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

Joseph M. Golden John J. Horton*

October 29, 2017†

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 assess its ads' effectiveness by running an experiment, suspending campaigns in some markets while keeping the status quo in others. However, a firm by itself typically cannot know the effects of its advertising on competitors, as competitors are, of course, not likely to share information with each other. And yet for firms that care about market share—such as firms in winner-take-all/most industries—the effects of their advertising on competitors might be a key consideration.

*Collage.com and NYU Stern, respectively.

†We thank Tom Blake, John Bound, Charles Brown, Dan Hirschman, Jason Kerwin, Eric Schwartz, Isaac Sorkin, Frank Stafford, Hal Varian, and Steve Tadelis for feedback on this research. 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.

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 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. To the best of our knowledge, this is the first paper we are aware of that contains an experimental estimate of the effects of advertising on a competitor.

Firm Vary and Firm Fixed were two online marketplaces for services, connecting buyers and sellers. Prior to merging, these companies were fierce business competitors. As an indication of how similar the two product offerings were, at the conclusion of the merger, buyers and sellers were simply migrated to Firm Fixed's platform. Before merging, they were battling search advertisers, often targeting their ads at the same search terms on Google and other search engines. In their 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 so-called "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 and during the experiment in the regions where Firm Vary continued to purchase Google search ads, Google searches for the term "Firm Vary" would generally show a brand ad for Firm Vary in the top position and an ad for Firm Fixed in the second position.

We examine two sets of experimental effects: (1) the effects of Firm Vary's experiment on Firm Fixed's search advertising campaign and (2) the efficacy of Firm Vary's ads as measured by the effects on its own business and Firm Fixed's business. For (1), we consider the effects of the experiment on both Firm Fixeds brand and non-brand advertising. Our experimental design is essentially identical to Blake et al. (2015) (BNT hereafter), which conducted a similar experiment with eBay's sponsored search ads. Aside from our unique ability to measure the effects on competitors, another innovation relative to BNT is that our experimenting firm was not nearly as well-known as eBay, and as BNT conjecture (and we demonstrate), how well-known the brand is could have a strong effect on the effectiveness of sponsored search advertising.[1]

For brand advertising, we find that when Firm Vary turned off its own ads, Firm Fixed's ads on the term "Firm Vary"—its brand advertising—moved into the top advertising position, as expected. However, surprisingly, Firm Fixed received nearly the same number of clicks in these 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. This result strongly suggests that even for a fairly unknown brand, queries with that 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. In addition to the number of clicks not changing, the costs per click did not fall, which we predicted, as the nature of the ad auction implies Firm Fixed's price per click was not determined by Firm Vary.

For non-brand advertising, the effects of Firm Vary's exit were minimal. The only detectable effect was that Firm Fixed's ads moved up in average position slightly. There was no detectable change in the number of clicks Firm Fixed received or their cost per click or any other metric. We predicted large changes in position because we thought that Firm Vary and Firm Fixed's search ad campaigns were aggressively competing head-to-head. The lack of effects on clicks is consistent with how small the change in position was.

One possible explanation for the lack of effects on position for non-brand advertising would be that Firm Vary's ads almost always had a worse position than Firm Fixed's, and thus the removal of Firm Vary's ads could do little. Google does not offer auction specific data, but the average position of Firm Vary's non-brand advertisements was in a *better* position (i.e., higher up) than Firm Fixed's non-brand advertisements, making this relative position explanation unlikely. The most likely explanation for the lack of effects is that both firms, despite being close competitors and having similar search ad advertising categories, were often not contesting the same keywords.

In terms of the business outcomes from Firm Vary's experiment, we focus on customer sign-ups, or "registrations" for both firms. In treated advertising markets, Firm Vary has about 23% fewer customer registrations. However, we find little

[1] Coviello et al. (2017) conduct an experiment identical to BNT, but using a retailer much less well-known compared to eBay. They find results similar to our own, in that paid ads are effective, though not fully efficient because some lost clicks are diverted to the organic channel.

evidence that Firm Fixed gained any new customers in areas where Firm Vary shut off its ads, as the point estimate 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 any appreciable number of Firm Vary's lost paid customers, but that they also did not receive more customers from users clicking on "organic" (i.e., unpaid) Firm Fixed search results.

