The Effects of Search Advertising on Competitors: An Experiment Before a Merger

Joseph M. Golden John J. Horton*

October 10, 2019†

Abstract

We report the results of an experiment in which a company, "Firm Vary," temporarily suspended its sponsored search advertising campaign on Google in randomly selected advertising markets in the US. By shutting off its ads, Firm Vary lost customers, but only 63% as many as a non-experimental estimate would have suggested. Following the experiment, Firm Vary merged with its closest competitor, "Firm Fixed." Using combined data from both companies, the experiment revealed that spillover effects of Firm Vary's search advertising on Firm Fixed's business and its marketing campaigns were surprisingly small, even in the market for Firm Vary's brand name as a keyword search term, where the two firms were effectively duopsonists.

1 Introduction

Firms that advertise would like to know if their ads are effective. Any firm that advertises in a sufficiently large number of distinct markets—and that can measure where customers or sales originate—can credibly assess the effectiveness of its ads by running an experiment, suspending campaigns in some markets while maintaining the status quo in others (Lewis et al., 2011). However, a firm by itself typically cannot know the effects of its advertising on competitors, as competitors are not

*Collage.com and MIT, respectively.

†We thank Susan Athey, Tom Blake, John Bound, Charles Brown, Daisy Dai, Avi Goldfarb, Dan Hirschman, Jason Kerwin, Eric Schwartz, Isaac Sorkin, Frank Stafford, Hal Varian, Steve Tadelis, Brad Shapiro, Zach Winston, and Florian Zettelmeyer for feedback on this research. We also thank numerous employees at Firm Vary and Firm Fixed for their support of this research, including HA, NM, KN, RP, VS, ST and OT. Special thanks to Lindsey Trainor Golden for help with legal analysis. All views expressed in this paper are solely ours and should not be attributed to any other entity. All remaining errors are our own. Slides: http://www.john-joseph-horton.com/papers/search_ad_exp.slides.html

likely to share information with each other. And yet for firms that care about market share—such as firms in winner-take-all/winner-take-most industries—the effects of their advertising on competitors might be a key consideration in their marketing strategy.

In this paper, we report the results of an experiment in which one firm, which we call, "Firm Vary," temporarily suspended its sponsored search advertising campaign on Google in treatment group of randomly selected advertising markets in the US. Sponsored search advertising is a type of advertising in which a firm's ads appear next to "organic" search results when certain keywords are used in a query conducted on a search engine. In addition to exploring the effects of the experiment on Firm Vary's business, we can also explore the effects of this experiment on Firm Vary's closest competitor, which we call "Firm Fixed." This competitor perspective is possible because the two firms—Firm Fixed and Firm Vary—eventually merged, allowing us to combine their marketing data.[1]

Firm Vary and Firm Fixed were two-sided online marketplaces for services, connecting buyers and sellers. Prior to merging, these companies were fierce business competitors. As an indication of how similar their two product offerings were, at the conclusion of the merger, customers were simply migrated to Firm Fixed's platform and Firm Vary's platform was shutdown. Before merging, they were battling search advertisers, often targeting their ads at the same search terms on Google and other search engines. In their search ad campaigns, both firms bid on a variety of keywords related to their business; they also bid on each other's brand name. We will frequently have reason to distinguish between "brand" advertising (i.e., the firms bidding on their own brand name or their competitor's brand name) and all other advertising, which we will refer to as non-brand advertising.

Prior to the experiment, Google searches for the term "Firm Vary" would generally show an ad for Firm Vary in the top position and an ad for Firm Fixed in the second position. Given the potential for business stealing—and the effects that bidding has on the cost paid by competitors—the competition on branded keywords is often characterized as a prisoner's dilemma (Desai et al., 2014). The prisoner's dilemma logic frequently appears in practitioner-focused articles on bidding on competitor brand keywords, without necessarily using that term.[2] In our setting, given

[1]The ability to measure the effects of advertising on a competitor is rare but not wholly unprecedented (Lewis and Nguyen, 2014).

[2]For example, Baadsgaard (2016) writes that a "con" of bidding on competitor keywords is the potential for retaliatory bidding by the competitor. He goes on to describe a past consulting

the potential cross-side networks effects in marketplace businesses, this prisoner's dilemma framing seems particularly apt—a gained customer might be worth \$1 in revenue, but a customer lost to a main rival might cost \$1 (or more). Regardless of the value of a customer, the first question is, empirically, to what extent do competitors affect each other's ad campaigns and the quantity of customers acquired?

We examine two sets of experimental effects: (1) the effects of Firm Vary's experiment on Firm Fixed's search advertising campaign, both brand and non-brand, and (2) the effect of Firm Vary's ads on its own business and Firm Fixed's business. Our experimental design is essentially identical to Blake et al. (2015) (BNT hereafter), who ran the same experiment with eBay's sponsored search ads, but without the ability to assess the effects on competitors.

Our experimental results for brand advertising are that when Firm Vary turned off its own ads, Firm Fixed's ads on the term "Firm Vary" moved into the top advertising position, as expected. However, surprisingly, Firm Fixed received nearly the same number of clicks in treated advertising markets as in the control—Firm Fixed did not measurably benefit, despite the increased prominence of its ads and having Firm Vary's ads out of the way. We show that had Firm Fixed (1) "inherited" the same click-through-rate as Firm Vary, (2) received some plausible fraction of Firm Vary lost clicks implied by pre-experiment campaign statistics, or (3) had a click-through-rate improvement implied by other literature (Goldman and Rao, 2016; Simonov et al., 2018), we would have more than enough to statistical power to detect an effect.

Our results strongly suggest that queries with the brand as a keyword were "navigational"—users are searching for the names of these sites in order to navigate to them rather than entering a URL and as such, Firm Fixed stood to benefit very little from a superior position.

For brand ads, it was easy to see the firms were competing on the same keywords—entering a search for one firm's name would bring up ads for both firms. For

arrangement in which he tried to negotiate a "cease-fire" between a client who had started bidding on a competitor's brand who later retaliated. Another example comes from Cummins (2018), writing on the blog by Wordstream.com, a Gannett company offering software and consulting to search advertisers, explaining that by bidding on competitors, you are "... basically starting a war" and that "your bidding on their brand names will make it more expensive for them to bid on their own name, the same goes the other way [making it] more expensive for you to convert on your own brand." Consistent with the negative sums nature of a prisoner's dilemma, there is even evidence that firms have engaged in "cooperation" to avoid it: see the "1-800-Contacts" case, in which the company was accused by the Federal Trade Commission of entering into a collusive agreement not to bid on each other's branded keywords (Federal Trade Commission, 2016).

non-brands, the situation is more complex, as both firms were bidding on literally thousands of keywords and so knowing, *ex ante*, the degree of competition was challenging. At the time the experiment was run, Firm Vary anticipated large changes in position for Firm Fixed because it thought that Firm Vary and Firm Fixed's search ad campaigns were aggressively competing on more or less the same keywords, as both were after the same customers. As our analysis is occurring several years later, post-merger, we are able to use some information not available at the time, including sources of market-level data on keyword competition (Decarolis and Rovigatti, 2018).

Comparing keyword-level data for Firm Vary and Firm Fixed, during the month of the experiment, Firm Vary competed on 5,672, out of the 30,000 keywords that Firm Fixed successfully bid on, or about 19% of Firm Fixeds keywords. The set of overlapping keywords was about 34% of Firm Varys total keywords. On average, Firm Vary's ads were above Firm Fixed's ads, though the fraction of overlapping keywords where Firm Vary was above Firm Fixed is only 58%. In short, they were clearly competing, with the most common cells in the joint distribution of positions by keyword were Firm Vary first and Firm Fixed second, and Firm Fixed first and Firm Vary second. However, they were not competing everywhere, and the ultimate impact of Firm Vary being out the way is ultimately an empirical question.

Our experimental results for non-brand advertising show that the effects of Firm Vary's exit were minimal. The only detectable effect was that Firm Fixed's ads moved up in average position by an amount consistent with the keyword level measure of overlap. There was no detectable change in the number of clicks Firm Fixed received, their cost per click or any other metric—all close to zero.

A key question is whether our close-to-zero effects are informative about various hypotheses about what Firm Vary's exit would do. To answer this question, we model the number of lost Firm Vary clicks and then "transfer" them to Firm Fixed under different assumptions. We then see whether we had sufficient power to detect the degree of transfer with that assumption, such as scaling the transfer by the degree of impression and keyword overlap. This analysis reveals that under various assumptions of the degree of transfer, we had sufficient power to rule out even a sizable fraction of these lost clicks going to Firm Fixed.

Our experimental results for business outcomes is that in treated advertising markets, Firm Vary has about 23% fewer customer registrations, or "signups," but essentially none of these lost customers went to Firm Fixed. The point estimate for

Firm Fixed customers is close to zero (slightly less than 1%): we can rule out Firm Fixed gaining more than 6% additional customers with 95% confidence. This implies not only that Firm Fixed did not obtain an appreciable number of Firm Vary's lost customers from paid clicks, but that they also did not receive more customers from users clicking on "organic" (i.e., unpaid) Firm Fixed search results as well.

In addition to estimating the total change in the number of signups, we can also calculate a measure of ad efficiency for Firm Vary. This calculation is possible because when a visitor arrives at a company's website from a search engine, the advertiser knows what link a visitor clicked on. In particular, they know whether the link was a sponsored search ad, or whether the link was an "organic" search result which occurred because of the search engine's algorithm. Companies pay for the former, but not the latter. The naive way of assessing ad effectiveness typically used by practitioners is to assume that all signups resulting from clicks on paid ads would not otherwise have occurred (Lewis et al., 2011). However, at least some of those users would have counterfactually just clicked on an organic result or otherwise found the site if there was no sponsored search ad shown. Comparing the causal effect of Firm Vary's search ads on its business with the naive measures of advertising efficacy, we find that the naive method overestimates the number of new customers who registered due to search ads, and thus overestimates the value of these ads. The experimental estimate is about 63% as large as the naive estimate. While this efficiency measure is still far from 100%, it is also far away from the near 0% that BNT find even for non-brand advertising.

Our key contribution relative to BNT is our ability to measure the effects of advertising on a competitor. With our two advertising campaigns observed over time, we can make predictions about what experimental effects "should" have been under various scenarios. For example, we can make predictions about how many clicks (brand and non-brand) or customers Firm Vary lost by turning off its advertising, and then see what we would have estimated as a treatment effect, had Firm Fixed received those lost clicks or customer signups. Using this prediction, our comparison to the actual experimental estimates can rule out various hypotheses. Using signups as a business outcome reflects the fact that both firms were focused on acquiring customer signups.[3]

[3] We do not explore the value of a customer (and hence the ROI for advertising), in part because we have no insight into this lifetime value question, but also because both companies viewed the winner-take-all/winner-take-most nature of their industry as making the value of any particular customer somewhat irrelevant during this phase of heated competition.

BNT shows that brand-based advertising is ineffective for eBay, which some have interpreted as a consequence of eBay's well-known brand. However, we also find brand advertising was probably ineffective for Firm Vary, a company far less well known than eBay: at the time the experiment ran, Firm Vary had less than 5% national brand recognition in the US. Furthermore, we are able to show that little if any of this "lost" traffic went to their close competitor, even though Firm Fixed moved up in position. BNT conjecture that the threat of customer poaching/business stealing might explain why a firm might rationally bid on its brand. Our paper shows that at least for Firm Vary, bidding to prevent poaching was likely unnecessary.

For non-brand search advertising, we reach different conclusions than BNT, albeit for explicable reasons. BNT find that even non-brand search advertising is ineffective at increasing sales. We find that although search ads are not fully efficient—some of the users that clicked on search ads and became new customers would have instead clicked on organic links and also become new customers—they are far from ineffective. Our best estimate is that paid search ads were 63% efficient, meaning that in the absence of ads, nearly two thirds of customers Firm Vary acquired would not have otherwise signed up.[4]

The main managerial insight offered by our paper is guidance for firms planning their search advertising and the allocation of marketing resources (Johnson, 2017).[5] The first insight is that sponsored search ads were effective at acquiring new customers, albeit not as much as a naive estimate would suggest. So for would-be ad buyers who find that the benefits they receive from those new customers exceed the costs, then sponsored search advertising is useful. The second insight is that there was no discernible business stealing in our context. To the extent this generalizes, it means that advertisers who think they are in a prisoner's dilemma with their competitors—both compelled to bid on their brand keywords, even though both would be better off not bidding—they very well might not be, even for not very well-known brands.

