page, shifting competitors’ ads to positions 2–4 and reducing traffic “stealing.” Thus, our results are consistent with the findings of BNT for a company like eBay, but show that eBay’s case as a very strong brand facing no competitors is not the norm. By using a diverse sample of firms,

we provide robust, practical guidance for marketing managers to optimize brand search expenditure.

We start with a case of a focal brand facing no competition on its branded queries. Such a focal brand can experimentally estimate ad effectiveness by pausing ads on branded queries at random times or in random geo-locations (eBay used both strategies). Our experiments produce a similar variation by randomizing at the user level. We estimate causal effects for 824 firms that consistently advertised on their own-brand queries using the experimental conditions “Cap 0” (no paid links) and “Cap 1” (one paid link). In our sample of firms, sponsored links on branded search queries drive significant incremental traffic: Total clicks to focal brands increase by 2%–3% on average. This effect is significantly larger for lesser known brands, while the strongest brands in our sample show effects closest to that of eBay. Even for smaller firms, the incremental traffic is still relatively modest, but the “crowd-out” of free clicks tends to be quite large: Addition of a brand ad shifts nearly half of the clicks on the focal brand’s organic link(s) to the paid ad, thus exceeding the causal effect by more than a factor of 10.¹

We now examine what happens when competitors advertise on branded queries. While prominent firms have fought legal battles to block competitors from bidding on their branded keywords, the courts have consistently upheld the legality of the practice. If one or more competitors clear the reserve price and the focal brand submits a bid as well, we observe that the focal brand almost always occupies the top slot.² The competitors appear below the top ad, but above the organic results. We measure the impact on the focal brand from two channels, i.e., (1) click “stealing” and (2) increasing the cannibalization of free clicks. We estimate both quantities by fixing the set of focal brands and comparing the “Cap 1” condition to “Cap 2,” “Cap 3,” and the control. We find that competitors can steal only 1%–5% of clicks, with the magnitude depending on the number of competitors present and brand attributes. The addition of competing ads has a much larger impact on crowd-out rates. When facing no competitors (Cap 1), 60% of total clicks to the focal brand’s website are paid (the rest are free). When we randomly add in competitors, the fraction of paid clicks increases by 10% for the first competitor added, 9% for the second, and 5% for the third, reaching 84% for a full slate of competitors.

The final case to consider is when competitors advertise and the focal brand does not. For a fully randomized comparison, we would need to remove focal brands while keeping the ads of competitors. Because the experiment always preserved the auction ranking for business considerations, we do not have such variation. Our approach is to use cases in which competitors advertise and the focal brand consistently chooses not to. To control for brand strength and characteristics, we place brands in three categories based on their

CTR in the no ads (“Cap 0”) condition. The high CTR group matches the case of those brands that choose to advertise on their own keywords, and here we find that a single competing advertiser in the top position acquires 15%–20% of searchers, suggesting a strong effect of ad position. This large impact of ad position is consistent with the recent finding that position is more important when consumers are less familiar with the firm (Narayanan and Kalyanam 2015); theoretical work also predicts this pattern (Jerath et al. 2011). Interestingly, for the other two CTR groups, the effect size is nearly identical to the high group. The key difference is that for the low CTR group (often corresponding to brands that are sold by licensed resellers), most of these clicks come at the expense of other firms on the results page, not the focal brand. The middle group lies between the high and low, with about half the clicks coming at the expense of the focal brand. In all scenarios, click stealing substantially increases as more competitors are exogenously added to the page, which is consistent with our first set of findings that additional competing ads lower CTR on the organic links.

Note that we use two different sets of companies to estimate position effects of competitors’ ads, with the presence of the focal brand’s ad being endogenously determined by the company. Thus, even though we are controlling for baseline clickability on the focal brand’s links, our estimates can be upward or downward biased. For example, focal brands will be more likely to advertise when competitors’ ads in the top position steal more traffic, in which case we underestimate the position effect. Alternatively, focal brands might know that their own brand ad will not get many clicks in the presence of competitors’ ads, in which case we overestimate the position effect. At the same time, given the relatively uniform level of traffic going to competitors in the top position across the brands that we study, we find the position effects explanation for the difference in competitors’ traffic to be more plausible than selection issues.

Our final finding is that competing firms tend to be much smaller than focal brands. For the 564 firms that consistently face competition, the median Alexa website rank (across all websites for U.S.-located visitors) for the focal brand is 8,000, whereas it is 80,000 for the top ranked competitor, 145,000 for the second, and 178,000 for the third. These already drastic differences become larger if converted to page views, since the page view distribution is “heavy tailed” (Kumar and Tomkins 2010). Typically, much smaller competitors are “piggy-backing” off the awareness of their larger rivals, i.e., brands with which they aspire to associate, in a form of targeting. For firms of similar size, it is a well known result that allowing firms to target each other’s customers with special offers enhances competition by putting firms in a prisoner’s dilemma

(Thisse and Vives 1988), which has been explicitly tied to the context of sponsored search as well (Desai et al. 2014). Given the repeated nature of the interaction, we might expect that cooperative “live and let live” strategies could be supported in equilibrium. Yet, in practice, we see that the size asymmetry greatly limits the focal brand’s ability to punish, which helps explain the continued prevalence of competing firms’ advertising in the marketplace.

Our core contribution is to assert that focal brand prominence and the level of competition on branded keywords are the key factors in determining returns to the focal brand. In the absence of competitors, the average focal-brand ad causally increases total page CTR by 2.27% and shifts about half of the free clicks to the paid link. This means that for each incremental paid click, the firm must pay for about 16 clicks that would have been free, though this varies with brand strength. If a firm is not facing any competition, our guidance to marketing managers is to conduct “ad pause” experiments to measure “cost per incremental click”; in our sample of firms it tends to be higher than the cost per click (CPC) they pay on searches in which they do not have a high organic ranking; this suggests that there may be a way to reallocate expenditure to increase return on investment (ROI). In the presence of competition, our results suggest that own-brand ads play an important “defensive” role as competitors stand to siphon off a considerable fraction of traffic if they are allowed to occupy the top position(s) above the focal brand’s organic link(s). The focal brand ads fend off much smaller brands that are using the awareness of the focal brand as a targeting mechanism. Increasing the number of competitors causally raises the CTR on the focal brand’s paid link at the expense of free clicks, raising the cost of defense. Nonetheless, our results suggest that the implicit ROI versus the counterfactual of not advertising appears to be positive, endorsing the practice of using brand ads defensively. For smaller brands, the evidence suggests that targeting customers via awareness of larger brands, especially if those brands do not have defensive positions, is a strategy marketing managers should seriously consider.

2. Context and Related Literature

We define brand keywords as queries that consist of a trademarked term, where the trademark holder occupies the top organic slot. Competitors using a trademarked term to guide their bidding is a contentious practice. Focal brands dislike the fact that their competitors can target a user who has expressed an explicit interest in them. Indeed, these firms may raise brand awareness with other forms of advertising with the goal of monetizing this awareness via search (Lewis and Nguyen 2015); competitors entering the equation

make this more difficult. Despite many trademark infringement lawsuits using these lines of reasoning, the courts have consistently upheld the legality of showing competing ads on branded queries. However, the use of trademarked terms in a competitor’s ad text is not allowed, though an exception was granted in 2009 to licensed resellers.³ Chiou and Tucker (2012) study this change and find that it did not damage the focal brand because it made the competing resellers less distinct.

In sponsored search, advertisers pay for “consideration,” as measured by clicks; thus, clicks are the central unit of analysis in sponsored search.⁴ Early experimental evidence on click substitution patterns comes from Reiley et al. (2010), who showed that organic links and ads are substitutes for each other. This substitution pattern is overwhelmingly present in our experimental data as well, and has also been found in structural work (Jeziorski and Segal 2015).⁵ Reiley et al. (2010) further show that more ads can increase total CTR for the ad placed in the top slot, because organic links act as slightly better substitutes for ads. As discussed in Section 1, Blake et al. (2015) ran a large experiment with eBay, and the results of the experiments led the firm to discontinue brand search ads. In the experimental period, eBay did not face competing ads.

The next relevant strand of the literature studies how the position on the page impacts user choice. Craswell et al. (2008) use fully randomized experiments of algorithmic results (exogenously shuffling links) to show that “position effects,” i.e., the causal influence of position on the page, can be large near the top of the page. Agarwal et al. (2011) conduct a field experiment with a retailer and find a strong causal effect of position on CTR. Narayanan and Kalyanam (2015) study position effects in the ad slate with a regression discontinuity approach, which uses the fact that position is determined by a continuous “rank score.” They find position effects can be quite large, especially at the top of the page, but vary considerably depending on a user’s experience with the focal brand, and are smaller for the focal brand than competitors, and stronger when the advertiser is smaller. Granka et al. (2004) presents eye-tracking evidence that most users use a “top-down cascade” approach to searching the page. Later work has shown that search is more complex than simple top-down traversals, although an overall top-down pattern still dominates (Dupret and Piwowarski 2008).

The final relevant piece of the literature models bidding in sponsored search. Desai et al. (2014) model interactions between competitors on brand keywords and discuss the prisoner’s dilemma nature of the interaction. Yang et al. (2013) model auction entry and show that increased competition tends to hurt incumbent advertisers but helps the search platform. Jerath et al. (2011) present a related theoretical model where there

are low- and high-quality firms, and show that inferior firms have a greater incentive, all else equal, to locate at the top of the page, yet the superior firms may get more clicks even though they occupy lower positions. In our setting, this occurs when the brand's organic link still gets the majority of clicks when there are rival ads above it.⁶

3. Empirical Setting

In this section, we provide details of the brand search process, experiments, data, and estimation.