For Firm Vary, we also examine the efficiency of its ads, which we can think of as the fraction of new customers who clicked a search ad link, then registered, who would not have registered without the presence of the ad. When a visitor arrives at a company's website from a search engine, the company knows what specific 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 is to assume that all registrations resulting from clicks on paid ads would not otherwise have occurred. However, it is likely that at least some of those users would have counter-factually just clicked on an organic result if there was no sponsored search ad shown.

We compare 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.

BNT that brand-based advertising is ineffective, which some have interpreted as a consequence of eBay's well-known brand. However, we also find brand-based advertising is 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 precisely none of this "lost" traffic went to their close competitor, even though Firm Fixed moved up in position. BNT speculate that the threat of poaching 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 unnecessary.

For non-brand search advertising, we reach very different conclusions than BNT, albeit for explicable reasons. BNT find that even non-brand search advertising is ineffective at increasing sales. It is important to note that our business measure is not sales, but new customer registrations. 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 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.[2]

Given the enormous sums spent by firms on digital advertising—estimated at $83 billion in 2017—of which search advertising makes up a large share[3], the insights offered by this paper have practical importance to firms regarding the efficient allocation of resources. The first insight is that sponsored search ads were effective at acquiring new customers (albeit not as much as a naive estimate would suggest), and 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, which 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—very well might not be the case, even for not very well-known brands. Our two firms certainly felt like they were in close competition and engaged in a prisoner's dilemma, whereas the post merger data implies they were—to push the prisoner's dilemma metaphor—not even talking to the same prosecutor. What seems probable is that customers using the brand name in a search almost always are familiar with the brand and their query is thus likely to be navigational.

The paper proceeds as follows. Section 2 discusses sponsored search advertising. Section 3 describes the empirical context and our experiment design. Section 4 shows the results of the experiment and we conclude in Section 5.

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

[2]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).

[3]http://www.adweek.com/digital/u-s-digital-advertising-will-make-83-billion-this-year-says-emarketer/

able to meet.[4] 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).[5]

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, and face similar constraints and objectives. 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.

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

[4] 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)?

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

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:

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

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. As discussed in Varian (2007) and Edelman et al. (2007), the GSP mechanism differs from the Vickrey-Clarke-Groves mechanism in that it has no dominant strategy equilibrium, and truth-telling is not an equilibrium. Rather, advertisers can have an incentive to bid below their valuation, because in some cases doing so results in a less prominent, but less expensive ad position. The GSP mechanism is thus subject to strategic manipulation, even when advertisers know their true valuations.

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.[6] However, others have argued that users entering brand names are often entering 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.

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%.[7] Revenues from trademarked/branded keywords as a share of Google's total revenue are plausibly higher today, as Google has both permitted and

[6] 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.

[7] 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.

encouraged more advertising on trademarked terms over the years.[8] 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. Both Firm Vary and Firm Fixed used sponsored search advertising to acquire new customers, focusing on "buyers" rather than "sellers." Firm Vary cared specifically about potential business being lost to Firm Fixed due to the potential winner-take-all dynamics of their competing network-based online marketplaces.[9] As such, Firm Vary historically bid on its own name—i.e., engaged in brand advertising—to keep Firm Fixed from poaching potential customers via search advertising. Firm Vary was spending about $10 million per year on sponsored search advertising, and Firm Fixed was spending a similar amount. Before the experiment, Firm Vary was spending about 10.9% of its marketing budget targeting competitors' brand keywords directly, though it was only spending about 1.1% on its own brand keywords.

Firm Vary and Firm Fixed were not very well known brands 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 related firms. Both Firm Vary and Firm Fixed had very little brand awareness, with both being less than 5%—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.

[8] 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.

[9] 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.

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[10] analyzing this industry—online marketplaces similar to Firm Vary and Firm Fixed— found that Firm Vary and Firm Fixed had nearly identical "industry" shares of the web page visits, with their cumulative fraction close to 50%. 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.

Although Firm Vary and Firm Fixed made up a large fraction of an identifiable industry, they actually competed with a much larger number of advertisers that were interested in the same keywords. The same market report discussed above reported that over its entire campaign for a month, 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. This out-of-industry 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.

3.1 Experimental design and internal validity

The design of the experiment is simple. During the experiment, Firm Vary shut off all of its Google search ads in half of the Direct Marketing Areas (DMAs) in the United States for a period of 28 days, starting on March 11, 2014. There are 210 DMAs, which subdivide the country into regions and were originally designed for television-based advertising purposes. Advertisers can now target Google ads geographically by DMA. Figure 1 shows the US DMAs in the experiment and their treatment assignment, with control DMAs in white, and treatment DMAs in gray. Treatment and control DMAs were selected at random.