The paper proceeds as follows. Section 2 discusses sponsored search advertising. Section 3 describes the empirical context and our experiment design. Section 4 reports the results of the experiment on Firm Fixed's ad campaigns. Section 5

[4] That paid search ads are effective matches recent work, also from a field experiment (albeit not on Google), in the context of Yelp (Dai and Luca, 2016).

[5] Enormous sums spent by firms on digital advertising—estimated at $83 billion in 2017—of which search advertising makes up a large share.

reports the results of the experiment on Firm Fixed's and Firm Vary's businesses. We conclude in Section 6.

2 Background on search advertising

Unlike other forms of advertising, the intent of sponsored search advertising is fairly straightforward: the search engine shows ads to search engine users with a revealed commercial need (as evinced by their search query) that the advertiser might be able to meet.[6] What fundamentally distinguishes sponsored search advertising from more conventional sources of advertising is the ease of targeting (Goldfarb, 2014). The ads themselves are too short and too unimpressive to do much more than claim a product exists that might meet the customer's revealed need. These are not ads that are likely to persuade would-be customers directly (Ackerberg, 2001). However, some have modeled consumers as inferring firm quality or "fit" from relative position of an ad on the page (Athey and Ellison, 2011; Armstrong et al., 2009).[7]

2.1 How sponsored search works

When users search on Google, it generates two separate sets of ranked results related to the search term: organic search results and paid ads. Other popular search engines, such as Bing, work similarly. We will describe the system that determines which ads get displayed, in which positions and at what cost to the advertisers. Search engines sell their ads via real-time auctions in which advertisers bid on search terms. Varian (2007) and Edelman et al. (2007) provide a general overview and analysis of search ad auctions.

Google sells its search ads via "generalized second-price" (GSP) auctions. These algorithmic auctions happen in real-time, nearly instantly, triggered by each search. Google's ad inventory consists of potential ad positions in which to show an ad impression, up to some maximum number per page. This inventory is highly heterogeneous, as advertisers target their ad copy and their bids, which they submit in advance, to specific search terms.

[6] A long-standing question in economics has been what, precisely, is advertising "for" (Nelson, 1974; Schmalensee, 1978; Milgrom and Roberts, 1986; Kihlstrom and Riordan, 1984)—is it to convey information directly (i.e., facts about products and prices) or perhaps indirectly (i.e., signal something about quality)?

[7] There have been some attempts to analyze bidding behavior to understand valuations and willingness to pay for position (Börgers et al., 2013; Varian, 2007; Yao and Mela, 2011).

Although Google's unit of inventory is an impression, advertisers generally submit bids not on impressions, but on *clicks*. These bids are known as cost per click (CPC) bids. These bids are not, however, determinative of position, as Google computes a bid-modifying quality score for each advertisement in an auction. This score is a function of various quality metrics (Varian, 2007), including Google's estimate of the ad's click-through rate for a given position, which is the percent of users who see the ad that click on it. Search engines, including Google, generally do not make public their exact methods for "scoring" an ad.

Ads are positioned by the ranking of their quality-adjusted bids. When a user clicks on an ad in position i, the GSP mechanism determines that the advertiser pays the minimum amount (or slightly more) that would keep their ad's score just above the ad in position $i+1$. That is, the advertiser pays approximately the following per click:

$$\mathrm{CPC}_i = \frac{\mathrm{bid}_{i+1} \times \mathrm{score}_{i+1}}{\mathrm{score}_i}. \tag{1}$$

Advertisers seek to maximize their surplus from the search ad auction as a function of their bid, including the choice to not participate by not bidding.

2.2 Brand ads

Companies may generally bid on ads for their own and their competitors' trademarked terms, such as their brand name. For example, Coca Cola can bid on the term "Pepsi" and Pepsi can bid on the term "Coke." Some search engine marketing experts claim that trademark owners and competitors must bid aggressively on brand terms to block competitors from poaching their potential traffic.[8] However, others have argued that users entering brand names are often conducting navigational queries—users are searching for the names of these sites in order to navigate to them rather than enter in a URL. To wit, some of the most popular search queries are the names of popular websites, such as "Google" and "Facebook." These navigational queries indicate little commercial intent by users, and in many cases, there are no associated ads. Bidding on branding keywords is commonplace among top

[8]For example, see this blog post at Search Engine Land, a popular blog and resource for search engine marketing professions, which calls allowing a competitor to outbid you on your own trademarked terms an "obviously untenable situation": `http://searchengineland.com/how-to-protect-brand-keywords-for-less-121566`. Similarly, in the legal literature, Gervais et al. (2013) argues that trademark owners bid on their own terms to block their competition.

retailers—all of the top 15 online retailers except Costco bid on their brand name and occupy the first position.[9]

How common navigational queries are with respect to brands presumably depends on the nature of the site. For our firms, which are both platform marketplaces, users are obtaining services delivered over time, and so they might make more frequent website visits. However, they might make matches on platform and then rarely return, in contrast to a more traditional e-commerce site. As such, it is not clear *ex ante* whether navigational queries are more or less common than in other settings.

As some of our results are about brand advertising, a natural question is how important this kind of advertising is industry-wide. A proxy for the importance of brand advertising is how much of search engine revenue comes from this kind of advertising. How much of Google's present revenue comes from these brand searches is unknown currently, but at the last time for which the figure is publicly available (April 2004), it was 7%.[10] Revenues from trademarked/branded keywords as a share of Google's total revenue are plausibly higher today, as Google has both permitted and encouraged more advertising on trademarked terms over the years.[11] Additionally, Google has introduced and refined software tools, such as its "Keyword Planner," to suggest and aid in the discovery of relevant keywords, including trademarked terms, for advertisers to consider for their advertising campaigns.

3 Empirical context and experiment design

Firm Vary and Firm Fixed were both online marketplaces for services. Both firms used sponsored search advertising to acquire new buyers. Firm Vary was spending about $10 million per year on sponsored search advertising, and Firm Fixed was

[9] See Table 4 in Appendix A.1.

[10] Rosetta Stone Ltd. v. Google, Inc., 676 F.3d 144, 155-156 (4th Cir. 2012) (citing Joint Appendix, Vol. IX, Tab 41, Ex 6, "Google Three Ad Policy Changes" at p. 4264-4265). Rosetta Stone initially filed this case in 2009 and the parties settled in 2012.

[11] In the US prior to April 2004, Google allowed trademark holders to, upon request, block other advertisers from both advertising on their trademarked terms, and from including these terms in their ad text. Later in 2004, Google changed its policy to no longer allow trademark holders to block ads on their trademarked terms. Then in 2009, Google began to allow advertisers to include trademarked terms in their ad text under certain circumstances. Google's current AdWords Trademark Policy grants resellers and informational sites limited permission to use trademarked terms in their ad text. Google's current policy is available here: `https://support.google.com/adwordspolicy/answer/6118?hl=en`. Accessed September 2, 2017.

spending a similar amount. Firm Vary cared specifically about potential business being lost to Firm Fixed due to the cross-side network effects in their respective marketplaces and thus the potential for winner-take-all dynamics.[12] As such, Firm Vary historically bid on its own brand name—i.e., engaged in brand advertising—to keep Firm Fixed from poaching potential customers via search advertising. Before the experiment, Firm Vary was spending about 11% of its marketing budget targeting competitors' brand keywords directly, though it was only spending about 1% on its own brand keywords.

3.1 Degree of brand awareness and market definition

Firm Vary and Firm Fixed were not very well-known brands among the general population of Internet users when the experiment was run. We ran a Google Survey several months prior to the experiment to learn what fraction of the US population had heard of Firm Vary and Firm Fixed, as well as other benchmark firms. Both Firm Vary and Firm Fixed had little brand awareness—Firm Vary had 4.3% and Firm Fixed had 2.6%. By comparison, in the same survey, 47% of respondents reported recognizing "LinkedIn," the professional networking social network. Although we do not have comparable brand awareness data for eBay, given its age and size, it was likely considerably higher than even LinkedIn. This low brand recognition likely reflects that neither brand was consumer-facing, with both catering to small and medium-sized businesses. Unfortunately, we have no way of assessing brand awareness among this narrower set of would-be customers.

The interpretation of some of the results depends on the larger industry Firm Vary and Firm Fixed are in, as well as their competitors in the search advertising space. Product market definition is challenging, but a third-party report prepared for Firm Vary and Firm Fixed had nearly identical "industry" shares of the web page visits, with their cumulative fraction close to 50%.[13] The next nearest true competitor had less than 5%. According to this same report, Firm Vary and Firm Fixed were also both getting about 50% of the paid clicks in this "industry," approximately splitting the total evenly between themselves. However, as we will see,

[12] Sayedi et al. (2014) propose a more complicated game to model the relationship between poaching in search advertising, and spending on traditional advertising, such as television and newspaper ads. In our setting, both Firm Vary and Firm Fixed primarily engaged in online advertising.

[13] Some aspects of the report—and the identity of the third-party preparing the report—would reveal proprietary information, and so we keep it anonymous. However, the third-party in question would be *very* well-positioned to discuss these issues.

even though each firm was getting similar numbers of clicks, it does not mean they splitting clicks at the keyword/auction level.

3.2 Pre-experiment search ad campaigns

Post-merger, we learned that prior to the experiment, Firm Vary and Firm Fixed had broadly similar search ad campaigns. Figure 1 plots daily time series for the brand campaigns (bidding on the Firm Vary brand name) and non-brand advertising campaigns of Firm Vary and Firm Fixed prior to the experiment. The reported series are averaged by Designated Marketing Areas (DMA), which is the level at which campaign results are reported. There are 210 DMAs, which subdivide the country into regions and were originally designed for television-based advertising purposes.

To preserve some confidentiality, the values in Figure 1 are divided by the max value obtained during the period by either firm. Note that the y-axis is on a log scale to better show the data, though it is labeled in levels. In the bottom row of the figure, we report average ad position un-scaled/transformed. We reversed the sign on mean position to put the series in the "correct" order corresponding to the first position being at the top.

For brand ads—Firm Fixed bidding on the Firm Vary brand name and Firm Vary bidding on the Firm Vary brand name—in the left column of the figure, we can see that Firm Vary had more impressions than Firm Fixed on its ads for its own brand name. The differences are substantial—overall, Firm Fixed had only 26% of the impressions as did Firm Vary. However, this gap narrowed in the weeks leading up the experiment, approaching 69%. Despite modest differences in impressions between Firm Fixed and Firm Vary (at least in the week leading up to the experiment), there were enormous differences in the number of clicks. The Firm Vary CTR was nearly 10x the Firm Fixed CTR. However, Firm Fixed had a higher CPC, likely reflecting a quality-adjustment that raised their costs due to their low CTR. Firm Vary total spend on brand advertising was greater than Firm Fixed. With respect to position, in the bottom panel we can see that Firm Vary was first and Firm Fixed was second.

For non-brand ads (in the right column of Figure 1), both firms were receiving very similar numbers of impressions and paid clicks and were spending similar amounts. Firm Vary had a higher CTR and a slightly lower CPC. In the bottom row indicating average position, we can see that Firm Vary typically had a higher position, implying Firm Fixed would have more to gain from Firm Vary being out of the way, on average. The average positions are lower, but as we will see, for most of the overlapping keywords, the two firms were near the top and the average was dragged down by a tail of low-positioned ads.