3.1. A Description of the Brand Search Process

Drawing on past work in marketing and economics, we present a stylized description of brand search to ground our investigation. We posit that consumers engage in search to achieve an end objective, such as buying a good that serves a particular function. In our study, we observe two key steps in this process: (1) Searching a branded term signals *awareness* of the focal brand and intention to find the brand online; (2) Clicking a link involves *consideration*. Awareness does not imply consideration; a consumer can choose to visit a competing brand or opt to not consider any of the firms. Because competing firms might satisfy a user's end objective, they have an incentive to intercept consumers with links in the search results. Indeed, this seems like it would be a useful form of targeting for lesser known products that have similar functionality. The main way to do this is by bidding on paid links for keywords containing branded terms since they are unlikely to have sufficient relevance to branded queries to appear high in the organic results.

We assume consumer i examines the first N_i links on the page and chooses the one with the highest expected utilities, where N_i can differ by consumers. Such a description of search is consistent with past work using eye-tracking (Granka et al. 2004, Dupret and Piwowarski 2008) and papers studying the causal influence of position on the page (Craswell et al. 2008, Narayanan and Kalyanam 2015). Incomplete search provides an additional incentive for firms to use paid listings because it allows their link to enter the awareness sets of more searchers and face a smaller number of competing alternatives in expectation. Furthermore, this setting suggests that if competitors can get clicks when the focal brand occupies the top ad slot, then at least some consumers opt to consider them despite their initial awareness of the focal brand. Position effects emerge from incomplete search; the fewer links users consider, the higher returns there are from "defensive positions" at the top of the page. In the absence of competing firms, a focal brand ad can impact choices by shifting the organic results of competitors down the page and through including any additional information that increases the chance of gaining consideration.

3.2. Experiment Description

The data in our study come from a series of randomized experiments on the Bing search engine. On Bing, the sponsored listings that appear at the top of the page, above the organic listings, are known as the "mainline." A maximum of four mainline ads are shown on a given query, the same practice used by Google. Absent experimentation, the number of ads and their composition is endogenously determined by firms' bids, the reserve price, and the hard cap at four. A cross-sectional regression that analyzed differences in the number of advertisements by keyword would conflate the true effectiveness of advertisements with differing environments across keywords. We use the experimental variation to control for these confounding factors.

The experiments were conducted on a small fraction of U.S.-located users over nine days in January 2014 with randomization at the user level. Four experiments took place, in which the maximum number of mainline ads was limited to 0, 1, 2, and 3. Each experiment had a balanced control group, which corresponded to the maximum of four mainline ads, the typical production setting. This is standard practice in online experimentation, as it provides a check that each experimental "line" was correctly executed.

The treatment limited the number of ads that could be shown, but often this cap was not binding. For instance, in the treatment group that limited mainline ads to a maximum of 3 ("Cap 3" to use the terminology we use throughout), if there were not enough bidders who met the reserve price to fill the three slots, then fewer than three ads were shown. We carefully control for this issue by selecting only queries that matched into bidding data in which an ad would have been shown in the absence of the experiment. See Appendix A for more details on this process.

Because of the nature of our experiment, we can estimate advertising effectiveness only for companies that are already advertising on their keywords on Bing, which we refer to as "treated." While we are unable to estimate treatment effects for an average company for which consumers search on Bing, our "treated" companies are interesting on their own as these are the brands that choose to advertise. Given that the decisions to advertise are endogenous, we would also expect the average treatment effect for the current advertisers to exceed the average treatment effect across all companies of similar size.

3.3. Data Description

To identify brands, we extracted 87,000 retailer and brand names from the Open Directory Project.⁷ A search is characterized as a brand query if and only if (1) the query is on this list, meaning it is a verified firm brand, and (2) the query matches the domain

name in the first organic position. We focus only on brands that are in the first organic link because this selects true brand queries. Queries for brands that are not in the first organic position might be searches of a different nature, perhaps not meant to get directly to the brand page, but to a broader set of sites. Figure 1 provides an example of a brand query. Queries are simplified using standard techniques, e.g., we treat “Macy’s,” “macys.com,” “macys,” and “macy’s” as the same query. We focus on searches with 0 or 1 clicks on the page, ignoring rare instances of 2 or more clicks.⁸

Table 1 gives the number of brands binned by the number of observations for those brands in all control conditions combined. Sixty-five percent of all brands in the control group have fewer than 11 exposures, but represent only 0.18% of all traffic, whereas 96% of traffic comes from the 1,045 brands that have 1,000 or more exposures. We keep the 2,517 companies with over 350 exposures, which cover 98.7% of the market activity.⁹ Of the selected 2,517 companies, 824 advertise on their own brand keywords more than 90% of the time.

In estimating the direct returns to brand search advertising, we focus on these 824 brands. More detailed information on all of the firms is given in Appendix A, Figure A.1.

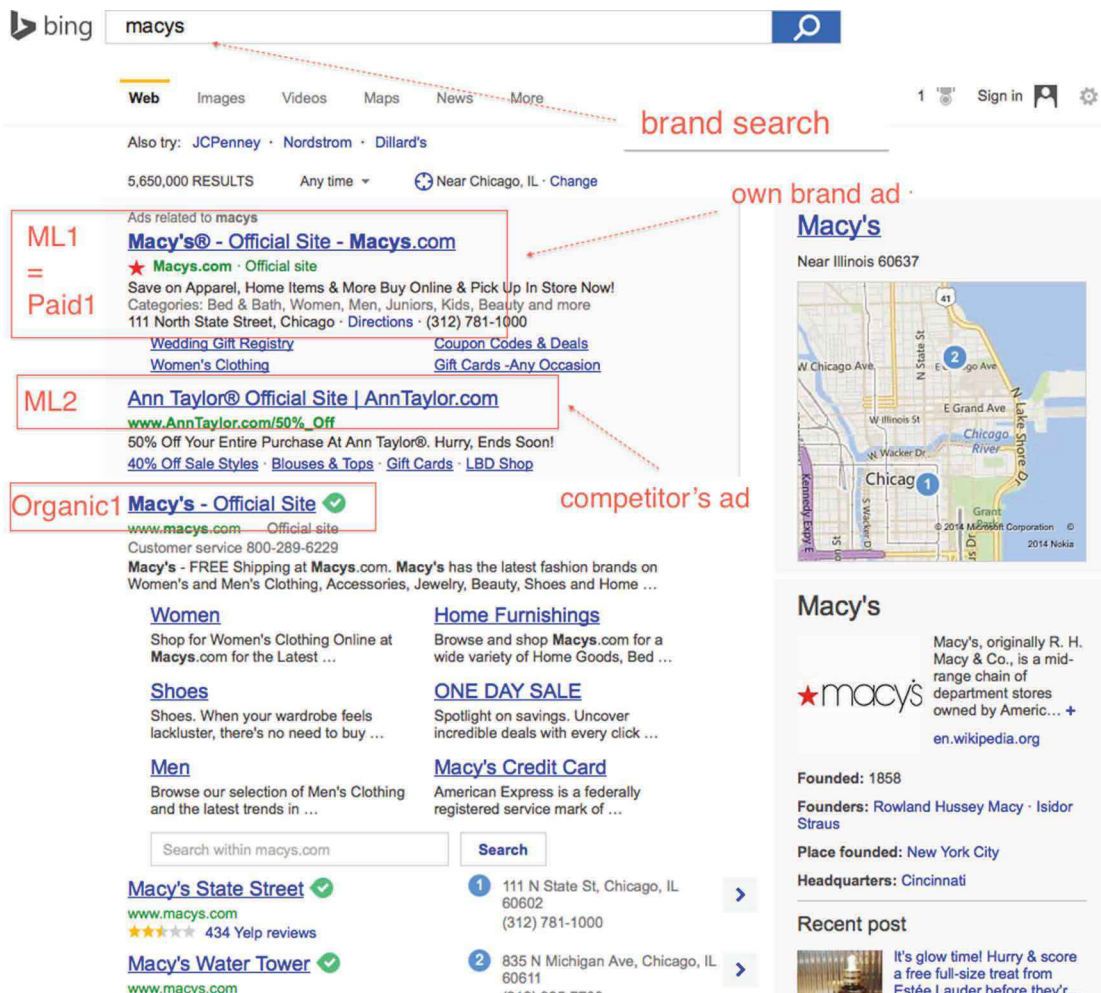
3.4. Estimation

For each brand j , we observe a number of brand searches in each experimental condition c , N_{jc} . For each search, among other things, we observe the URLs of organic links shown on the page, URLs of paid links shown, and click decisions of consumers. We classify the URLs as belonging to the focal brand if it matches the brand name and as belonging to competitors otherwise. We estimate the probability of clicking on a focal brand’s link across experimental conditions using a simple frequency estimator

$$\widehat{\Pr}(\text{click } j \text{ in } c) = \frac{1}{N_{jc}} \sum_i I(i \text{ clicks } j \text{ in } c),$$

where $I(i \text{ clicks } j \text{ in } c)$ follows a Bernoulli distribution. The estimator has expectation of p_{jc} and the variance

Figure 1. (Color online) Brand Search Example



Note. This example has two mainline ads, i.e., own brand ad in mainline 1 and a competitor’s ad in mainline 2.