After the experiment, Firm Vary resumed bidding on search ads in the entire

[10] 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.

Figure 1: Treatment and control direct marketing areas (DMA)

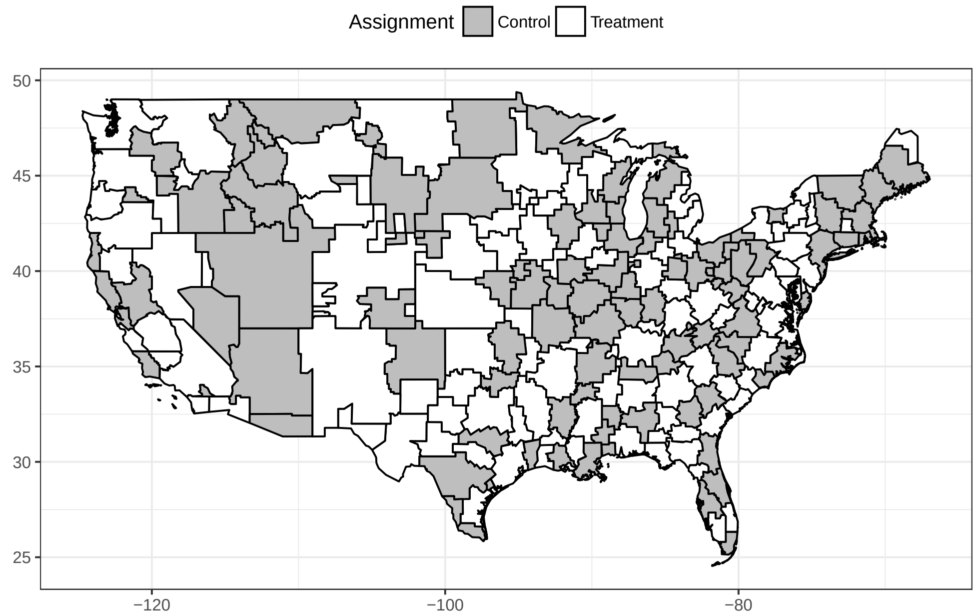

Notes: This figure shows the US directed marketing areas (DMA) and their allocation to either the treatment—in which Firm Vary turned off all search advertising—or the control, where Firm Vary kept its search advertising campaigns unchanged.

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.

3.2 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 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.[11] Although there is no evidence that they did, Firm Fixed could have learned that something had happened independently. Fortunately during the experiment, Firm Fixed neither changed its bidding behavior overall nor did

[11] 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.

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 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.

3.3 Sample size and usable data

The experiment's duration was limited by business concerns. A simulation-based a priori power analysis suggested that the experiment should run for a minimum of two weeks, a duration we expected would yield somewhat imprecise, but usable results to assess the overall performance of Firm Vary's ad campaign. 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. In the case of our experiment, ideally, we would be able to randomize at the individual level, rather than the DMA level, to gain more power; however, doing so is not possible with currently available ad targeting and tracking options.

Unfortunately, the 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.[12]

[12] 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.

3.4 Advertising predictions

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.

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.[13] For Firm Vary's brand ads, since

[13] Firm Vary's internal market research also identified Firm Fixed as its primary competitor.

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. However, as the two companies did not target the exact same set of keywords, we expected 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.

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 average position of non-brand ads for both firms), 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. We further expected Firm Fixed to get more clicks due to the absence of Firm Vary's ad, independently of the effects of clicks due to position and impressions.

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). We know the expected direction of the effects of the experiment on Firm Fixed's CPC and position based on the search ad auction's features. While the details of the auctions are important, the notion that we can think of CPC as a price subject to the forces of supply and demand seems well-supported (Goldfarb and Tucker, 2011).[14] 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 with their search ad campaigns.

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. 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.

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

[14] 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.

position of Firm Fixed's ads.

3.5 Effects on Firm Vary and Firm Fixed Business

For business outcomes—namely new customer registrations—we are interested in both the effects of the experiment on Firm Vary and Firm Fixed, whereas for the sponsored search campaign metrics, all of outcomes were Firm Fixed outcomes. Of course, the effects on Firm Fixed's business depend on how much of an effect Firm Vary's experiment had on Firm Fixed's ad campaign.