Figure 1: Daily campaign attributes (normalized) for brand advertising (both firms advertising on the "Firm Vary" brand keyword) and non-brand advertising campaigns of Firm Vary and Firm Fixed prior to the experiment

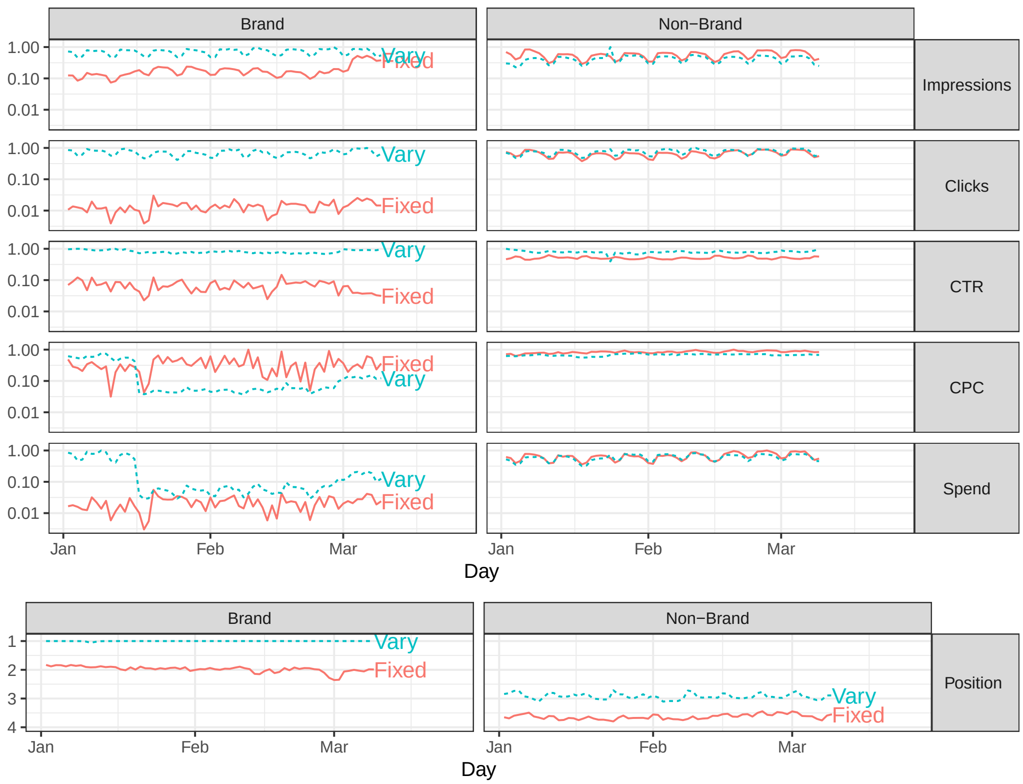

Notes: This plot shows daily normalized values for various brand and non-brand campaign attributes for both Firm Vary and Firm Fixed prior to the experiment. Each series by type (e.g., CPC, CTR) is divided by the max value obtained by either firm, for either type, during the panel covered by the data. The y-axis is on a log scale for all outcomes except for the position. The bottom panel shows the average position, which is not normalized. The brand campaigns are for the firms bidding the "Firm Vary" brand keyword.

3.3 Non-brand keyword competition

Although Firm Vary and Firm Fixed each viewed the other as their main competitor, they actually competed with a much larger number of advertisers that were interested in the same keywords, if not the same customers. In a report prepared for Firm Vary by a third-party, Firm Vary's ads appeared along nearly 10,000 other distinct "domains" i.e., other firms bidding on the same search terms and having their ads appear next to Firm Vary's ads.[14] This ad competition reflects the fact that many of the search terms Firm Vary bid on were also of interest to other firms that are not product market competitors with Firm Vary. For example, consider the search term "accounting"—searches containing this term could be of interest to firms directly offering accounting services, firms offering accounting software, authors selling books on accounting, marketplaces for accounting services, and so on.

To give a sense of how closely Firm Vary and Firm Fixed were competing over keywords, we analyze historical keyword data for our two firms, with data obtained from SEMRush. SEMRush has apparently been periodically "scraping" information on what shows up in response to a large number of keyword queries (Decarolis and Rovigatti, 2018). With the SEMRush data, we know, for each keyword, the display ads that appeared, the domain they linked to, and the ad position. We also know the timestamp when the data was collected, but not the actual DMA—we strongly suspect that SEMRush only collects data from a single DMA. Neither firm had DMA-specific advertising strategies, but there is a possibility that keyword data from the scraped DMA is not representative.

To explore what the campaigns looked like right before the experiment, we use data from "March 2014" (scraping for this month actually spanned from 2013-11-18 to 2014-01-25); the number of distinct keywords Firm Vary successfully bid on is 16,861. Although this technically overlaps with our experiment, it appears that the SEMRush scraper was launched from a control DMA, and so data is available for both Firm Vary and Firm Fixed. Of these, Firm Vary competed on 5,672, out of the (at least) 30,000 that Firm Fixed successfully bid on.[15] The SEMRush data does not show the degree of impression overlap, so these estimates are lower

[14] This figure comes from the same well-positioned third party described earlier.

[15] The number of Firm Fixed keywords was exactly 30,000 which suggests some kind of truncation on the part of SEMRush rather than Firm Fixed bidding on exactly this many keywords. In other months near the experiment, the number is slightly lower, suggest the 30,000 is likely a reasonable approximattion.

bounds—it could be the case that Firm Vary and Firm Fixed would co-appear in other presentations to users but happened to not co-appear in the particular scrape captured by SEMRush.

Using only overlapping keywords, we plot the positions for each keyword in Figure 2, "jittering" points to prevent overplotting. The position of Firm Vary's ad for that keyword is the x-axis, and the position of Firm Fixed's ad for that keyword is the y-axis. Along the diagonal, there are no keywords, as each firm can only have one position in a particular scrape instance. At the margins, histograms indicate the fraction of keywords in each position, along with percentage for that position (these histograms reflect the jittering, which is why they are not completely flat within a position level).

Figure 2: Non-brand keyword overlap among Firm Vary and Firm Fixed prior to the experiment

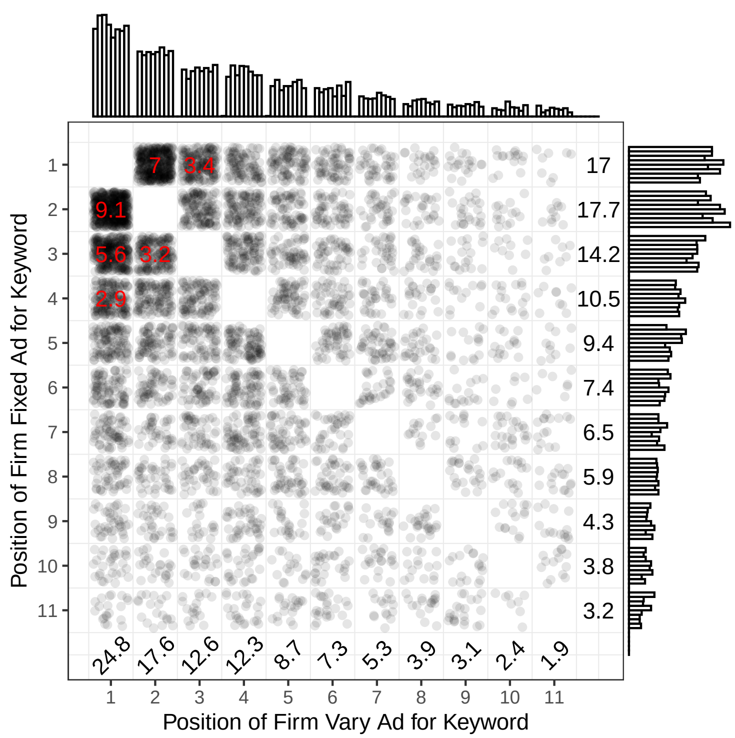

Notes: This figure shows, for the collection of keywords bid on by both Firm Fixed and Firm Vary, the joint distribution of positions on the page.

From the marginal distributions, we can see that for both firms, the most common position for a keyword is to be in the first position, the second most common is the second position, third most common is in the third position, and so on. Among overlapping keywords, for Firm Vary, 24.8% of its keywords were in the first position; for Firm Fixed, 17% of keywords were in the first position.

As we can see, much of the point mass in the joint distribution is where Firm Fixed and Firm Vary are in the top two positions. For "cells" with more than 2.5% of the total joint point mass, the percentage is labeled. We can see that 7% of all overlapping keywords had Firm Fixed in the first position and Firm Vary in the second position; 9.1% of all overlapping keywords had Firm Vary in the first position and Firm Fixed in the second position. Firm Fixed had something to gain, on average, from Firm Vary being out of the way.

It is important to note that Firm Vary did not use SEMRush data to plan the experiment or forecast effects. Our collection and analysis of this keyword data was done five years later. However, this keyword data would have proven quite useful in forecasting the size of experimental effects, especially with respect to the change in Firm Fixed's non-brand ad position, as we will show.

3.4 Experimental design

The design of the experiment is simple. During the experiment, Firm Vary shut off all of its Google search ads in half of the DMAs in the United States for a period of 28 days, starting on March 11, 2014. Advertisers can target Google ads geographically by DMA, which is what enables the experiment. Figure 3 shows the US DMAs in the experiment and their treatment assignments, with control DMAs in white, and treatment DMAs in gray. Treatment and control DMAs were selected at random. In Appendix A.2, we show that the DMAs are well-balanced with respect to pre-treatment ad campaign attributes.

After the experiment, Firm Vary resumed bidding on search ads in the entire US, as it had before the experiment. Prior to the experiment, Firm Vary did not target or vary search advertising purchases within the US geographically.

3.5 Internal validity

The internal validity of the experiment would be jeopardized if Firm Fixed reacted to Firm Vary's experiment. In particular, Firm Fixed could have potentially disrupted

Figure 3: Treatment and control direct marketing areas (DMA)

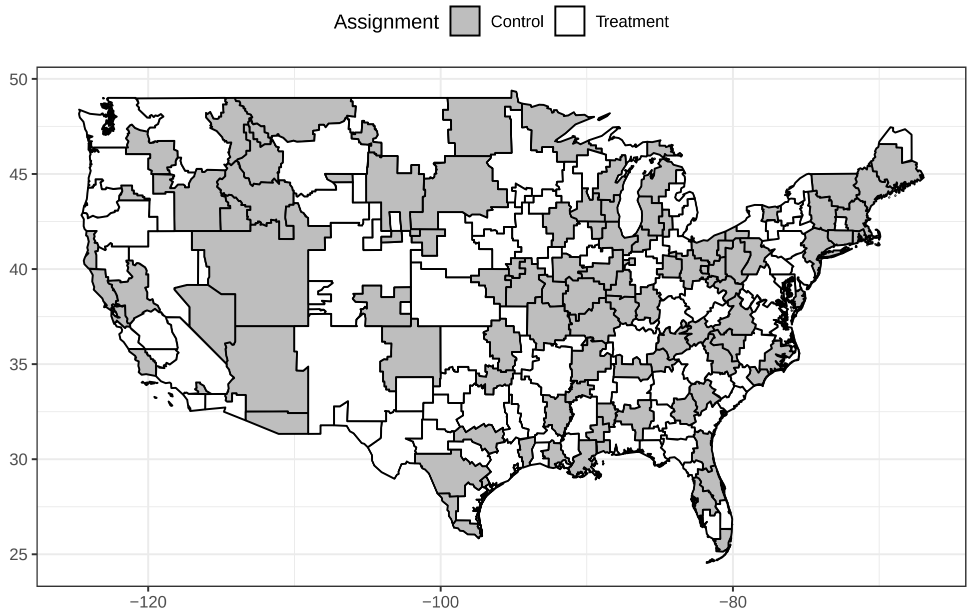

Notes: This figure shows the US directed marketing areas (DMA) and their allocation to either the treatment—in which Firm Vary turned off all search advertising—or the control, where Firm Vary kept its search advertising campaigns unchanged.

the experiment by changing its bidding behavior in the treatment and control DMAs during the experiment. For this reason, Firm Vary did not announce this experiment publicly.[16]

Even without being told, Firm Fixed could have learned that something had changed, jeopardizing internal validity. Fortunately during the experiment, Firm Fixed neither changed its bidding behavior overall nor did it specifically target its bids geographically within the US. The only changes Firm Fixed made to its bids during the experiment in the US applied to the entire US. These changes were minimal, and followed an overall bidding strategy that did not change during the experiment.