Table 1. Ad Coverage in the Control Condition

Number of exposures in control	Number of brands in control	Percentage of brands (%)	Percentage of traffic (%)	Percentage of own ads in ML1 (%)	Percentage of competitor's ads in ML1 (%)
1	4,869	23.1	0.02	3	30.6
2	2,773	13.1	0.02	4.1	32.4
3	1,686	8	0.02	6.3	30.8
4–10	4,315	20.5	0.12	10.2	34.5
11–100	4,200	19.9	0.64	19.8	34.6
101–1,000	2,202	10.4	3.6	42.64	28.5
> 1,000	1,045	5	95.6	43.8	13.6
Total	21,090	100	100	14.4	31.4

Notes. The percentage of ads is computed across companies. For example, companies with 4 exposures and companies with 10 exposures are given the same weight in group 4–10. The total frequency is also computed across companies, unweighted.

of $p_{ic}(1 - p_{jc})/N_{jc}$, where p_{jc} is the true probability of a click. Similarly, we estimate the probability of clicking on competing firms, j' , after searching for brand j in the experimental condition c as $\widehat{\text{Pr}}(\text{click } j' \text{ in } c)$.

We compute $\widehat{\text{Pr}}(\text{click } j \text{ in } c)$ for each combination of brand j and experimental condition c . Experimental conditions were balanced to compare treatment and control groups. In Section 4, we compare treatment conditions to each other, e.g., comparing Cap 0 to Cap 1 allows us to isolate the effect of own brand advertising in the absence of competitors in positions 2–4. Differences in $\widehat{\text{Pr}}(\text{click } j \text{ in } c)$ across experimental conditions correspond to the average treatment effects of brand j 's or competitors' advertisement on the traffic to j 's website coming from branded search. To make sure different conditions can be compared without bias, we check that the associated control conditions do not differ from each other. See Online Appendix B for the click probability estimates for the focal brand's/competitors' organic/paid web links, along with the 95% confidence intervals around these estimates. We do not find any statistically significant differences in these comparisons.

4. Results

In this section, we separate our analysis by competitive scenario, starting with the case of no competitors present, moving to the case where a focal brand and competitors are present, and then to the case where only competitors are present. We then present data on costs to infer ROI, and close by examining competitor attributes.

4.1. Ad Effect Without Competitors Present

We examine the effectiveness of advertising in the absence of competing firms by comparing focal brand click probabilities in the Cap 0 and Cap 1 conditions. This corresponds to the probability that an individual arrives at the website of the searched brand. Figure 2(a) plots the average estimate of these probabilities across

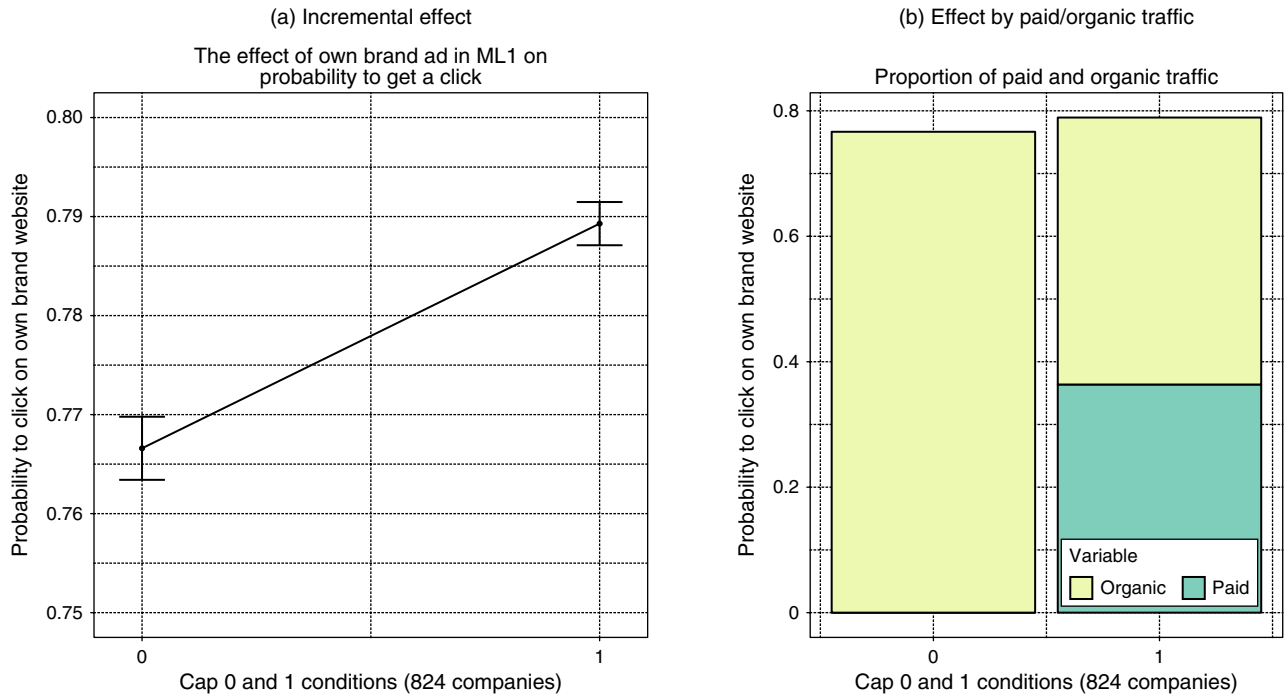
824 firms, as well as the corresponding 95% confidence interval. The figure shows that advertising on one's own keyword drives an incremental 2.27% of traffic. This estimate is overwhelmingly statistically significant; note, however, that the y -axis is "zoomed in."

Figure 2(b) shows the traffic to the focal brands' website by link type. In Cap 0, all traffic navigates to the focal brand's website through the organic links. In Cap 1, about half of the traffic goes through the paid ad on the top of the page, reflecting "cannibalization" or "crowd-out." In Online Appendix D, we show that most of this traffic would have gone through the first organic result, but we document statistically significant crowd-out for the first six slots (the focal brand often occupies many of the organic results); the effect size declines with position. The difference in the overall bar height represents incremental clicks, which shows that while there is a causal effect of the focal brand's ad, the majority of paid clicks are those that were crowded out from the organic channel.

We next examine if ad effectiveness differs across brands. We focus our analysis on a subsample of 493 brands that have a sufficient amount of traffic for reliable brand-specific estimates.¹⁰ Figure 3 shows a histogram of the brand-specific estimates with an overlaid normal density calibrated to the data. The empirical distribution has heavier tails than the normal density, and we formally test and readily reject the hypothesis that the observed heterogeneity is driven by sampling variation alone.¹¹

We decompose this heterogeneity using brand prominence, which we proxy for with the website ranking among U.S.-located users from Alexa.com, a widely used website ranking service. We also considered log global rank, "bounce rate" (the probability of a visit fewer than 30 seconds), the fraction of traffic from search engines, pages viewed per day, and time spent per day. (See Table S1, Online Appendix C for a summary.) Specification (1) in Table 2 regresses the advertising causal effect estimate on log U.S. rank. Brands

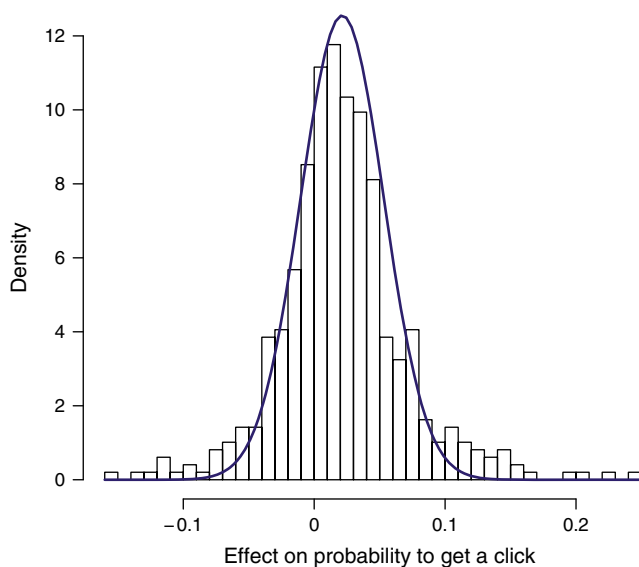
Figure 2. (Color online) The Effect of Advertising on Focal Brand’s Own Keyword



Notes. Point estimates are computed for each brand using the frequency estimator defined in Section 3.4. Results are averaged across brands. The error bars give \pm two standard errors.

with a higher ranking (closer to 1) tend to have a smaller advertising effect. On average, the regression predicts the ad effect for a very well known company is 2% lower than for the median company in our sample, meaning we would predict a near-zero effect for such firms.

Figure 3. (Color online) The Distribution of Brand-Specific Heterogeneity



Notes. Point estimates are computed for each brand using the frequency estimator defined in Section 3.4. The blue line corresponds to the implied normal density.