As both Firm Vary and Firm Fixed primarily used search ads to attract new buyers to their marketplaces, the number of new customer registrations is the metric we use to quantify the effects of Firm Vary's ads on both businesses. For Firm Vary registrations, we focus on estimating ad efficiency. Previous research (Lewis et al., 2011) conducted randomized controlled trials in advertising across a variety of online settings and demonstrated by comparing observational and experimental estimates that observational methods can drastically overestimate the efficacy of online advertisements.

For Firm Vary, we expected registrations 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 registrations to increase in DMAs where Firm Vary ceased advertising, but by an unknown amount.

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. 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.[15]

Overall, we thought that Firm Vary would lose some amount of new business

[15] They report a field experiment in which a retailer varied their ad creative and position rank. These factors do affect click through rates, highlighting the importance differentiation with respect to rivals.

by turning its ads off. 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 registrations, we expected an increase in this metric for Firm Fixed.

4 Results

We focus on the effects of Firm Vary's experiment on Firm Fixed's advertising campaigns. Then, we turn to the effect on both Firm Vary's and Firm Fixed's business.

4.1 Effects on Firm Fixed's advertising campaign

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 2 plots the daily difference between the treatment DMAs and control DMAs for a collection of outcomes, for both brand search ads (dashed line) and non-brand ads (solid). The data show 40 pre-period days and the full post experiment period.[16] 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 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 visually discernible effect on the position for non-brand advertisements, highlighting the usefulness of a regression-based approach.

In the remaining panels, the outcomes are the log cumulative number 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

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

Figure 2: Difference in average daily Firm Fixed search ad campaign outcomes by DMA treatment assignment