Coincidentally, the two companies concluded the process of merging during the

[16]Prior to the conclusion of the merger, the companies needed to operate as separate, competing entities, and as such, Firm Vary did not inform Firm Fixed about the experiment.

experiment, one week prior to the end of the experiment. In preparation for the anticipated conclusion of the merger, neither Firm Vary nor Firm Fixed made any substantial changes to their bidding during the experiment period prior to the merger (aside from Firm Vary running this experiment). Following the conclusion of the merger, no major changes were made to Firm Fixed's bidding strategy until after the end of the experiment period. This was done both to facilitate completion of the experiment and to allow the new company enough time to formulate an updated search advertising strategy.

3.6 Sample size and usable data

The experiment's duration was limited by business concerns. These concerns proved to be well-founded, as we will show. A simulation-based *a priori* power analysis suggested that the experiment should run for a minimum of two weeks, a duration expected to yield usable results to assess the overall performance of Firm Vary's ad campaign with respect to signups.[17] The realized sample is somewhat smaller than expected due to two independent problems, both unrelated to the experiment or its results, but which coincidentally occurred on consecutive days during it. As such, there is a gap of seven days (the "Omitted Period") without useful data during the experiment, so the effective experiment length is 21 days.[18]

3.7 Predicted effects of the experiment on Firm Fixeds ad campaigns

We predicted that Firm Vary's experiment could affect Firm Fixed's ads with respect to: (1) the position of ads, (2) the number of impressions received, (3) the clicks received, (4) the cost per click and (5) the total campaign cost. Table 1 describes the metrics and our predictions of the effect Firm Vary's experiment would have on Firm Fixed's campaign, both for brand and non-brand keywords.

[17]As noted across a variety of ad campaigns by Lewis and Rao (2015), attaining a large enough sample size to achieve the statistical power necessary to evaluate an advertising campaign can be challenging.

[18]The first problem is unexpected Firm Vary bidding behavior (as described by the Firm Vary marketing department) that occurred for three days resulting in substantially fewer ad impressions those days. The second is a denial of service attack against Firm Vary which resulted in four days of lost or unusable data. Days in which either of these two problems occurred, along with the day before the start of the experiment (during which Firm Vary tested out the experiment for part of the day) are omitted from our analysis.

Table 1: Description of advertising metrics and predicted effects of Firm Vary's experiment on Firm Fixed's Google sponsored search ad campaign.

Outcome	**Prediction for Brand ads**	**Prediction for Non-Brand ads**	**Description**
Position	Decrease by ≈ 1	Decrease by < 1	The average position of the ad on the search results page. 1 indicates the top position on the page, which is most likely to get clicked.
Impressions	Remain the same	Increase slightly	Count of total times an ad is shown.
Clicks	Increase substantially	Increase substantially	Count of total times an ad is clicked on.
Cost-per-click (CPC)	Remain the same	Decrease substantially	Total clicks divided by total cost.
Cost	Increase substantially	Unsure	Total cost of the ad campaign.

This table shows the predicted effects of turning Firm Vary's search ads off on Firm Fixed's brand ads (i.e., ads where the keyword is "Firm Vary") and the non-brand Firm Fixed ads, for each of five search ad metrics.

Position. As such close business competitors, we believed that Firm Fixed and Firm Vary were fierce competitors for search ads. For Firm Vary's brand ads, since Firm Vary was in position 1 and Firm Fixed was in position 2 generally, we expected Firm Fixed's position to decrease by 1 i.e., move up the page, becoming the first result.

For non-brand ads, although the direction of the effect should be the same, the magnitude of the effect on position is ambiguous. For any particular non-brand keyword, Firm Fixed's position would decrease (i.e. improve) by exactly one if Firm Vary's ad would outrank Firm Fixed's, and otherwise would remain unchanged. If they were, on average, "tied" then we would expect an average position change of about 0.5.

As the two companies did not target the exact same set of keywords, we expected non-brand position to improve by somewhat less than 0.5, but still by a substantial amount. The effect would be smaller if Firm Vary's ads were generally below Firm Fixed's ads when they both appeared. With access to the SEMRush data, we know (post experiment) that the actual overlap and the fraction of keywords where Firm Vary was above Firm Fixed implies an improvement of just -0.071 spots on average.

Impressions. Firm Fixed's ad impressions count would generally go up in cases where the Firm Vary ad would have been the worst-ranked ad shown, and Firm Fixed's ad was the highest ranked ad that was not shown. In other words, impressions would increase for Firm Fixed because it would be included in search results where they counter-factually would not have been because of the presence of Firm Vary's ads. However, this specific situation is rare overall (given the average position of non-brand ads for both firms, as we saw in Figure 2), and was non-existent for the brand ads (as they were in positions 1 and 2), so we expected a very small overall increase in non-brand Firm Fixed impressions, but no increase in impressions for brand ads.

Clicks. Clicks generally increase with more impressions and better position on the page, so we expected clicks on non-brand Firm Fixed ads to increase overall, primarily due to their improved position and somewhat due to an increase in impressions. For brand ads, we expected clicks to increase entirely due to the improvement in position from 2 to 1. Ads in position 1 generally receive substantially more clicks than ads in position 2 (Jansen et al., 2013), so we expected a particularly large increase in clicks on Firm Fixed's brand ads. Goldman and Rao (2016) report on the effects of movement from position 1 to position 2 for an "off-brand" ad (i.e., in our context, Firm Fixed's ad appearing on a query using the "Firm Vary" brand name) reduces CTR to 60% of the 1 position CTR. We further expected Firm Fixed to get more clicks due to the absence of Firm Vary's ad, independent of the effects on clicks due to position and impressions.

Simonov et al. (2018) offers experimental evidence on the effects of position with respect to brand advertising. Using experiments conducted by Bing that altered the number of paid ad slots available on an auction level, they study the flow of clicks to organic results and competitors. They find that for a focal brand holding the top position, having competitors below can siphon off a substantial number of clicks (1-4%) They cannot experimentally replicate our scenario of the focal brand leaving the market—the "focal" brand not showing up at all, though they can do a non-experimental analysis. An important difference that Simonov et al. (2018) highlight is that most of the competitor firms hoping to siphon off traffic are far, far smaller than the focal brand. The situation is quite different from our setting of two equally-sized rivals with similar overall campaigns.

Given the importance of clicks to advertisers and the fact that a better position is costly, several studies have focused on analyzing the effects of position in the sponsored search context. "Micro" empirical studies of click behavior show that position clearly matters; but empirical reality does not closely match a model of consumers as cascading sequentially from top to bottom, with ads in other positions being irrelevant (Jeziorski and Segal, 2015; Gomes et al., 2009). However, as a stylized fact from the literature, it is well-established that click through rates decline in position—Ghose and Yang (2009), analyzing data for a single retailer bidding on multiple keywords, find that position and click through rates are negatively correlated.

Cost-per-click (CPC). For brand advertising, the CPC for Firm Fixed was not set by Firm Vary, which was occupying the first position—it was set by the ad in the third position (or whatever the reservation price was). As such, we expected Firm Vary's experiment to have no effect on Firm Fixed's brand ad CPC.

For non-brand advertising, we expected CPC would go down whenever Firm Vary's ads would otherwise have occupied the ad position one below (i.e. worse) than Firm Fixed's, and otherwise Firm Fixed's CPC would remain unchanged. While the details of the auctions are important, we can think of CPC as a price subject to the forces of supply and demand (Goldfarb and Tucker, 2011).[19] The exit of Firm Vary is thus a negative demand shock on the keywords Firm Vary was competing for. We thus expected Firm Fixed's CPC to go down overall for non-brand advertising, but only to the extent that Firm Vary's ad was determinative of the price paid by Firm Fixed. As with clicks, despite the general prediction that Firm Fixed's costs should fall for non-brand advertisements, the magnitudes of these effect sizes depend on how aggressively Firm Fixed and Firm Vary were competing with each other and

[19] Goldfarb and Tucker (2011) exploit a natural experiment—laws regarding the advertising by "ambulance chaser" lawyers—to show that in states where some offline channels are forbidden, related sponsored search advertising terms are about 5% to 7% higher, showing that at least at an industry level, prices are sensitive to demand.

with other search ad campaigns. If we use the SEMRush keyword data, the fraction of overlapping keywords where Firm Vary would be "marginal" for Firm Fixed is 18.7%. However, we know little about how this would translate into price, as it depends on factors we do not observe.

Costs. We made no prediction about the overall cost of Firm Fixed's ad campaign. For non-brand ads, we expected more clicks because of higher position (increasing cost), but lower CPC (reducing cost), and had no general expectation about the relative size of these two effects. For brand ads, we expected CPCs to stay constant (as Firm Vary's ads were not setting the price for Firm Fixed's ads), but we expected costs to increase substantially with the increase in clicks, because of the improved position of Firm Fixed's ads.

3.8 Predicted effects on Firm Vary and Firm Fixed customer signups

As both Firm Vary and Firm Fixed primarily used search ads to attract new buyers to their marketplaces, the number of new customer signups is the metric we use to quantify the effects of Firm Vary's ads on both businesses. For Firm Vary signups, we can also estimate ad efficiency, as we know the click origin of signups.

Given that organic search results seem like obvious substitutes for paid advertisements—and hence are a relevant consideration for any would-be advertiser—there is a literature focusing on the interplay between organic and paid search advertising in customer acquistion. Using data from the keyword search advertising campaign of a single retailer, Agarwal et al. (2015) find that organic search results are substitutes for keyword search advertisements, but have a complementary effect on revenue because the organic results improve click-through rates. Yang and Ghose (2010) also present evidence of complementarities between paid and organic listings. Animesh et al. (2011) consider competition between rivals in the online sponsored search market.[20]

For Firm Vary, we expected signups to decrease substantially, but by somewhat less than the number of new customers who clicked on a search ad. The reason is that some of those customers who came through ads would have, in the absence of ads, come by clicking on an organic link. For Firm Fixed, we expected signups to increase in DMAs where Firm Vary ceased advertising, but by an unknown amount.

[20] They report a field experiment in which a retailer varied their ad creative and position rank.

Firm Vary dropping out of the ad auction would have to at least weakly help Firm Fixed's business by increasing its ads' exposure and number of clicks. Based on our predicted effects on Firm Fixed's overall ad campaign as described in Table 1 and a pre-merger estimate by Firm Vary that search ads accounted for a large—but non-majority share—of Firm Fixed's new customer signups, we expected an increase in this metric for Firm Fixed.

Even if Firm Fixed siphoned off a large fraction of traffic that "belonged" to Firm Vary, it is an empirical question how many would turn into customers. Many of these poached clicks from brand ads could be of low quality and not show up in customer signups. Simonov and Hill (2018) offers experimental evidence that this is the case and that many "poached" clicks quickly lead to the user hitting "back" suggesting they were looking specifically for the focal brand.

4 The Experiment's Effects on Firm Fixed's advertising campaigns

To begin our analysis of the experiment, we simply compare the daily time series of Firm Fixed's search advertising outcomes before and during the experiment. Figure 4 plots the daily difference between the treatment DMAs and control DMAs for a collection of outcomes, for both brand ads (dashed line) and non-brand ads (solid). The data show 40 pre-period days and the full post experiment period.[21] The start of the experiment is indicated with a vertical line, and the x-axis are days relative to this date. For position, the outcome is the difference in average impression-weighted positions. For all other measures, the difference in the log sum of that measure is plotted.

Starting with position—shown in the top panel of the figure—we can see that before the experiment, the treatment and control DMAs show no gap for ads overall, and no systematic gap for brand ads, though because of the smaller samples, there is more day-to-day variation. After the experiment began, Firm Fixed's brand advertisements move up by 1 (a decrease in position), as expected. However, there is no discernible effect on the position for non-brand advertisements, though as we will show, a regression-based approach does indicate an effect.

In the remaining panels of Figure 4, the outcomes are the log cumulative number

[21] The number of pre-period days to include is of course somewhat arbitrary—we vary the number as a robustness check and find no substantive differences in effects.