More prominent brands could gain less from advertising because they occupy more space on the page due to richer and more numerous organic results. Specification (2) examines if the effectiveness of advertising is correlated with the number of organic links for the focal brand. We do not find a statistically significant relationship. Specification (3) examines the correlation between the effectiveness of advertising and the space occupied by the top organic result on the page. This space is measured by the number of “deep

Table 2. Relationship Between Brand Capital and Advertisement Effectiveness

	Dependent variable			
	Effect of focal brands ad			
	(1)	(2)	(3)	(4)
log(U.S. Alexa website rating)	0.004** (0.001)	0.004*** (0.001)	0.002 (0.001)	0.002 (0.001)
Number of own organic links		0.001 (0.002)		-0.001 (0.002)
Deep links			-0.002 (0.001)	-0.002 (0.001)
Detail card			-0.002*** (0.0005)	-0.002*** (0.0005)
Observations	493	493	493	493
R ²	0.013	0.014	0.062	0.063
Adjusted R ²	0.011	0.010	0.057	0.055

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

links” (sublinks below the first organic results) and “detail cards” (detailed informational panels including things such as maps and ratings). Figure 4 provides an example of deep links and detail cards on the search page. Brands with detail cards have a significantly lower advertising effect. The coefficient on log U.S. rank becomes insignificant in this specification, suggesting that the mechanism behind the higher effectiveness for smaller brands is the size and nature of the content of the first organic result. Specification (4) includes all measures of how much space the focal brand’s results occupy on the page; once we account for the nature and number of organic search results, the coefficient on website rank drops by 50% and loses statistical significance. Finally, note that the differences across brands we observe could also be driven by the types of users who choose to search for a particular brand. Because users are not randomly assigned to queries, we are unable to tease apart differences driven by user versus brand characteristics.

4.2. The Role of Competitors in Positions 2–4

In this section, we shift our focus to the case when one or more competitors clears the auction’s reserve price while the focal brand still occupies the top slot.¹² In this case, the role of the focal brand’s ad is defensive:

In the absence of the ad, competitors would occupy the top paid position on the page. For the 824 companies analyzed above, 564 face 3 competitors advertising in positions 2–4 at least once in all experimental conditions. For these companies, we can compare the traffic to the focal brand’s website without any competing firms present in the paid links (Cap 1) to the case with one competing firm (Cap 2), two competing firms (Cap 3), and three competing firms (Controls) present. We use the bidding data to keep only the cases when competing firms would have advertised in the absence of the experiment.

Figure 5(a) displays the impact of competing firms in positions 2–4 on overall traffic for the focal brand. The first point displays the probability of navigating to the brand’s website when only the brand ad is present. The second through fourth points display the traffic to the brand’s website when adding in competitors to slots 2 through 4, respectively. Focal brands are split by the median website traffic ranking. The relationship is downward sloping and significant. However, the magnitude is modest: The full slate of competitive ads reduces the traffic to the focal brand by an average of 4.3%. The lower brand capital firms show a level shift down in CTR but a similar pattern with respect to competing ads.

Figure 4. (Color online) Deep Links and Detail Cards Example

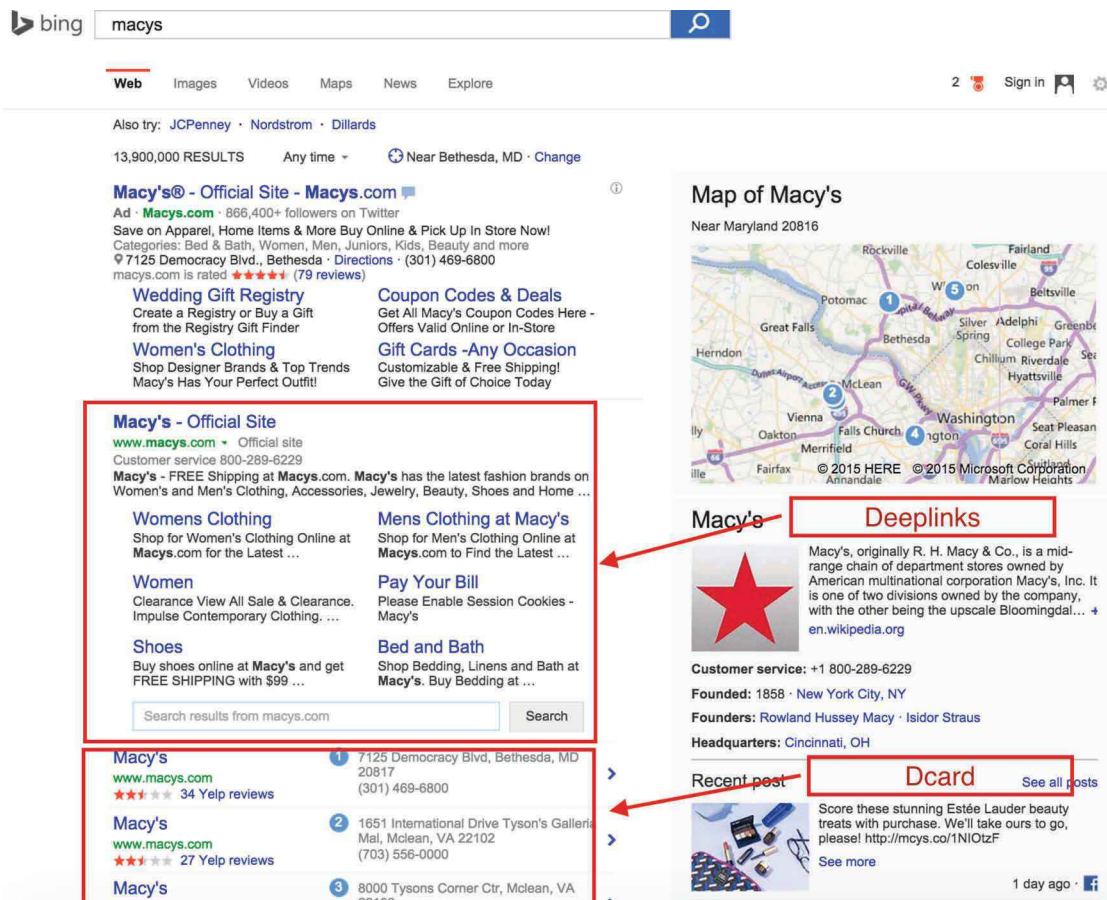
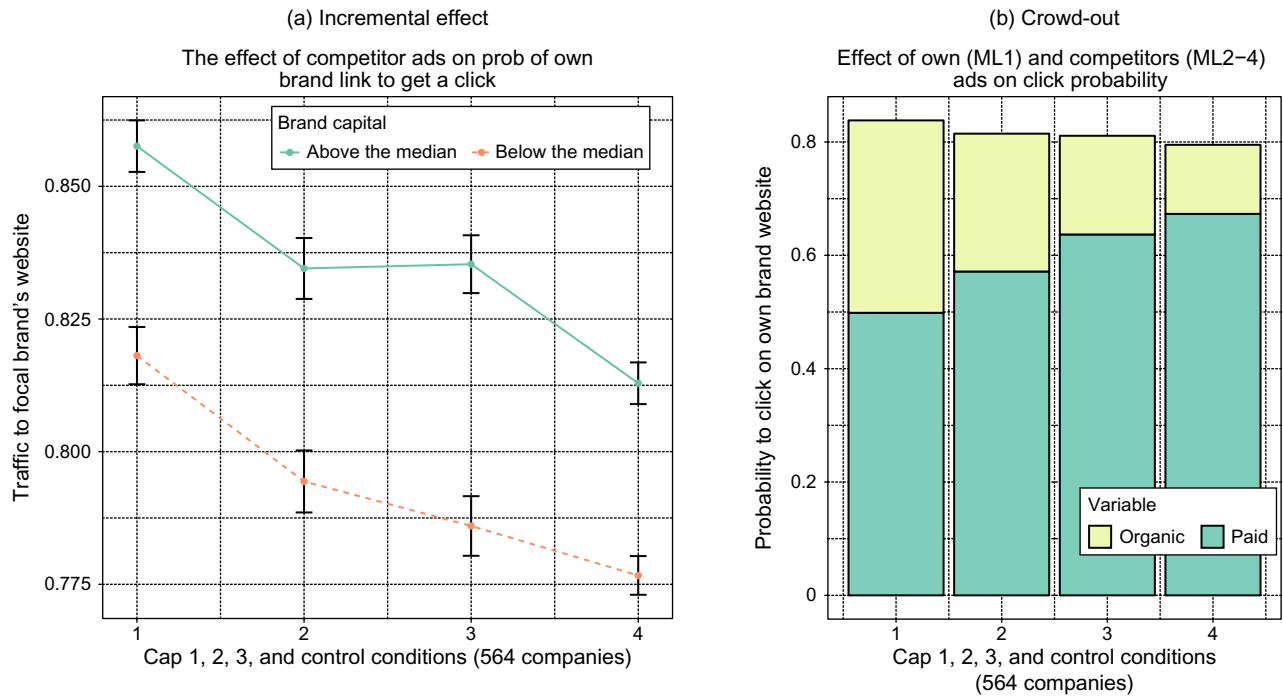


Figure 5. (Color online) The Effect of Competitive Ads in Mainline Slots 2–4



Notes. Point estimates are computed for each brand using the frequency estimator defined in Section 3.4. Results are averaged across brands. Error bars are \pm two standard errors.

Figure 5(b) shows the effect of competing firms on the crowd-out of organic traffic. When only the focal brand’s ad is present, 60% of the traffic navigates through the paid link.¹³ This fraction increases to 70%, 78%, and 84% with one, two, and three competing firms, respectively.