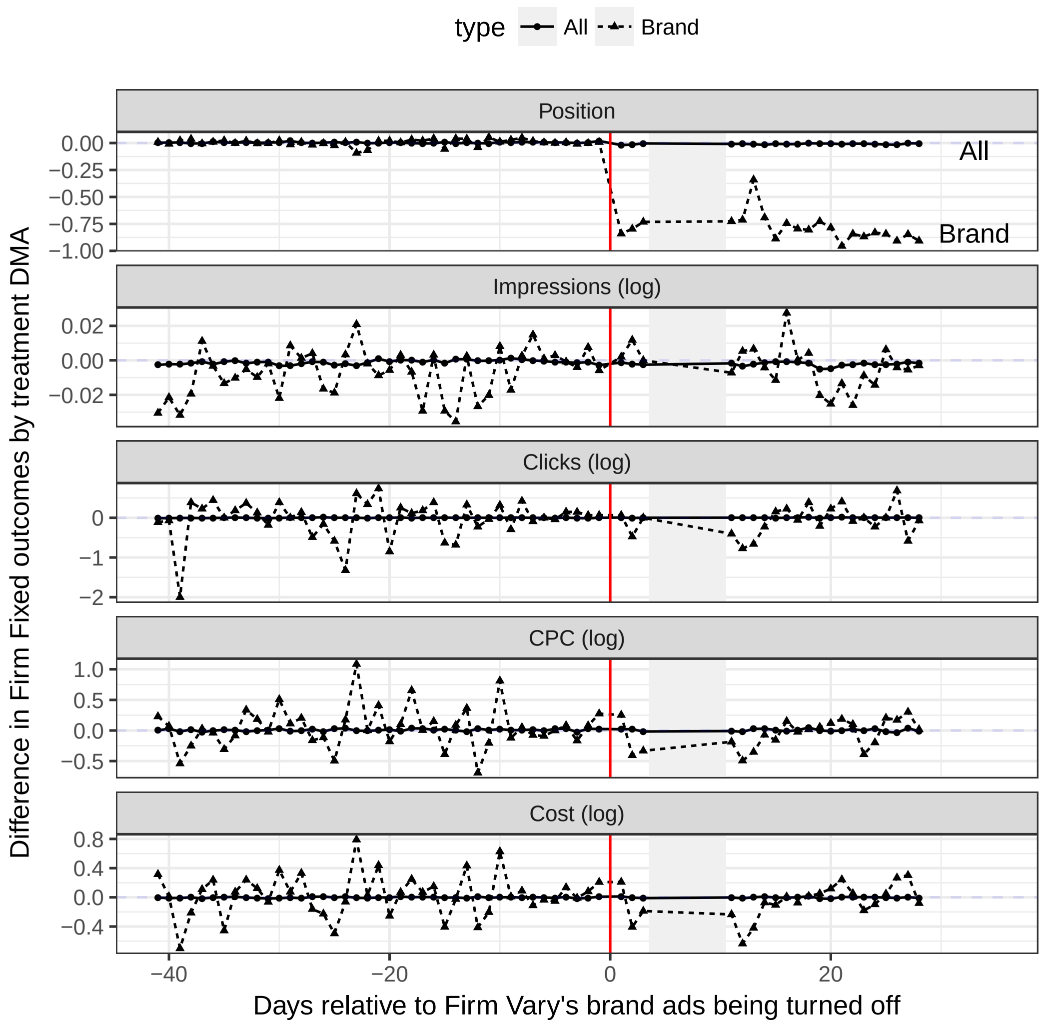

Notes: This figure shows the by-day gap in search ad outcome metrics for Firm Fixed, for both brand ads i.e., ads appearing when the search term was "Firm Vary" (indicated with a dashed line) and for all, non-brand search advertising (solid line). For position, the outcome is difference in average impression-weighted positions. 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 those where it continued. Data are omitted for the seven day "Omitted Period," as described in Section 3.1, and indicated by gray rectangles.

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—later, we will use different specifications, including by-day outcomes. In all results, we cluster standard errors at the DMA level.

For variables where we wish to estimate a percentage change, we use $f(x) = exp(x)$ and estimate the results using the Poisson quasi-maximum likelihood estimator (QMLE).[17] We prefer this estimator to taking the log of Y_{it} because some of our outcome observations are equal to 0.[18] 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 of Firm Fixed's advertising campaign are reported in Table 2. Panel A reports brand estimates, while Panel B reports non-brand estimates. Each regression is an estimate of Equation 2 using the different ad outcome metrics described in Table 1. In all cases, the treatment group is the set of DMAs where Firm Vary turned its ads off.

Starting in Panel A, 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).

[17] 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).

[18] 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	CPC	Cost
	(1)	(2)	(3)	(4)	(5)
AdsOff_{it}	-0.938^{**}	0.009	-0.184	0.230^{*}	-0.171
	(0.016)	(0.049)	(0.119)	(0.104)	(0.107)
N	408	420	420	420	241

Non-brand ads:

	Position	Impressions	Clicks	CPC	Cost
	(1)	(2)	(3)	(4)	(5)
AdsOff_{it}	-0.055^{**}	-0.018	0.028	-0.070^{**}	-0.007
	(0.008)	(0.018)	(0.020)	(0.025)	(0.024)
N	408	420	420	420	408

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 1. 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 $p \leq .01 : **$.

Furthermore, 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. Given the base click rates for these ads (which we intentionally do not report), if Firm Fixed had instead captured all of the search ad clicks Firm Vary lost by not running its ads, Firm Fixed would have received approximately 10,000% more clicks, leading to a coefficient of approximately 100 in this regression specification.

In Column (4), the outcome is CPC, which we predicted would be unaffected by the treatment. The point estimate is actually positive and conventionally significant. However, as we will show later, this significant effect is likely due to sampling variation, as this effect is sensitive to the regression specification. The outcome in Column (5) 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.

Panel B 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 2. There is no strong evidence of a change in impressions or clicks. In Column (4), we can see that CPC decreased (as expected) by a statistically significant 7.0%. However, as we will show, this effect is also likely due to sampling variation, as it is not robust to various alternative specifications.

4.2 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 3, 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 3: Effects of Firm Vary's search ad suspension on Firm Fixed's campaign