Figure 4: Difference in average daily Firm Fixed search ad campaign outcomes by DMA treatment assignment

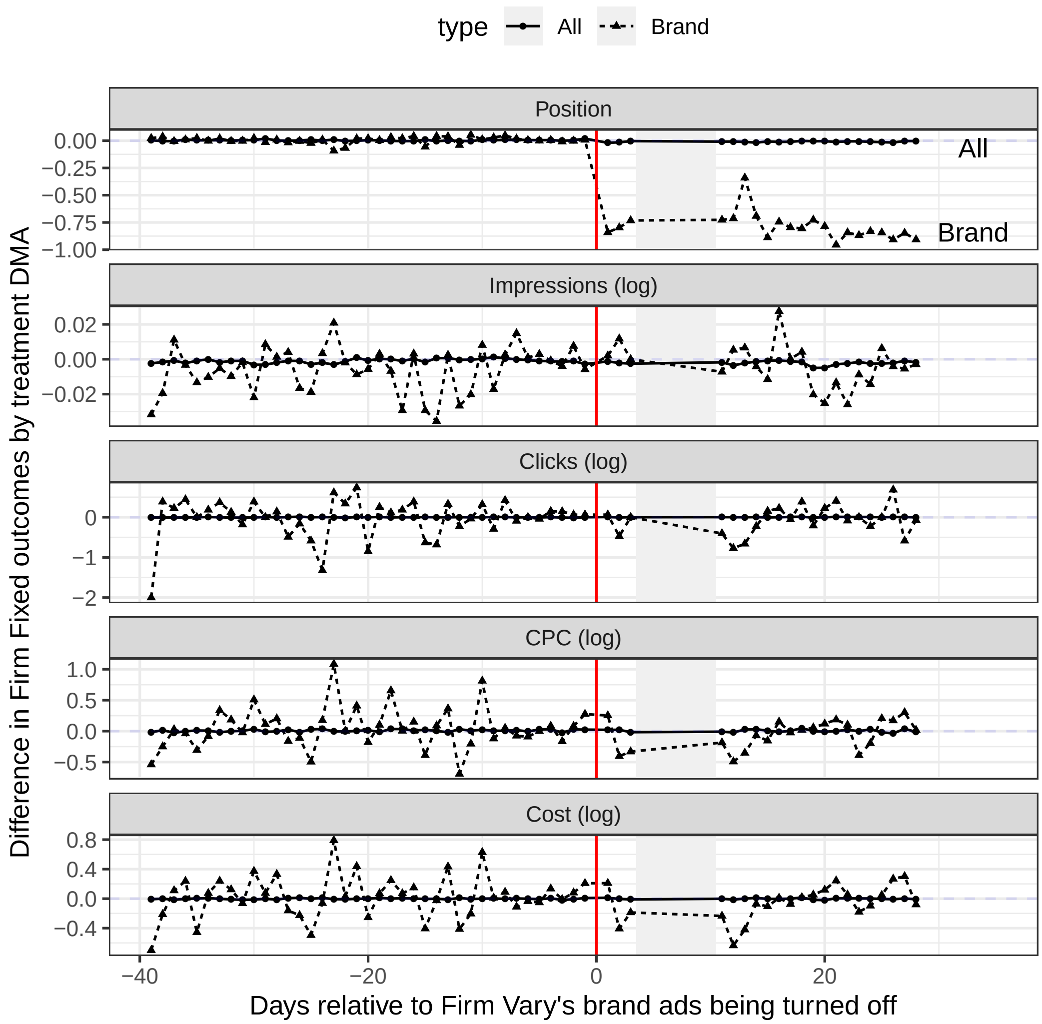

Notes: This figure shows the by-day gap in search ad outcome metrics for Firm Fixed, for both brand ads i.e., ads appearing when the search term was "Firm Vary" (indicated with a dashed line) and for all, non-brand search advertising (solid line). For position and CPC, the outcome is difference in impression-weighted averages. For all other measures, the difference in the log sum of that measure is plotted. The vertical red line indicates when Firm Vary stopped advertising and where it continued. Data are omitted for the seven day "Omitted Period," as described in Section 3.6, and indicated by gray rectangles.

of impressions, the log number of paid clicks, the log cost per click, and the log of the total cost. What is striking across panels is how little evidence there is for any kind of treatment effect for either brand or non-brand search advertising campaigns. We will explore these effects more formally with regressions, but there is little evidence that Firm Vary's departure from the market did much of anything to Firm Fixed.

To gain precision by accounting for pre-experiment differences across the different DMAs, and to use a more appropriate transform for the outcome, we switch to a regression framework. We evaluate all results using a difference-in-difference approach via the following regression:

$$Y_{it} = f\left(\beta_1 \textsc{AdsOff}_{it} + \delta_t + \gamma_i + \epsilon\right), \tag{2}$$

where Y_{it} is the outcome variable, i indexes the 210 different DMAs, t indexes time periods, $f(\cdot)$ is a link function and $\textsc{AdsOff}_{it}$ is an indicator for whether Firm Vary had its ads turned off in DMA i at time t. The time and DMA fixed effects are, respectively, δ_t and γ_i. We aggregate the results into two time periods: "before" and "during" the experiment. We also use different specifications, including by-day outcomes in Appendix A.3. In all results, we cluster standard errors at the DMA level.

For variables where we wish to estimate a percentage change, we use a Poisson model, with $f(x) = exp(x)$ as a link function and use QMLE.[22] We prefer this estimator to taking the log of Y_{it} because some of our outcome observations are equal to 0.[23] For the ad position, where we expect a linear change, we use $f(x) = x$, i.e., just use the average position as the outcome, and estimate via OLS.

The effects of Firm Vary's suspension on Firm Fixed's advertising campaign using Equation 2 are reported in Table 2. The top panel reports brand campaign estimates, while the bottom panel reports non-brand campaign estimates. In all cases, the treatment group is the set of DMAs where Firm Vary turned its ads off.

[22] Silva and Tenreyro (2006) and Wooldridge (2002) describe and motivate the use of this estimator. This estimator does not assume that $Var(Y|X) = E(Y|X)$, as the name Poisson might misleadingly suggest, for consistency or asymptotic normality, and it has nice efficiency and robustness properties (Wooldridge, 2002).

[23] In practice, for some outcome variables there are few observations equal to zero, and in these cases, we get similar results when we estimate using OLS, and either use $log(Y_{it}+1)$ as our outcome variable, or drop observations where $Y_{it} = 0$.

Table 2: Effect of Firm Vary's ad campaign suspension on Firm Fixed's sponsored search advertising campaigns

Brand ads:

	Position	Impressions	Clicks	CTR	CPC	Cost
	(1)	(2)	(3)	(4)	(5)	(6)
$\textsc{AdsOff}_{it}$	-0.931^{**}	0.009	-0.184	0.001	0.230^{*}	$-0.171^{\dagger}$
	(0.017)	(0.049)	(0.119)	(0.003)	(0.104)	(0.093)
N	408	420	420	408	241	420

Non-brand ads:

	Position	Impressions	Clicks	CTR	CPC	Cost
	(1)	(2)	(3)	(4)	(5)	(6)
$\textsc{AdsOff}_{it}$	-0.052^{**}	-0.018	0.028	0.0004^{*}	-0.070^{**}	-0.007
	(0.007)	(0.018)	(0.020)	(0.0002)	(0.025)	(0.019)
N	408	420	420	408	408	420

Notes: Each column shows the estimated impact of Firm Vary shutting down its Google ads on different aspects of Firm Fixed's Google ad campaign. All estimates include DMA and time-period fixed effects and are estimated using Equation 2. Standard errors are clustered at the DMA level and are in parentheses. Columns (1) and (5) contain fewer than 420 samples because position and CPC are only defined if there are any impressions and clicks, respectively, in a time period-DMA observation. Significance indicators: $p \leq 0.10 : \dagger$, $.p \leq 0.05 : *$, and $p \leq .01 : **$.

4.1 Experimental Effects on Firm Fixeds Brand campaigns

Starting in the top panel, labeled "Brand ads", as expected, Firm Fixed's ad position for the keyword "Firm Vary" improved by almost exactly 1, which we can see in Column (1)—the coefficient is negative, as going from the second position to the first position is a decrease. This was a movement of Firm Fixed from the second position to the first position (recall the brand ad position results from Figure 1). The sample size is not exactly 420 (210 DMAs x 2 periods) because some DMAs had no impressions for brand ads; the CPC regression sample size is also smaller because we only observe CPCs if a click occurs. From Column (2), we see that impressions remained almost exactly the same, as predicted.

For clicks, the prediction was that they should increase due to the improved position and Firm Vary's ads being out of the way. However, from Column (3), we can see that this prediction was not borne out—the point estimate is negative (-18%), though not statistically significant. In Column (4), the outcome is the click-through rate, or CTR. This was expected to rise, and while it is positive, the effect is closer to zero and not conventionally significant.

An obvious question is whether the results above are surprising, given the degree of statistical power. We can explore counterfactuals for these click results by predicting the number of lost clicks by Firm Vary and then "transferring" them to Firm Fixed under various assumptions. As we have no reason to think Firm Fixed would lose clicks from Firm Vary's exit, simply adding clicks is unlikely to double count, given our near-0 estimates on the effects on clicks. We will discuss the various counterfactuals in depth below, but the main conclusion from this exploration is that although each of the six counterfactuals would seem to be reasonable predictions for the effects on clicks, in fact, the actual effect was statistically and substantively smaller than what any of these counterfactual assumptions would have predicted.

To begin exploring counterfactuals, we first have to make an imputation for lost clicks. Our imputation of lost Firm Vary clicks is done fitting a QMLE model of Firm Vary brand clicks with DMA and Day fixed effects (mirroring Equation 2) using DMAs and days for which the treatment was not "on":

$$\textsc{Clicks}_{it} = f\left(\delta_t + \gamma_i\right) | \textsc{AdsOff}_{it} = 0$$

and then predicting $\widehat{\textsc{Clicks}}_{it}$ for the Firm Vary treated DMAs. With this specification, the DMA specific effects are learned from the pre-experiment period and the day effects are learned from the post period for control units.

In Figure 5, in the panel labeled "Brand Clicks" we report estimated treatment effects for this outcome under various alternative hypotheses about the transfer of imputed clicks to Firm Fixed's ads targeting Firm Vary's brand name. In the bottom of the panel, we report the actual effect (indicated with a red dashed vertical line), which was close to 0 (indicated with a blue dotted vertical line). The rest of the estimates are ordered from largest effect to smallest and reflect different assumptions about how many clicks would be transferred to Firm Fixed. Note that the scale on the x-axis is the square root of the absolute value, times the sign i.e., $\text{sign}(x)\sqrt{|x|}$ in order to better show the data—a log scale would "zoom' in on estimates closer to zero and a linear scale would make it difficult to show the full range of estimates.

In the "+Imputed lost Firm Vary clicks" regression, the outcome is the Firm Fixed clicks plus all the implied clicks not obtained by Firm Vary. Note that because of how the experiment worked, the number of brand clicks is mechanically zero in treated DMAs in the experimental period. We can see from the point estimate and CI that had a full transfer of clicks occurred, we would have easily detected it.

As we showed when comparing brand ad differences in impressions counts (in Figure 1), full transfer of clicks is overly optimistic. For one, Firm Fixed had a smaller fraction of impressions on the brand keyword—26% overall and 69% in the last week before the experiment. However, if we use the impression ratio that prevailed right before the experiment to transfer clicks, we would have still easily detected it (labeled "+Imputed lost Firm Vary clicks (best case—69%)"). We also would have detected the effects had the average impression overlap of the pre-experiment period prevailed, as we can see in "+Imputed lost Firm Vary clicks (worst case—26%)."

The next point estimate down is "Inherits Firm Vary CTR" where we take Firm Fixed's impressions but calculate clicks using Firm Vary's CTR for the first position. We impute this CTR using the same method as we did for lost clicks using a generalized linear model with a logit link function (CTRs have to be in $[0, 1]$). Again, we can see with this inherited CTR approach, the point estimates are far larger than what we observed. In "Simonov, Rao & Blake scaling" we use the 4% estimate of siphoned clicks from Simonov et al. (2018), which was the high estimate. In "Goldman & Rao CTR scaling" we use the results from Goldman and Rao (2016) on effects of movement from position 1 to position 2 for an "off-brand" ad (i.e., in our context, Firm Fixed's ad appearing on a query using the "Firm Vary" brand name) reduces CTR to 60% of the 1 position CTR.