4.3. Impact of Competitors in the Absence of a Focal Brand Ad

We now examine the case when competing firms occupy the top paid position. Because of the nature of

the experiment, this occurs for a different set of focal brands: “Cap” conditions exogenously remove ads from below but not from above. We identify brands for which competing firms occupied the top paid position more than 90% of the time during our sample period. There are 181 such brands in the sample. By construction, this set of brands does not overlap with the 824 brands we have used for the analysis above. Figure 6 presents histograms of click probability in the Cap 0 condition, which removes the impact of the ads themselves, for brands with competitors in the top position

Figure 6. (Color online) Heterogeneity in Firms That Bid and Do Not Bid on Their Own Brand Keywords

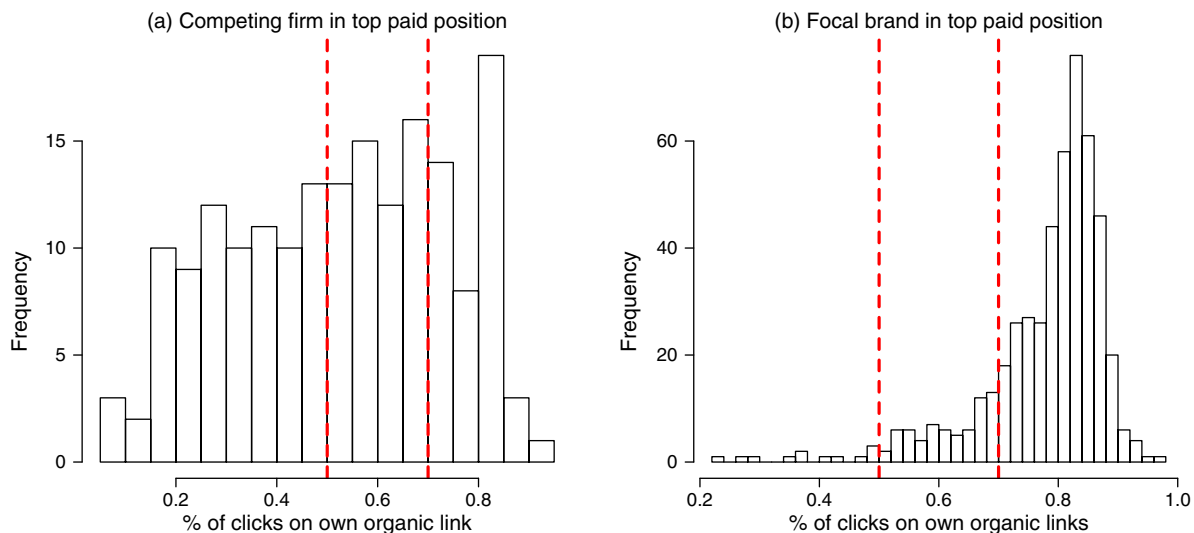
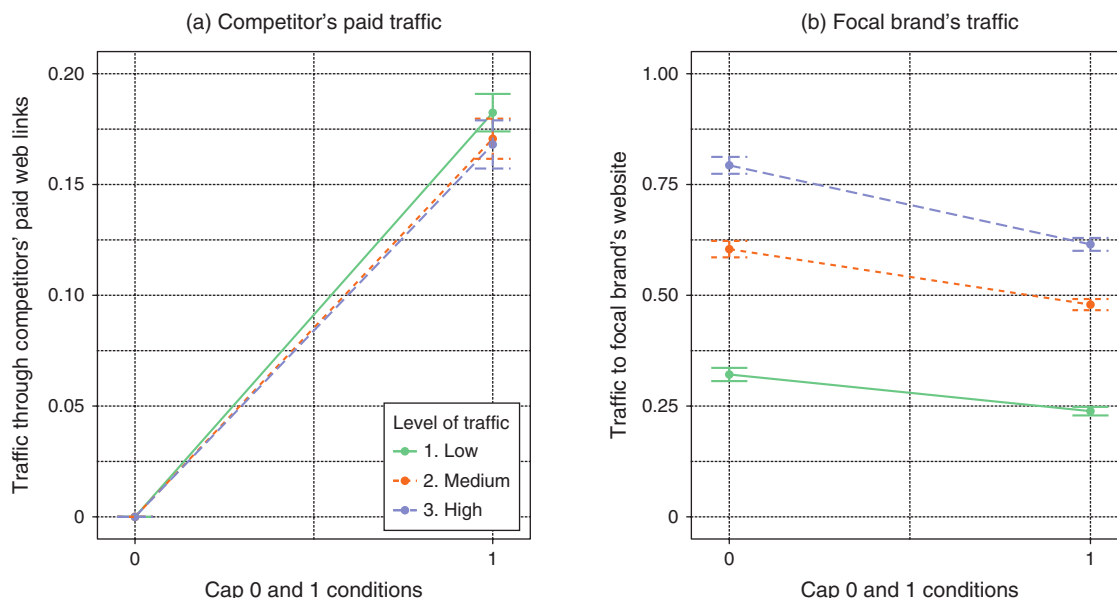


Figure 7. (Color online) Effect of Competing Firms Advertising in Top Paid Position on the Page by Level of Traffic to Focal Brand's Website



Notes. Point estimates are computed for each brand using the frequency estimator defined in Section 3.4. Results are averaged across brands. Error bars give \pm two standard errors.

and focal brands in the top position. For brands that decide to advertise, we observe click probabilities that usually exceed 70%, whereas lower organic click probabilities are more common when a competitor occupies the top slot. To make more reliable comparisons, we categorize brands by the amount of traffic they get in the absence of any ads: (1) “low” CTR segment for CTR less than 50% (80 firms), (2) “medium” segment when CTR is between 50% and 70% (56 firms), and (3) “high” segment for CTR greater than 70% (45 firms).

Figure 7(a) shows the increase in traffic to competitors in position 1 by firm type. Interestingly, competing firms get 17%–18% of traffic for all three types of focal brands. Figure 7(b) shows that this traffic comes from different places compared to the counterfactual of no ads, given by our Cap 0 condition. For the high traffic firms, almost all 18% come at the expense of the focal brand. For the medium and low CTR segments, 12.5% and 8.5% come at the expense of the focal brand, respectively, reducing the incentive for “defense.” Indeed, we see that it is less common to advertise in these cases. Averaged across all firms, focal brands lose 12% of clicks. All figures quoted in this paragraph are statistically significant well beyond the 0.01 level.

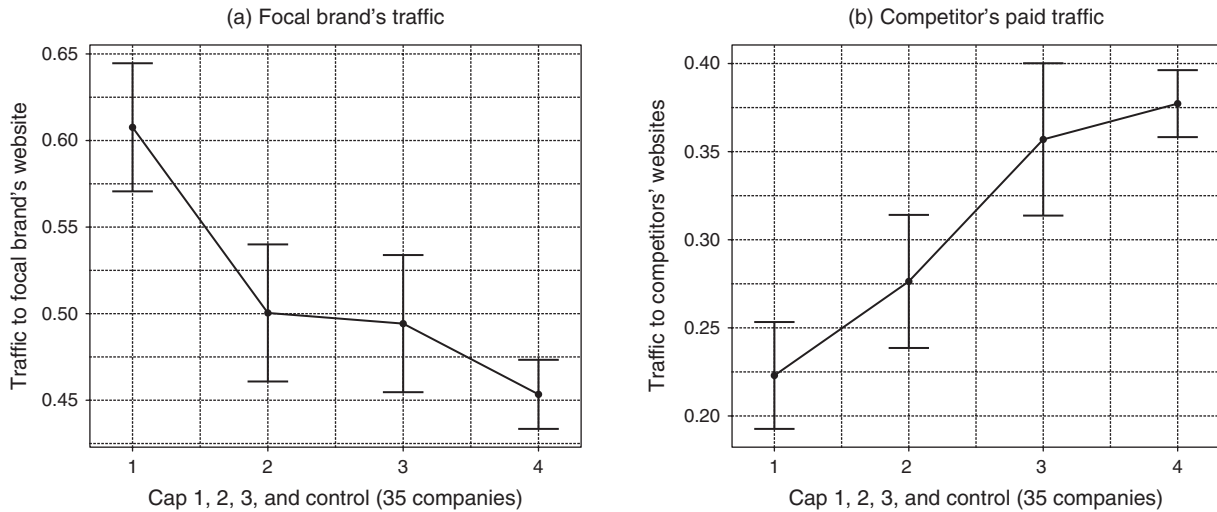
To get a rough idea of how additional competitors affect traffic to the focal brand, we use the 35 focal brands in the high CTR segment that faced up to 4 competitors at least once. Figure 8(a) shows that the probability that a consumer navigates to the focal brand’s website decreases from 61% in the case of 1 competitor to 50%

in the case of 2 competitors, 49% in the case of 3 competitors, and 45% in the case of 4 competitors, resulting in competitors intercepting 42% of traffic compared to the no ads condition. Figure 8(b) shows the analogous change in the fraction of traffic to competitors’ websites.

Overall, we find strong click stealing effects of competitors’ ads when the focal brand does not advertise on its branded keyword. One explanation for these results is the position effects of competitors’ ads: A single competitor in the top position on the page, on average, steals 18% of clicks from a high traffic brand, but a competitor following a focal brand’s ad steals only 1%–2% of clicks. Alternatively, this disparity can be driven by the difference in the set of companies we study, with selection bias possibly going in either direction. For example, focal brands will be more likely to advertise when competitors’ ads in the top position steal more traffic, in which case we underestimate the position effect. At the same time, focal brands might know that their own brand ad will not get many clicks in the presence of competitors’ ads, in which case we overestimate the position effect.

Our interpretation of the results above leans towards the earlier explanation of strong position effects. First, as presented in Figure 7(a), competitors can get almost the same amount of clicks across different types of focal brands that we study. This indicates that consumers are likely to click on the top link on the page regardless of the nature of the focal brand for which they have searched. Second, these results are robust if we

Figure 8. Effect of Second, Third, and Fourth Competing Firms Advertising in the Top Paid Position on the Page for Companies with High CTR



Notes. Point estimates are computed for each brand using the frequency estimator defined in Section 3.4. Results are averaged across brands. Error bars give \pm two standard errors.

focus only on companies with high brand capital: We find companies with a high level of brand capital in the set of 35 brands used in Figure 8 and re-estimate the amount of traffic stolen by competitors' ads.¹⁴ Results are presented in Online Appendix E, Figure S5, and are similar to the above, with 4 competitors stealing 45% of the focal brand's traffic. This indicates that results hold regardless of the nature of the competitors' brands that try to intercept the focal brand's traffic. Based on this, we consider the position effects explanation for the difference in intercepted traffic by competitors more plausible than selection issues, but we acknowledge that we cannot rule out selection completely as it can be driven by the unobserved factors.