Figure 5: Effect of Firm Vary's search ad suspension on Firm Fixed's customer signups under various counterfactuals

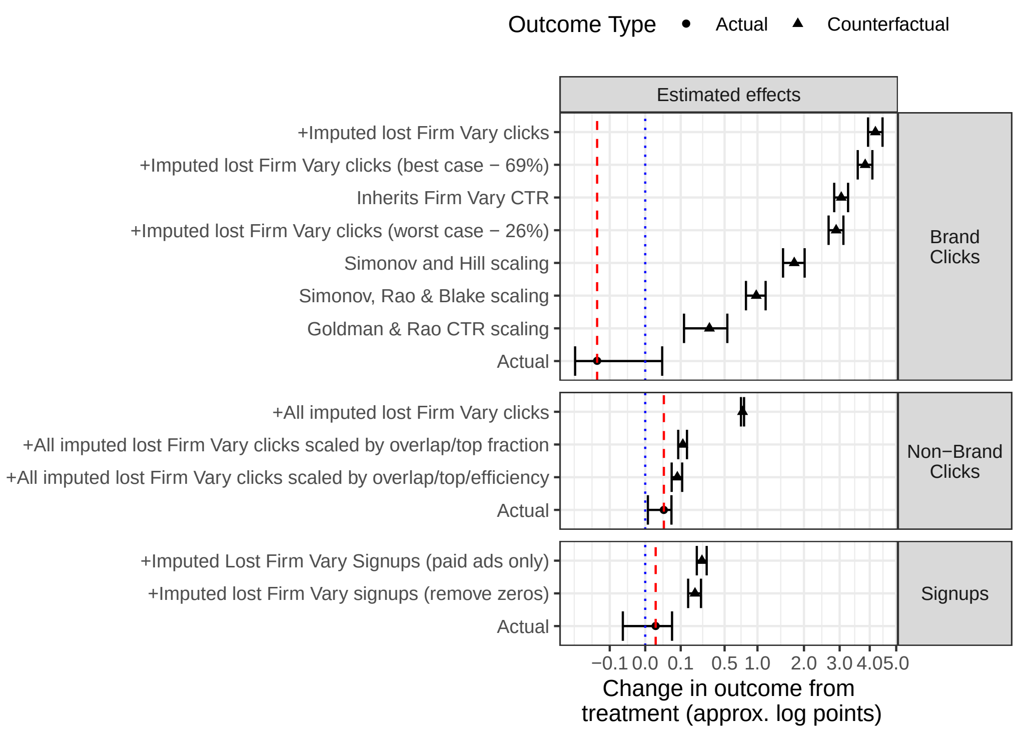

Notes: This compares our actual estimates of the effects of Firm Vary's ad suspension on Firm Fixed's signups to what we would estimate under various counterfactuals. The scale on the x-axis is the square root of the absolute value, times the sign i.e., $\text{sign}(x)\sqrt{|x|}$ in order to better show the data.

We also report a counter-factual based on Simonov and Hill (2018), which used experimental variation in the presence of a competitor bidding on the focal brand's brand keywords. They find that when in the 2nd position, the competitor steals about 1 to 2% of the traffic, but when moving to the first position, steals 6% - 15%. We use the average for both percentages, and so assume that Firm Fixed show have gotten about $0.105/0.015 \approx 7$ times as many brand clicks when in a treated DMA during the experiment. This counter-factual, labeled "Simonov and Hill scaling" is clearly far from our point estimate.

Returning to our table of results, in Column (5) of Table 2, the outcome is brand CPC, which we predicted would be unaffected by the treatment. The point estimate is actually positive and conventionally significant. However, as we show in Appendix A.3, this significant effect could easily be due to sampling variation, as this effect is sensitive to the regression specification.

The outcome in Column (6) is the total cost, which we predicted would rise because of the predicted increase in clicks (with a constant CPC). However, as we observed an insignificant change in both clicks and CPC, with the click effect being the “wrong” sign, any change in total cost would be difficult to interpret. As it is, we find no significant effect on total costs.

4.2 Experimental Effects on Firm Fixeds Non-Brand campaigns, with counterfactuals

The bottom panel (labeled *Non-brand*) of Table 2 reports the effects on Firm Fixed’s non-brand ad campaign. Position did decrease as expected, but the estimated, statistically significant effect is small, which is why we could not see it in Figure 4. We can, *ex post*, use the SEMRush keyword data to try to predict the change in average position.

Among the overlapping keywords, if we adjust Firm Fixed position ranks by eliminating Firm Vary i.e., moving the rank up if Firm Vary was above and keeping them unchanged if Firm Vary was below, the predicted position improvement is -0.33. However, if we scale this by the fraction overlapping, the predicted improvement is only -0.063, which is quite close to the experimental estimate of -0.052 (and the predicted improvement is well within the confidence interval).

This predicted mean is only based on total keywords, with no weighting based on the number of impressions of that ad, which is the case for our experimental estimate. The SEMRush data includes an estimate of search volume, and if we instead weight our prediction this way, the predicted position improvement is -0.37, which is still in our experimental CI. Again, if we scale this by the fraction overlapping, the predicted improvement in position for Firm Fixed is -0.071, which is close to the experimental estimate.

The is some limited evidence of an improved CTR, in Column (4) from having a better position. . In Column (5), we can see that CPC decreased (as expected) by a statistically significant -7%. However, this CPC effect could easily be due to sampling variation (see Appendix A.3).

As to whether our results on clicks can rule out meaningful alternative hypotheses, we can do the same kind of imputation as we did for brand clicks. In Figure 5, in the panel labeled, "Non-Brand Clicks" we report treatment effects under various counterfactuals about the transfer of clicks. Compared to brand clicks, we can see that the collection of non-brand click counterfactuals are much closer to what we actually estimated, though all are still outside the CI for the actual estimate.

In the top of "Non-Brand Clicks" panel, in the regression labeled "+All imputed lost Firm Vary clicks" the outcome variable is the clicks Firm Fixed got, plus all imputed lost non-brand Firm Vary clicks. However, the two firms did not compete on all search terms, and in only a smaller fraction was Firm Vary above Firm Fixed. To account for the lack of overlap, we use these two factors to scale down the transfer in "+All imputed lost Firm Vary clicks scaled by overlap/top fraction." This shrinks the estimate considerably, but it is still outside the CI for the actual estimate.

Another factor that would limit how many clicks Firm Fixed could obtain is the extent to which paid Firm Vary clicks were turned into organic clicks rather than "lost." We do not know this fraction directly for clicks, but we can use as a proxy the fraction estimated for signups, which we report in Section 5. This counterfactual is reported in the line labeled "+All imputed lost Firm Vary clicks scaled by overlap/top/efficiency fraction." As expected, this more realistic counterfactual is closer to our actual point estimate, but still outside the CI for the actual estimate.

5 The Experiment's Effects on Firm Vary's and Firm Fixed's Businesses

The effects of the shutdown of Firm Vary's ads on Firm Vary's and Firm Fixed's businesses, as measured by signups, are reported in Table 3, estimated by using Equation 2. The outcome in Column (1) is Firm Vary signups, while in Column (2) it is Firm Fixed signups. These signups include both organic and paid signups. We do not observe the source of signups i.e., whether they came from a branded or non-branded keyword. The regressions are estimated with Poisson QMLE.

The point estimate in Column (1) implies that Firm Vary lost approximately 23% of its new buyers by turning its search ads off. The experiment did not introduce separate exogenous variation in advertising on brand and non-brand terms, so we are unable to identify separate impacts of advertising for these two groups of terms. From Column (2), where the outcome is Firm Fixed signups, we see there is no evidence that Firm Fixed benefited from having Firm Vary out of the way. Given the Column (2) point estimate and standard error, we can rule out a positive increase in Firm Fixeds signups as large as 6% with 95% confidence.[24]

Table 3: Effects of Firm Vary search ad suspensions on new user registrations for Firm Vary and Firm Fixed

	Dependent variable:	
	Firm Vary Registrations	Firm Fixed Registrations
	(1)	(2)
$\textsc{AdsOff}_{it}$	-0.229^{***}	0.009
	(0.028)	(0.032)
Implied Efficiency	0.63	
Efficiency CI	$[0.42, 0.72]$	
N	420	420

Notes: Each column shows the estimated impact of Firm Vary shutting down its Google ads on new registrations. In Column (1), the outcome is new customer registrations for Firm Vary, while in Column (2) it is new customers registering with Firm Fixed. All estimates include DMA and time-period fixed effects and are estimated using Equation 2. The estimates are quasi-poisson maximum likelihood estimates. Standard errors are clustered at the DMA level and are in parentheses. Significance indicators: $p \leq 0.10 : \dagger$, $.p \leq 0.05 : *$, and $p \leq .01 : **$.

The apparent effects of the experiment on Firm Fixed's business is minimal. A natural question is whether this is surprising. We explore counterfactuals in Figure 5, in the panel labeled "Firm Fixed Signups." For our first counterfactual, we transfer all lost *paid* signups for Firm Vary in treated DMAs to Firm Fixed. Firm Vary mechanically had zero signups in treated DMAs in the experimental period, but we can impute the lost signups using the same method used for clicks, as all signups are labeled with their source. We report the regression results using this approach in "+Imputed Lost Firm Vary Signups (paid ads only)" line of Figure 5. We can see that had Firm Fixed picked up all of Firm Vary's lost customers, we would have easily detected it.

Transferring all paid signups to Firm Fixed is likely to overestimate what is possible, as some paid Firm Vary signups turned into organic signups in treated DMAs. As such, we impute the total number of customers Firm Vary lost of all

[24] $\Phi(0.062; \mu = 0.009, \sigma = 0.032) \approx 0.095$

kind i.e., not restricting to paid ads, and then transfer this number. However, for this imputed number, we have to subtract the number of actual signups received by Firm Vary and only transfer the difference (which could be negative if Firm Vary received more signups than predicted even if ads are off). We do this transfer, but if the difference brings Firm Fixed signups below 0, we truncate at 0. This is in the line labeled "+Imputed lost Firm Vary signups (remove zeros)." As with our other method, had Firm Fixed picked up all of Firm Vary's lost customers, we would have easily detected it.

5.1 Advertising efficiency measure

We can use the experiment to assess the efficiency of Firm Vary's ads—a key consideration for any would-be advertiser. We define efficiency as the fraction of all signups tracked to clicks on paid ads that *would not* have otherwise occurred without the ad. Let this fraction be e. If $e = 0$, it means that every paid signup would, in the absence of ads, simply come through the organic channel i.e., the new user signing up after clicking on an ad would have instead clicked on an organic search result. In contrast, $e = 1$ would imply every signup attributable to a click on a paid ad would not otherwise have occurred if that ad was not available.

We can calculate the efficiency of Firm Vary's ads from the experiment, as we know the source of customer signups. First, note that the total number of signups for a DMA, in the control group where ads are running, is simply the sum of organic and paid signups, or

$$Y^C_{ALL} = Y^C_{ORG} + Y^C_{PAID}.$$

If that same DMA had been in the treatment, the fraction $1 - e$ of its paid signups would come through the organic channel, and so the total number of signups observed in the treatment would be

$$Y^T_{ALL} = Y^C_{ORG} + (1 - e)Y^C_{PAID}.$$

As such, we can obtain an estimate of the efficiency as

$$
\begin{aligned}
\hat{e} &= \mathbf{E}\left[\frac{Y_{ALL}^T - Y_{ALL}^C}{Y_{PAID}^T - Y_{PAID}^C}\right] \\
&= \mathbf{E}\left[\frac{Y_{ALL}^C - Y_{ALL}^T}{Y_{PAID}^C}\right] \\
&\approx 0.76. \qquad (3)
\end{aligned}
$$

While intuitive, this simple method is not reasonable in practice, as different DMAs have large differences in the number of signups, leading to high variance in this estimate. As in our regressions, it is much better to transform the outcomes and perform the estimate in a regression framework where we can include DMA-specific effects.

For a meaningful interpretation of the regression results, it is useful to assume that in the control, the number of paid signups is proportional to the number of organic signups, i.e., $Y_{PAID} = zY_{ORG}$. The coefficient β_1 from Equation 2 is interpretable as the efficiency times the fraction of all signups attributable to a click on a paid ad in the treatment, or

$$
\begin{aligned}
\beta_1 &= \Delta \log Y_{ALL} \qquad (4) \\
&= \log\big(Y_{ORG}(1+z)\big) - \log\big(Y_{ORG}(1+zx)\big) \\
&= \log(1+z) - \log(1+z(1-e)) \\
&\approx \log(1+z) - (1-e)\log(1+z) \\
&\approx -e\log(1+z) \\
&\approx -ez.
\end{aligned}
$$

To identify e, we need to scale the estimated β_1 coefficient by the inverse of z, which we can estimate at the DMA level using data from the pre-experiment period and from control DMAs in the post period. Computing this fraction z with the experimental data, the point estimate for efficiency, $\hat{e}$, is 0.63. As there would also be sampling variation in z and well as β_1, to compute the standard error of $\hat{e}$ we conduct a block bootstrap of the panel, sampling DMAs with replacement and then re-labeling the index, giving a 95% bootstrap confidence interval of $[0.42, 0.72]$, with 500 replications. This point estimate of the advertising efficiency and the associated standard error is reported in Table 3 for Firm Vary.

6 Conclusion

Our results show that Firm Fixed did not gain a significant amount of the search ad traffic Firm Vary lost when it stopped bidding on its own search term. This is strongly contrary to the claim that companies must bid on their own terms to prevent competitors from reaching their customers or would-be customers. This discrepancy raises the question of whether our results generalize to other advertisers bidding on their own brand keywords. If it does, then many advertisers are needlessly spending money defending their brand terms by bidding on them.

We suspect this result does in fact generalize in many cases, because Firm Vary and Firm Fixed both behaved like typical advertisers on Google and there was nothing particularly unusual about their competition over search ads for each other's trademarked terms. Both companies were large search advertisers during the experiment, but were far from being the largest. Neither had conducted a randomized controlled trial with their search ads prior to this experiment. However, it is important to note that our two firms were of similar size and were peer competitors—in Simonov et al. (2018) there was much more poaching on brand ads, though most of it was done by much smaller competitors.

For non-brand advertising, our key finding is that search ads were effective for Firm Vary, but not as effective as a naive estimate would have implied. The lack of effects on Firm Fixed—despite conceiving of themselves as the closest of close competitors—suggests that firms can likely think of ads in a relatively simple, non-strategic way. When deciding whether to buys ads, the firm can consider whether they are worth it in terms of the customers they bring in; they do not have to consider the effects on competitors.

The results in BNT suggest that eBay's search ads were not effective. The authors attribute this result to eBay being a well-known brand—customers who clicked on a search ad and subsequently made a purchase on eBay would likely have made their purchase if not shown the ad because they already knew about eBay. In explaining our different findings, it is critical to note that Firm Vary was substantially less well known. Our results complement BNT and together provide evidence consistent with the view that the gap between the naive and causal estimates of a company's search ad campaign effectiveness increases in the company's size.

In our experimental setting, Firm Fixed did not respond to Firm Vary stopping some of their ads. In a more general setting, a competitor might respond to a business stopping some or all of their search ad purchases by bidding more or less aggressively, either overall, or only on some keywords, especially over a longer time horizon than our short experiment period.[25] Suppose Firm Vary and Firm Fixed remained competitors, and that Firm Vary stopped bidding on search ads. For the term "Firm Vary," Firm Fixed should not have responded, but rather would have passively moved from position 2 to position 1, and unhappily found out that their traffic did not increase substantially from this change.

For its non-brand campaign terms, Firm Fixed's best move to re-optimize depends on the other remaining bidders, so we are unable to predict the best response to Firm Vary dropping out of the auction. However, because of the relatively small impact Firm Vary's participation had on Firm Fixed's campaign, we suspect Firm Fixed's optimal bid changes in this scenario would be small. If Firm Fixed bid to spend a fixed marketing budget, which is a common—though not universal—practice among search advertisers, then Firm Fixed's lower CPC would allow it to increase its bids and acquire more clicks. However, even CPC effects were small, likely because there was less overlap in terms they were competing over. Of course, if Firm Vary were to drop out of the auction, it would lose the substantial amount of new business these ads generated for it, regardless of exactly how Firm Fixed would change its bids in response.

References

Ackerberg, Daniel A, "Empirically distinguishing informative and prestige effects of advertising," *The RAND Journal of Economics*, 2001, pp. 316–333.

Agarwal, Ashish, Kartik Hosanagar, and Michael D. Smith, "Do organic results help or hurt sponsored search performance?," *Information Systems Research*, 2015, *26* (4), 695–713.

Animesh, Animesh, Siva Viswanathan, and Ritu Agarwal, "Competing "creatively" in sponsored search markets: The effect of rank, differentiation strat-

[25]Note that the GSP auction is not incentive compatible (Edelman et al., 2007), making these kinds of strategic considerations relevant.

egy, and competition on performance," *Information Systems Research*, 2011, *22* (1), 153–169.

Armstrong, Mark, John Vickers, and Jidong Zhou, "Prominence and consumer search," *The RAND Journal of Economics*, 2009, *40* (2), 209–233.

Athey, Susan and Glenn Ellison, "Position auctions with consumer search," *The Quarterly Journal of Economics*, 2011, *126* (3), 1213–1270.

Baadsgaard, Jacob, "Bidding on the competition: Is it really worth it?," `https://searchengineland.com/bidding-competition-really-worth-258398` 2016.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Consumer heterogeneity and paid search effectiveness: A large-scale field experiment," *Econometrica*, 2015, *83* (1), 155–174.

Börgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek, "Equilibrium bids in sponsored search auctions: Theory and evidence," *American Economic Journal: Microeconomics*, 2013, *5* (4), 163–187.

Cummins, Evan, "Bidding on Competitors' Brands: Pros, Cons & Common Mistakes," `https://www.wordstream.com/blog/ws/2015/06/12/bidding-on-competitor-brands` 2018.

Dai, Daisy and Michael Luca, "Effectiveness of paid search advertising: Experimental evidence," *Working paper*, 2016.

Decarolis, Francesco and Gabriele Rovigatti, "Concentration and Internet advertising: The rise of buyer power," *Working paper*, 2018.

Desai, Preyas S, Woochoel Shin, and Richard Staelin, "The company that you keep: When to buy a competitor's keyword," *Marketing Science*, 2014, *33* (4), 485–508.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American Economic Review*, March 2007, *97* (1), 242–259.

Federal Trade Commission, "FTC Sues 1-800 Contacts, Charging that It Harms Competition in Online Search Advertising Auctions and Restricts Truthful Advertising to Consumers," https://www.ftc.gov/news-events/press-releases/2016/08/ftc-sues-1-800-contacts-charging-it-harms-competition-online 2016.

Gervais, Daniel, Martin Holmes, Paul Kruse, Glenn Perdue, and Caprice Roberts, "Is profiting from the online use of another's property unjust? The use of brand names as paid search keywords," *IDEA: The IP Law Review*, 2013, *53*, 131.

Ghose, Anindya and Sha Yang, "An empirical analysis of search engine advertising: Sponsored search in electronic markets," *Management Science*, 2009, *55* (10), 1605–1622.

Goldfarb, Avi, "What is different about online advertising?," *Review of Industrial Organization*, 2014, *44* (2), 115–129.

_ **and Catherine Tucker**, "Search engine advertising: Channel substitution when pricing ads to context," *Management Science*, 2011, *57* (3), 458–470.

Goldman, Matthew and Justin Rao, "Experiments as Instruments: Heterogeneous Position Effects in Sponsored Search Auctions.," *EAI Endorsed Trans. Serious Games*, 2016, *3* (11), e2.

Gomes, Renato, Nicole Immorlica, and Evangelos Markakis, "Externalities in keyword auctions: An empirical and theoretical assessment," in "International Workshop on Internet and Network Economics" Springer 2009, pp. 172–183.

Jansen, Bernard J., Zhe Liu, and Zach Simon, "The effect of ad rank on the performance of keyword advertising campaigns: The Effect of ad rank on performance of keyword advertising campaigns," *Journal of the American Society for Information Science and Technology*, October 2013, *64* (10), 2115–2132.

Jeziorski, Przemyslaw and Ilya Segal, "What makes them click: Empirical analysis of consumer demand for search advertising," *American Economic Journal: Microeconomics*, 2015, *7* (3), 24–53.

Johnson, Lauren, "U.S. Digital Advertising Will Make $83 Billion This Year, Says EMarketer," https://www.adweek.com/digital/u-s-digital-advertising-will-make-83-billion-this-year-says-emarketer/ 2017.

Kihlstrom, Richard E. and Michael H. Riordan, "Advertising as a Signal," *Journal of Political Economy*, 1984, *92* (3), 427–450.

Lewis, Randall A. and Justin M. Rao, "The unfavorable economics of measuring the returns to advertising," *The Quarterly Journal of Economics*, 2015, *130* (4), 1941–1973.

_ , _ , and David H. Reiley, "Here, there, and everywhere: correlated online behaviors can lead to overestimates of the effects of advertising," in "Proceedings of the 20th international conference on World wide web" ACM 2011, pp. 157–166.

Lewis, Randall and Dan Nguyen, "A Samsung ad for the iPad? Display advertising's competitive spillovers to search," *Working paper*, 2014.

Milgrom, Paul and John Roberts, "Price and advertising signals of product quality," *Journal of Political Economy*, 1986, *94* (4), 796–821.

Nelson, Phillip, "Advertising as information," *Journal of Political economy*, 1974, *82* (4), 729–754.

Sayedi, Amin, Kinshuk Jerath, and Kannan Srinivasan, "Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising," *Marketing Science*, July 2014, *33* (4), 586–608.

Schmalensee, Richard, "A model of advertising and product quality," *Journal of Political Economy*, 1978, *86* (3), 485–503.

Silva, Joao Santos and Silvana Tenreyro, "The log of gravity," *The Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Simonov, Andrey and Shawndra Hill, "Competitive Advertising on Brand Search: Traffic Stealing, Adverse Selection and Customer Confusion," 2018.

_ , Chris Nosko, and Justin M Rao, "Competition and crowd-out for brand keywords in sponsored search," *Marketing Science*, 2018, *37* (2), 200–215.

Varian, Hal R., "Position auctions," *International Journal of Industrial Organization*, December 2007, *25* (6), 1163–1178.

Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data*, MIT Press, 2002.

Table 4: Top 15 Online retailers and whether they bid on own brand name

Company	e-Sales (mil.)	e-Share	Brand ad?	Position	CPC
Amazon.com	$79,268	74.10%	Yes	1	0.02
Wal-Mart Stores Inc.	$13,484	2.80%	Yes	1	0.04
Apple	$12,000	5.10%	Yes	1	0.03
Staples	$10,700	55.50%	Yes	1	0.12
Macy's	$4,829	17.50%	Yes	1	0.06
The Home Depot	$4,267	5.00%	Yes	2	0.16
Best Buy	$3,780	9.40%	Yes	1	0.06
QVC	$3,722	42.70%	Yes	1	0.03
Costco Wholesale	$3,618	3.10%	No	N/A	0.58
Nordstrom	$2,699	18.90%	Yes	1	0.03
Target	$2,524	3.40%	Yes	1	0.23
Gap Inc.	$2,519	15.60%	Yes	1	0.06
Williams-Sonoma	$2,501	50.70%	Yes	1	0.09
Kohl's	$2,367	12.40%	Yes	1	0.08
Sears Holdings	$2,057	7.90%	Yes	1	0.24

Notes: This table shows the top 15 online retailers and whether they bid on their own brand name. Source: `https://wwd.com/business-news/financial/amazon-walmart-top-ecommerce-retailers-10383750/` combined with July 28th, 2018 data from SEMRush.