4.4. Costs and ROI

In the absence of competing firms, a focal brand produces 2.27 extra clicks and 36.4 paid clicks per 100 searches. Because a firm is paying for about 16 clicks per incremental click, "cost per click" (CPC, the standard pricing metric reported by online advertising platforms) will sharply diverge from the true "cost per incremental click" (CPIC). We define CPIC as the cost of getting 1 incremental click, and compute it as follows:

$$CPIC = CPC \times \frac{\Pr(i \text{ clicks paid link of focal brand})}{\left[\Pr(i \text{ clicks links of focal brand} \mid \text{Ad}) - \Pr(i \text{ clicks links of focal brand} \mid \text{No Ad}) \right]}$$

where the numerator is the probability a click goes through the paid link of a focal brand and the denominator is the incremental effect of a paid link on the focal brand's traffic.

The CPIC/CPC ratio is a natural measure of crowd-out. Our informal estimate of 16 is actually a lower

bound on the average ratio because CPIC/CPC is a convex function in ad effect size. Jensen's inequality tells us that a convex function evaluated at average values is strictly less than the average value of the function. While focal brands are effectively paying a high multiple on their nominal CPC, it turns out that their CPCs tend to be very low because of how the generalized second price (GSP) auction rewards relevance. To get a better idea of effective costs, we focus on the subset of companies that rarely face competing firms' advertising and with enough data to compute the brand-specific advertising effects, leaving us with 268 brands.¹⁵ Table 3 presents a detailed summary of the costs paid per click by these brands and competing

Table 3. Cost of Clicks for Focal Brands and Competitors

Measure	Unweighted average		Weighted average ^a	
	All brands	Significant	All brands	Significant
$\overline{CPC}_{\text{brand}}$ (\$)	0.15	0.15	0.09	0.06
$\overline{CPC}_{\text{compet}}$ (\$)	0.86	0.78	0.61	0.47
$\overline{CPC}_{\text{brand-other}}^b$ (\$)	1.36	1.43	1.14	0.92
$\overline{CPIC}_{\text{brand}}$ (\$)	3.50 ^c	1.12	2.52 ^c	1.42
$CPIC_{\text{brand}} > CPC_{\text{compet}}$ (%)		51.2		92.7
$CPIC_{\text{brand}} > CPC_{\text{own-other}}$ (%)		55.9		90.6
N	268	43	268	43

^aWeighted by number of searches.

^bCPC paid by the brand on other keywords when they are not high in the organic listing. It is computed for 85% of the companies with data coverage.

^cLower bound computed with average values as described in the text.

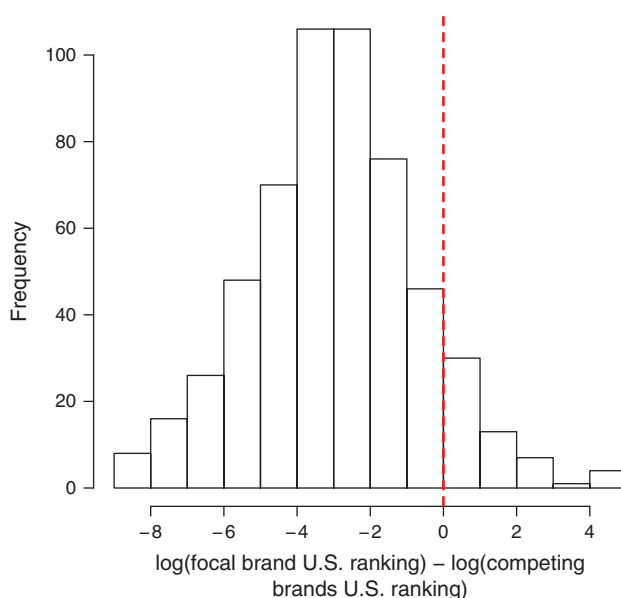
firms on the branded keywords. The average nominal CPC for focal brands is between \$0.06–\$0.15 depending on the sample (all versus those with a statistically significant ad effect) and the weighting. Competitors pay much higher prices, between \$0.47–\$0.86 on average, despite occupying lower positions.

For firms with a significant ad effect, the average CPIC is \$1.12 unweighted and \$1.42 weighted by searches. The lower bound on CPIC for all firms is \$3.50 unweighted and \$2.52 weighted. Both are substantially higher than the CPC of competitors. A second natural comparison point is what the CPC focal brands pay on keywords when they do not occupy a high organic position, and thus nearly all clicks are marginal. Even for the firms with a significant ad effect, CPIC exceeds the relevant comparison 93% of the time weighted by searches and 51% unweighted. We are unable to make firm-level comparisons when the effect size is not significant, but note that the lower bound given by evaluating the CPIC at average values exceeds the comparison measures of CPC on average. Taken together, the evidence indicates that brands that do not face competition tend to pay more for incremental clicks on their own brand keywords than they do elsewhere or their competitors pay on average. While these conditions do not always hold and are not always direct evidence of mistakes per se, they do suggest that this type of expenditure should be critically examined.

Our results suggest that advertising by competitors completely changes the story. A single competitor in the top position on the page, on average, steals 18% of clicks from a high traffic brand, but a competitor following a focal brand’s ad steals only 1%–2% of clicks. If this difference is due to strong position effects and not selection issues, focal brand ads have a strong ROI. This is because the defense is highly effective (the total CTR returns almost to the case when there is no advertising): Even though the focal brands must pay for 50 clicks to get 18 incremental clicks, their CPC is about 10 times less than they pay on other queries. Putting the pieces together, the implied CPIC is in an attractive range and, indeed, better than usual. Additional competing firms shift the focal brand’s organic link further down the page, which significantly increases crowd-out rates, but also the click stealing that would be expected if the focal brand was not present; the numbers work out so that the ROI still appears to be positive.

Note that ROI is positive relative to a counterfactual of a competitor stealing clicks. If competitors are not present, the returns to advertising are much lower. It is thus clear why brands have fought legal battles to ban competing firms from bidding on their trademarked terms. Smaller rivals use the focal brand’s awareness as a form of targeting, which then creates the need for defensive positions. As discussed in theoretical models, this enhances competition in a way that smaller

Figure 9. (Color online) Distribution of the Difference in Log Rank Between Focal Brand and Competitors



Note. Based on 564 firms that consistently face competition.

firms, who get very little search traffic themselves, and the platform tend to like and larger firms tend to dislike (Yang et al. 2013, Sayedi et al. 2014).

4.5. Strength of Focal Brand and Competitors

To understand the relative strength of competing firms versus focal brands, we plot the difference in the log(U.S.) website rank in Figure 9. The distribution is centered around -3 , which corresponds to the focal brand having e^3 (≈ 20) times higher rank. The red vertical line indicates when competitors have popularity equal to that of the relevant focal brand; very little mass lies to the right of this point. For the 564 firms that consistently face competition, the median Alexa rank of website popularity (across all websites for American visitors) is 8,000, whereas it is 80,000 for the top competitor, 145,000 for the second, and 178,000 for the third. It has been previously shown that website traffic displays heavy tails (Kumar and Tomkins 2010), meaning that having 20 times the lower rank corresponds to a much larger differential in terms of page views. Overall, the evidence clearly points to competitors consistently being much less prominent than focal brands.

5. Discussion and Conclusion

Our results provide a deeper understanding of why firms should advertise in brand search, and why they should not. In the absence of competing firms, focal brands can still improve the positions of their web links on the page by running a brand search advertisement. On average, focal brands get just over 2 extra clicks of 100 searches. The size and content of the first organic

result modulates the effect size: Because smaller firms tend to have less rich search results, they benefit more from brand ads. Thus, our results are consistent with the findings of BNT for a company like eBay, but show that eBay's case as a very strong brand is not the norm.

Competing firms greatly affect the market. When a focal brand occupies the top position, competitors steal a modest 1–5 clicks of 100 searches. Because of the nature of our experiment, we cannot remove the focal brand's ad while keeping the competitors, so we cannot directly estimate the position effect of competitors' ads. However, for a different set of high traffic brands, even a single competitor, on average, steals 18 clicks of 100 searches. Even though the comparison across competitive scenarios involves different firms that choose whether to advertise, we highlight that the impact competitors have with no focal brand present is 10 times higher than when the focal brand blocks competitors by occupying the top slot and the magnitude of click stealing is stable across firm's CRTs and brand prominence. This corresponds to competitors being relatively good substitutes for the focal brand and users examining a small number of links, which combine to produce strong position effects. These factors imply that focal brand ads play a defensive role, which greatly limits the ability of competitors to exploit strong position effects.

While we find it intuitive to discuss cases with and without competition on focal brand's search separately, it is important to keep in mind that the presence of competitors on the page is endogenous to the focal brand's advertising decisions. Thus, a company advertising on its brand keyword search and facing no competition should consider a potential entry of competitors when deciding to stop advertising on its keyword. Our results suggest that such competitive entry will lead to a large click stealing effect from the focal brand. However, the nature of our experiment does not allow us to directly estimate the effect of competitive entry and limits us to studying the advertisement effectiveness only under the current level of competition. We advise marketing managers to monitor the competitive response and its effect on the focal brand's traffic.

We discussed that while CPC is widely used and easy to understand, it often does not capture the true marginal costs of traffic. Indeed, CPC is similar for focal brands in the presence and absence of competitors, but our results indicate that true ROI hinges on their presence. The alternative metric we propose, CPIC, is economically sound, but computing it requires the relevant counterfactual, which may not be easily observable, can change over time or may be an unfamiliar concept to decision makers. We do, however, observe evidence that many firms broadly understand these complexities. For example, advertising on brand keywords is much more common when competitors are

present. Nonetheless, when we have the resolution to study behavior at the firm level, there is evidence that the depth of these complexities is not fully understood. Our guidance to marketing managers is to run search pause or geo-randomized experiments to measure CPIC to compute ROI.

Finally, our results show that firm heterogeneity is critical to understanding sponsored search. Smaller focal brands benefit from ads even in the absence of competitors. More strikingly, competitors are far less prominent than their associated focal brand; these small firms use brand search as a way of directly competing with their larger, more well known rivals. Because these competing firms tend to get little search traffic themselves, there is minimal fear of a reprisal in the form of bids on their keywords. Marketing managers at such firms should seriously consider using this tactic to target users, especially when focal brands are not defending their keywords. Based on our comparison across different experimental conditions, the evidence suggests that focal brands can almost, but not entirely, eliminate traffic stealing by placing an ad in the top position. Even in this case, given how small competing firms are, the traffic they attract could still be meaningful, revealing brand search to be an important form of competitive advertising.

Acknowledgments

The authors thank Christian Perez, Matthew Goldman, Giorgos Zervas, the participants of the 2015 Quantitative Marketing and Economics Conference and the 2017 UT Dallas Bass Conference, two anonymous reviewers, the associate editor, and the senior editor for helpful comments. Most of this work was completed when the first author was an intern and the third author was a full-time employee at Microsoft Research. All opinions represent the authors' own and not those of their employers.

Appendix A. Experiment and Auction Data

In the experiment, we randomly restrict the number of possible paid links on top of the search page. The control group corresponds to the default, which is a maximum of 4 advertisements in the mainline (Cap 4). There are 4 experimental conditions: Cap 0, 1, 2, and 3. The idea is similar to the control: Cap 0 does not allow any advertisements in the mainline, and Cap 3 allows at most 3 advertisements in the mainline.

This design of the experiment restricts us to studying only the cases where an advertisement is *eligible* to be shown in the mainline, which means that in the absence of our experiment advertisement will be shown in the mainline. For example, we cannot study the effect of an advertisement for a company that does not advertise on its own keyword; there will be no own brand advertisements in both Cap 0 and Cap 1 conditions.

Thus, we restrict our attention to cases where companies advertise. We are still facing a challenge: If a company advertises only 50% of the time on a given query and selects

search traffic where the effect will be higher, e.g., using geo-targeting, we cannot compare occasions with the advertisement to the treatment condition where the ad will be removed. For example, if we would like to estimate the effect of own brand advertisement in mainline 1 when in 50% of the cases the company advertises, and in 50% of the cases there is no advertisement, comparing occasions in Cap 1 conditions with a paid link shown to the entire Cap 0 conditions will bias the estimates.

To find the right treatment group, we need to allocate the occasions where the ad was actually removed from the mainline. In the example above, we would like to compare occasions with own brand paid link in Cap 1 to occasions in Cap 0 when own brand paid link would have been shown. To find such occasions, we collect the auction data for the search queries in the experiment. The allocation of positions in the mainline follows the standard GSP auction rules: Players submit the bids for a price of a click, the platform computes the “rankscore” of a given player, and players are allocated to positions in the mainline based on their rankscores. Given that the reservation level is cleared, a company with the highest rankscore gets position 1, a company with the second highest rankscore gets position 2, etc. Rankscore is proportional to the bid and a probability of click on the ad as computed by the platform

$$RS_j \propto b_j pclick_j^\alpha, \quad (A.1)$$

where b_j is a bid of company j , $pclick_j$ is a probability of company j to get a click, and α is the tuning parameter.

This implies that knowing the rankscores of bidders and the reservation level for a search query facilitates identifying which advertisement would be shown in the mainline in the absence of the experiment. To get this information, we exploit auction data collected by the advertising team. The experiment that we use was designed by removing the potential advertising slots from the mainline; bidding data was still collected.

Table A.1 presents the summary of matching experimental data and collected auction data for Cap 0–4. For Cap 1, 2, 3, and 4, around 56.3% of the search queries in the experimental data were matched with the auction data. A search query will not be recorded in the auction data if no advertiser submitted a non-trivial bid, as defined by the platform, so the unmatched data can correspond to queries with no bidders. The majority of unmatched queries correspond to occasions where no advertisements were shown, which supports this explanation.

The Cap 0 condition has a higher percentage of unmatched search queries. This indicates the problem with the matching,

given that the experiment was constructed to be balanced between the treatment and control groups. We further find that the percent of advertisements eligible for the mainline 1 position in Cap 0 is substantially different from the percent of advertisements eligible for the mainline 1 position in Cap 1, 2, 3, and 4.

This creates a potential problem for using the occasions with eligible brand advertisements for the Cap 0 condition. Consider the case of estimating the effect of own brand advertisement in mainline 1. Using the matched data, we would like to compare occasions with the own brand advertisement from the Cap 1 condition to occasions with the eligible own brand advertisement from Cap 0 conditions. We know that some occasions with the eligible own brand advertisement are missing from Cap 0. If this mismatch is correlated with the probability of a click on the own brand weblink, our estimate of the advertisement effect will be biased.

To check if there is a selection problem in Cap 0 matching, we estimate the effect of own brand advertisement in mainline 1 for companies that always advertise in mainline 1 on their own keyword.¹⁶ For these companies, the comparison of Cap 0 to Cap 1 provides the causal effect of own brand advertisement: We know that, if not for the experiment, search results in Cap 0 will have their own brand advertisement in mainline 1. We can also estimate the effect using only eligible own advertisement occasions in Cap 0 and Cap 1. If the estimates of the effect based on the two methods are different, we can confirm that the occasions in Cap 0 that have the eligible own brand advertisement in mainline 1 are correlated with the probability of click on the focal brand’s website.

Table A.2 presents the estimation results. The ad effect estimate based on all traffic is 1.68%. The effect estimated using only traffic with the eligible own advertisement is 0.63%. The difference in the two estimates is statistically significant.¹⁷

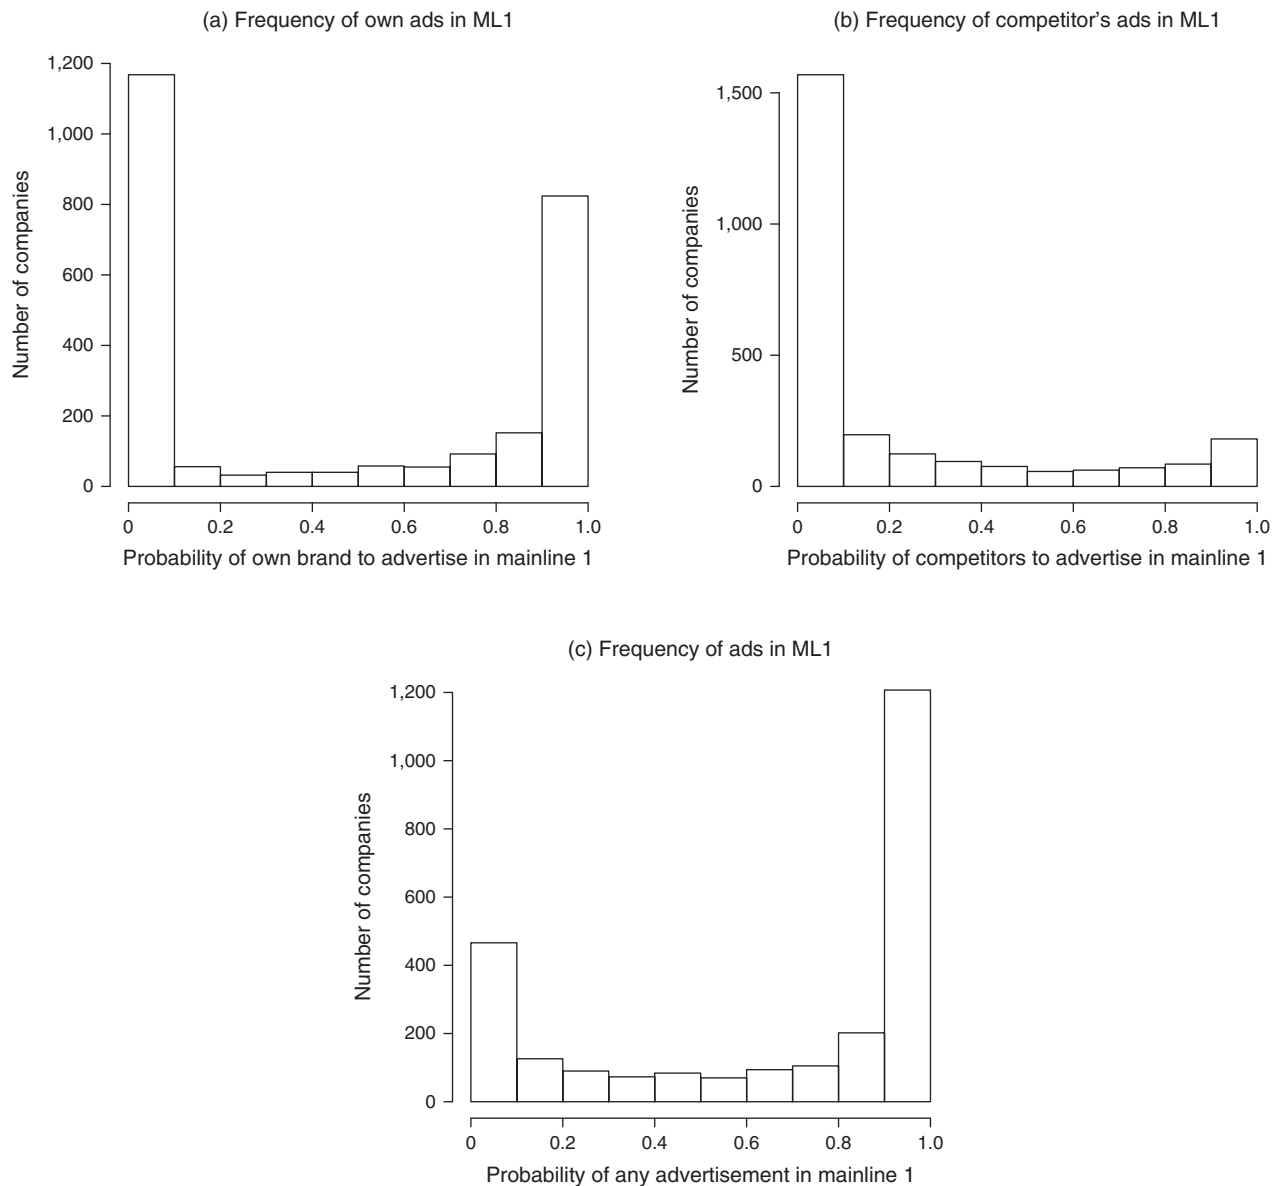
Table A.2. Effect of Own Brand Ad in Mainline 1 is Significantly Underestimated When Using Eligible Ads