Yang, Sha and Anindya Ghose, "Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence?," *Marketing Science*, 2010, *29* (4), 602–623.

Yao, Song and Carl F Mela, "A dynamic model of sponsored search advertising," *Marketing Science*, 2011, *30* (3), 447–468.

A Appendix

A.1 How common is bidding on own brand name?

Table 4 shows that is it exceedingly common for firms to bid on their own brand name keywords alone. All of the top 15 online retailers except Costco bid on their brand name and occupy the first position. Perhaps not surprisingly, the CPC for "costco" is not very low, as it is for the other brand names.

A.2 Randomization

In Table 5, we report the results of regressions where the outcome is pre-experiment DMA-level Firm Fixed ad campaign attributes. The key independent variable is the *future* experimental assignment of that DMA. Consistent with effective randomization, there is no evidence that the treatment assignment predicts these pre-experiment attributes.

Table 5: Placebo check for whether Firm Vary's ad campaign suspension affected Firm Fixed's sponsored search advertising campaigns *before* start of the experiment

Brand ads:

	Position	Impressions	Clicks	CPC	Cost
	(1)	(2)	(3)	(4)	(5)
$\textsc{AdsOff}_{it}$	−0.003	−0.033	0.150	−0.074	0.058
	(0.014)	(0.339)	(0.353)	(0.067)	(0.349)
N	204	210	210	163	210

Non-brand ads:

	Position	Impressions	Clicks	CPC	Cost
	(1)	(2)	(3)	(4)	(5)
$\textsc{AdsOff}_{it}$	−0.007	−0.007	−0.003	0.005	−0.002
	(0.010)	(0.331)	(0.342)	(0.014)	(0.362)
N	204	210	210	204	210

Notes: This table reports regressions where the outcomes are pre-experiment DMA-level ad campaign attributes for Firm Fixed. All estimates include DMA fixed effects. Standard errors are clustered at the DMA level and are in parentheses. Columns (1) and (5) contain fewer observations because position and CPC are only defined if there are any impressions and clicks, respectively, in a period-DMA observation. Significance indicators: $p \leq 0.10 : \dagger$, $.p \leq 0.05 : *$, and $p \leq .01 : **$.

A.3 Alternative specifications of the treatment effects

There are several ways to analyze the experimental outcomes. Although we think the QMLE approach used in Table 2 is preferable, there are reasonable alternatives. For one, instead of collapsing data into pre and post periods, we could also use each day as the unit of analysis, since campaign metrics are reported at that frequency. In Figure 6, we plot the treatment effect of Firm Vary turning off ads using a number of different specifications. The effects for brand ads are shown in the left column, and for non-brand ads in the right column. Note that each panel has an outcome-specific scale on the y-axis.

Figure 6: Effects of Firm Vary's search ad suspension on Firm Fixed's campaign

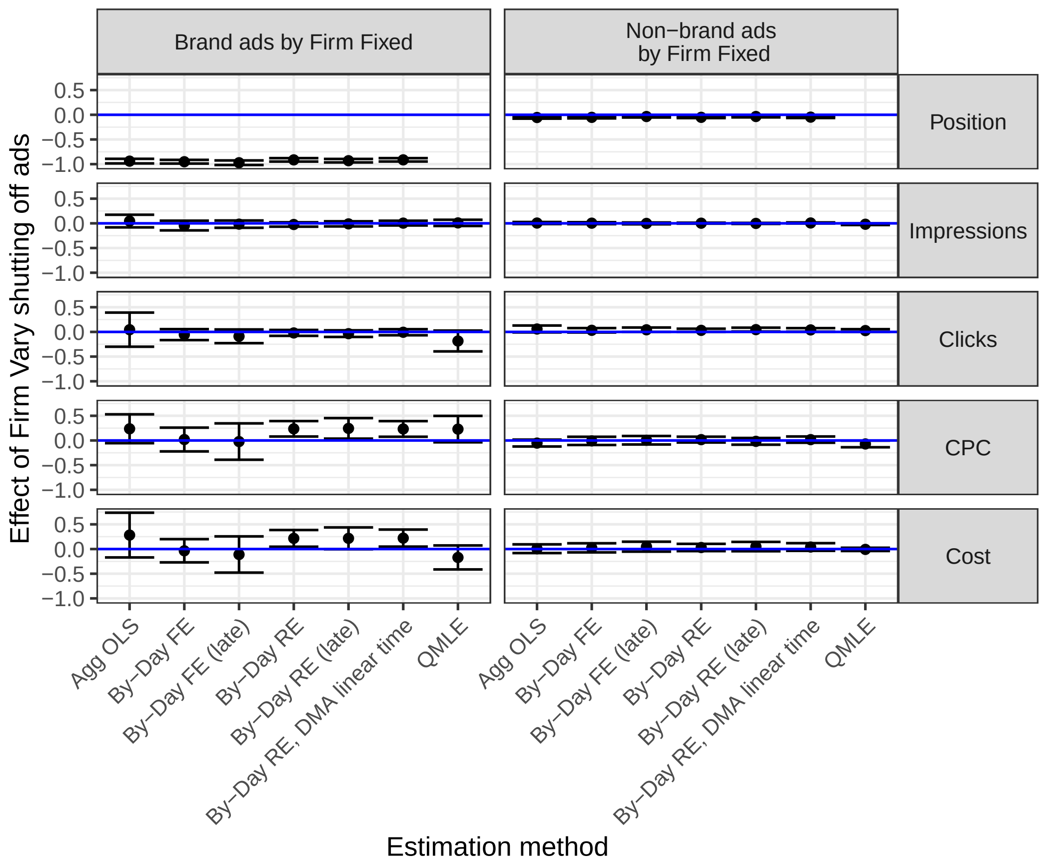

Notes: This figure shows a collection of estimates for the effects of Firm Vary's experiments on Firm Fixed's outcomes. The aggregate OLS sample uses the outcome, or its log transform in stead of the QMLE. The other estimates use a DMA-day level of analysis. The specifications, from left to right, are (1) collapsed (same pre/post set-up as the QMLE but with the log outcome); (2) by-day, with day and DMA-specific fixed effects and DMA clustered SEs; (3) same as 1, but only using a symmetric window around the experiment; (4) by-day, with day-and-DMA specific random effects; (5) same as 4, but with the addition of DMA-specific linear time trends; and (6) collapsed Poisson QMLE (the same estimate as reported in Table 2)

For each non-position outcome (except Poisson QLME regressions), we use the log of the outcome, dropping observations with a value of zero from the data. For the position outcome, we do not include the Poisson QLME estimate. The specifications, from left to right, are (1) collapsed (same pre/post set-up as the QMLE but with the log outcome); (2) by-day, with day and DMA-specific fixed effects and DMA clustered SEs; (3) same as 1, but only using a symmetric window around experiment; (4) by-day, with day-and-DMA specific random effects; (5) same as 4, but with the addition of DMA-specific linear time trends; and (6) collapsed Poisson QMLE (the same estimate as reported in Table 2).

From Figure 6, we can see that the specification generally does not seem to matter very much; all the point estimates and associated standard errors are reasonably similar to each other. However, the borderline significant results from Table 2 seem likely to be attributable to sampling variation. In particular, the finding of significant effects on Firm Fixed CPC in Table 2 does not hold up for either kind of ad—for both, the point estimates are of different sign depending on the specification. This is reassuring in the case of brand ads, where we *ex ante* expected no effect, and even for non-brand ads, given how small the observed effects are on position. Although small in the case of non-brand ads, the effects on position seem quite robust.

A.4 Brand position actual position

Both Firm Fixed and Firm Vary were bidding on Firm Vary's brand keyword. Figure 7 plots the frequency of average positions by day, for treated and control DMAs, by the pre- and experiment period. We can see that prior to the experiment, in both treatment and control DMAs, Firm Fixed's ad generally occupied the second position, though there was some variation.

Following the experiment, with Firm Vary out of the way in treatment DMAs, Firm Fixed's ad generally took the first position. The treatment effect on position we found in Table 2 for brand ads was not just one position—it was generally a move from the second to the first position. As such, we see no evidence that the competetive fringe—unlike Firm Fixed—tried to take advantage of Firm Vary's exit and "move up," as in the model of Sayedi et al. (2014), where smaller firms can more effectively poach. This could just be a short-run result, but given that Firm Fixed seemed to get almost no incremental clicks, it seems unlikely that any firm would do much better, suggesting there probably would not be a long-run effect.

Figure 7: Average daily Firm Fixed search ad position for Firm Vary brand keyword, by DMA treatment assignment

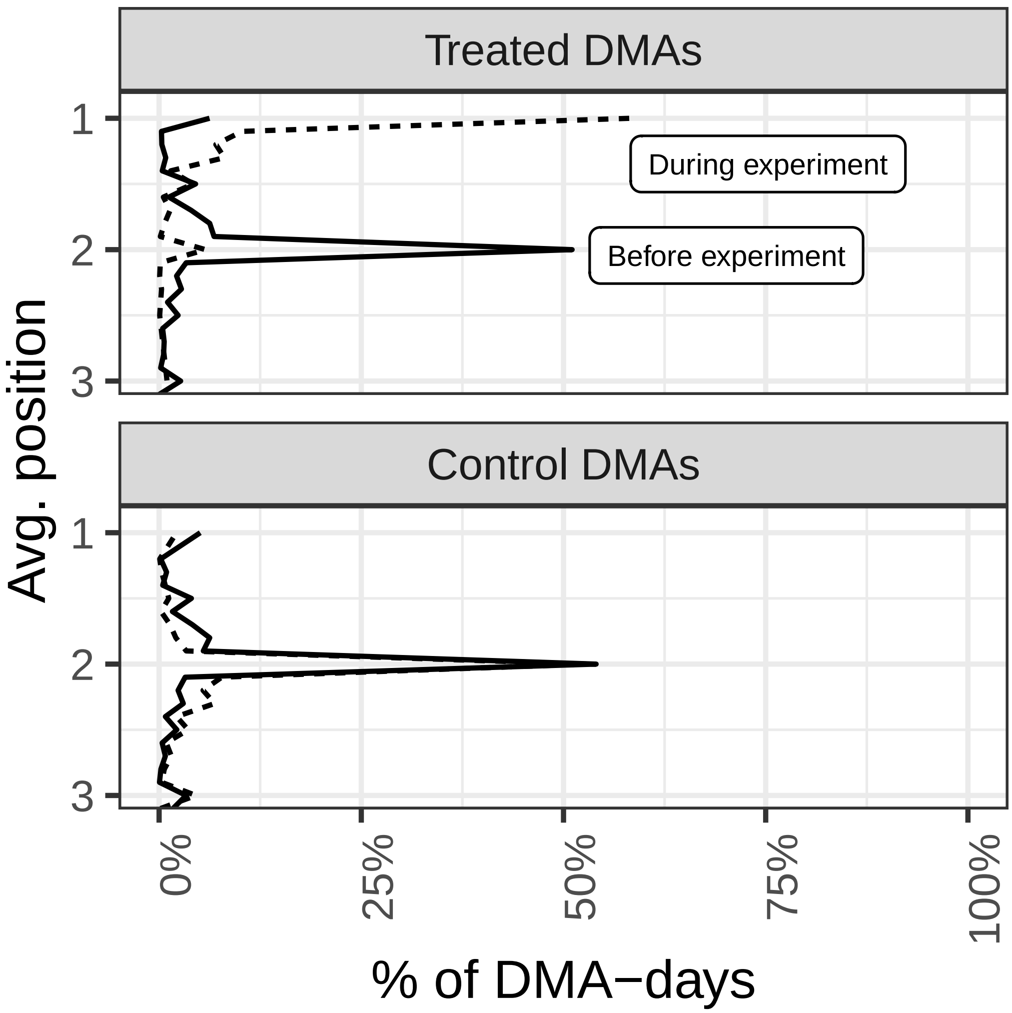

Notes: This figure plots average daily position for Firm Fixed's ad on Firm Vary's brand keyword. The kernel density estimate, by treatment assignment, is plotted in a heavy line.