	All queries	When own ad is eligible
N_{comp}	391	391
\hat{p}_{own0}	0.7867 (0.0022)	0.8115 (0.0029)
\hat{p}_{own1}	0.8035 (0.0015)	0.8179 (0.0015)
$\hat{p}_{own1} - \hat{p}_{own0}$	0.0168 (0.0027)	0.0063 (0.0033)

Notes. \hat{p}_{own0} is the probability of a click on own brand link in Cap 0. \hat{p}_{own1} is a similar probability in Cap 1.

Table A.1. Summary of Matching of the Experimental and Auction Data

Condition	Searches		Searches	% of eligible ads in			
	Total	Matched	Matched (%)	ML1	ML2	ML3	ML4
Cap 0	3,162,615	1,506,827	47.6	30.6	9.2	4.6	2.4
Cap 1	6,342,073	3,568,054	56.3	41.8	11.2	4.7	2.9
Cap 2	6,338,914	3,568,918	56.3	41.8	10.9	6.1	3.2
Cap 3	6,348,311	3,577,819	56.4	41.9	11	5.7	4.1
Control	22,209,220	12,506,083	56.3	41.9	11	5.7	3.5

Figure A.1. Frequency of Ads in Mainline 1 for 2,517 Most Popular Brand Queries

We thus confirm that the occasions of eligible own brand ads in mainline 1 in Cap 0 are correlated with the probability to get a click on the own brand weblink. This restricts us from using the eligible advertisements occasions to compare Cap 0 and Cap 1. Instead, we focus only on companies that have a paid link in mainline 1 more than 90% of the time. For these companies, a comparison of Cap 0 and Cap 1 gives a causal effect of advertisement.

Figure A.1 shows that around 50% of the companies advertise at least 10% of the time, with around 33% advertising more than 90% of the time. Restricting the analysis to the latter group gives us 824 companies that always advertise on their keyword.

Endnotes

¹ Focal brands in our sample always occupy the top organic slot in brand search results.

² It is exceedingly rare that a brand bids and does not occupy the top slot. This is due to the bidding behavior of focal brands: High relevancy and CTR are rewarded by the scoring function in the GSP auction.

³ For example, while the travel website Expedia is free to bid on “priceline,” it cannot include, for example, a “better than priceline” phrase in their ad text.

⁴ Our study focuses on clicks. Past work has consistently linked clicks to conversions, though the conversion rate may vary by position on the page (Rutz and Bucklin 2011, Agarwal et al. 2011, Goldman and Rao 2014).

⁵ One paper, Yang and Ghose (2010), sharply diverges from this result. Using a structural model, they assert that there is a positive interdependence between organic ranking and search CTR.

⁶ We do not, however, observe brands occupying the lower advertising slots. The context of brand search is somewhat different than generic product search, so we would not expect all of the predictions to be borne out.

⁷ dmoz.org. The project uses volunteer annotators to “classify the web.”

⁸ In these occurrences, the searcher often visits all advertisers, making it less interesting to study. Furthermore, search engines often refund clicks from such patterns.

⁹ With this selection rule, we are balancing the number of firms against the inclusion of brands that do not provide meaningful information because they are so small. We have done substantial robustness testing around this threshold, and there is no material impact on the results.

¹⁰ We keep companies with more than 80 exposures in each condition. Results are not sensitive to the choice of the threshold.

¹¹ We perform a series of standard tests for normality, including Shapiro–Wilk, Jarque–Bera, D’Agostino and other tests. All of them reject normality of the distribution. See Figure S4 in the online appendix for the empirical CDF of brand-specific estimates and further confirmation of these tests.

¹² It is exceedingly rare for the focal brand to occupy slots 2–4. Auction data reveal that if they choose to bid, they win the auction easily.

¹³ This is higher than the estimates for all 824 companies; this is because brands that face competitors tend to be less prominent; for them the effect of brand search advertising is stronger, leading to higher crowd-out.

¹⁴ We measure brand capital as the log(U.S.) rating on Alexa. The average brand capital of 35 brand keywords where a high CTR focal brand does not advertise but competitors do is lower (Average log(U.S.) ranking = 10) than the average brand capital of 564 brand keywords where both the focal brand and competitors advertise (Average log(U.S.) ranking = 8.86).

¹⁵ Competitors are present less than 20% of the time in control conditions; the results are not sensitive to this threshold.

¹⁶ This refers to companies that have their own brand advertisement in mainline 1 more than 99% of the time.

¹⁷ The difference in estimates is 0.0105, with a standard error of the difference being 0.0043, which corresponds to a *t*-stat of 2.46.

References

- Agarwal A, Hosanagar K, Smith MD (2011) Location, location, location: An analysis of profitability of position in online advertising markets. *J. Marketing Res.* 48(6):1057–1073.
- Blake T, Nosko C, Tadelis S (2015) Consumer heterogeneity and paid search effectiveness: A large scale field experiment. *Econometrica* 83(1):155–174.
- Chiou L, Tucker C (2012) How does the use of trademarks by third-party sellers affect online search? *Marketing Sci.* 31(5):819–837.
- Craswell N, Zoeter O, Taylor M, Ramsey B (2008) An experimental comparison of click position-bias models. *Proc. 2008 Internat. Conf. Web Search Data Mining* (ACM, New York), 87–94.
- Desai PS, Shin W, Staelin R (2014) The company that you keep: When to buy a competitor’s keyword. *Marketing Sci.* 33(4):485–508.
- Dupret GE, Piwowarski B (2008) A user browsing model to predict search engine click data from past observations. *Proc. 31st Annual Internat. ACM SIGIR Conf. Res. Development Inform. Retrieval* (ACM, New York), 331–338.
- Goldman M, Rao JM (2014) Experiments as instruments: Heterogeneous position effects in sponsored search auctions. Working paper, University of California San Diego.
- Granka LA, Joachims T, Gay G (2004) Eye-tracking analysis of user behavior in WWW search. *Proc. 27th Annual Internat. ACM SIGIR Conf. Res. Development Inform. Retrieval* (ACM, New York), 478–479.
- Jerath K, Ma L, Park Y-H, Srinivasan K (2011) A “position paradox” in sponsored search auctions. *Marketing Sci.* 30(4):612–627.
- Jeziorski P, Segal I (2015) What makes them click: Empirical analysis of consumer demand for search advertising. *Amer. Econom. J.: Microeconomics* 7(3):24–53.
- Kumar R, Tomkins A (2010) A characterization of online browsing behavior. *Proc. 19th Internat. Conf. World Wide Web* (ACM, New York), 561–570.
- Lewis R, Nguyen D (2015) Display advertising’s competitive spillovers to consumer search. *Quant. Marketing Econom.* 13(2):93–115.
- Narayanan S, Kalyanam K (2015) Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Sci.* 34(4):612–627.
- Reiley DH, Li S-M, Lewis RA (2010) Northern exposure: A field experiment measuring externalities between search advertisements. *Proc. 11th ACM Conf. Electronic Commerce* (ACM, New York), 297–304.
- Rutz OJ, Bucklin RE (2011) From generic to branded: A model of spillover in paid search advertising. *J. Marketing Res.* 48(1):87–102.
- Sayed A, Jerath K, Srinivasan K (2014) Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising. *Marketing Sci.* 33(3):388–407.
- Thisse J-F, Vives X (1988) On the strategic choice of spatial price policy. *Amer. Econom. Rev.* 78(1):122–137.
- Yang S, Ghose A (2010) Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence? *Marketing Sci.* 29(4):602–623.
- Yang S, Lu S, Lu X (2013) Modeling competition and its impact on paid-search advertising. *Marketing Sci.* 33(1):134–153.