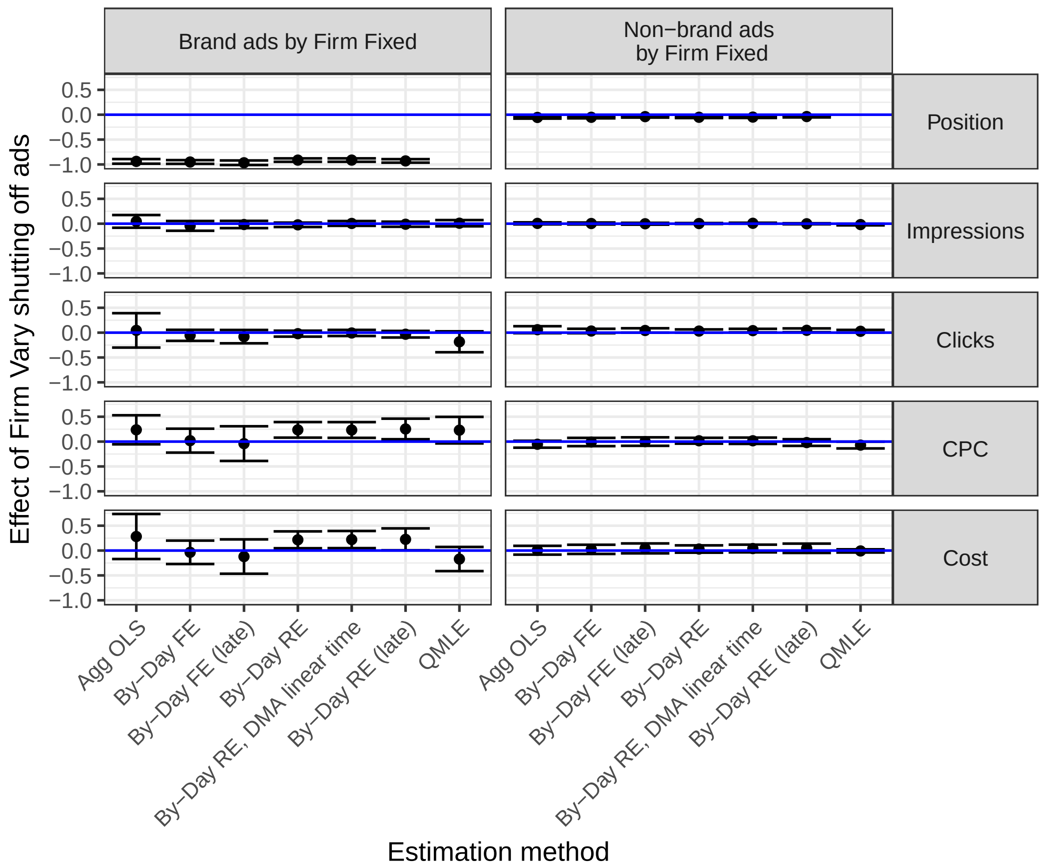

Notes: This figure shows a collection of estimates for the effects of Firm Vary's experiments on Firm Fixed's outcomes. The aggregate OLS sample uses the outcome, or its log transform in stead of the QMLE. The other estimates use a DMA-day level of analysis. The specifications, from left to right, are (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 3, 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 in either case—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 exected 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.

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

We next analyze the effect of shutting down Firm Vary's ads on Firm Vary's and Firm Fixed's businesses, as measured by sign-ups. Table 3 shows the business impact of the experiment as estimated by using Equation 2. The outcome in Column (1) is Firm Vary, while in Column (2) it is Firm Fixed registrations. These registrations include both organic and and paid sign-ups. 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. Perhaps unsurprisingly, given that Firm Vary's experiment had little discernible effect on Firm Fixed's search ad campaign, from Column (2), we see there is no evidence that Firm Fixed registered more new customers with Firm Vary out of the way.

We can use the experiment to assess the efficiency of ads—a key consideration for any would-be advertiser. We define efficiency at the fraction of all registrations 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 registration would, in the absence of ads, simply come through the organic channel i.e., the new user

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	$[0.41, 0.71]$	
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 $p \leq .01 : **$.

signing up after clicking on an ad would have instead clicked on an organic search result. In contrast, $e = 1$ would imply every registration 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 ads the from the experiment. First, note that the total number of registrations for a DMA, in the control group where ads are running, is simply the sum of organic and paid registrations, 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 registrations would come through the organic channel, and so the total number of registrations 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

$$\begin{aligned}\hat{e} &= \mathbf{E}\left[\frac{Y^T_{ALL} - Y^C_{ALL}}{Y^T_{PAID} - Y^C_{PAID}}\right] \\ &= \mathbf{E}\left[\frac{Y^C_{ALL} - Y^T_{ALL}}{Y^C_{PAID}}\right] \\ &\approx 0.76. \end{aligned} \tag{3}$$

While intuitive, this simple method is not reasonable in practice, as different DMAs have large differences in the number of registrations, 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 registrations is proportional to the number of organic registrations, i.e., $Y_{PAID} = zY_{ORG}$. The coefficient β_1 from Equation 2 is interpretable as the efficiency times the fraction of all registrations 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 $[0.41, 0.71]$, with 500 replications. This point estimate of the advertising efficiency and the associated standard error is reported in Table 3.

5 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.

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.[19] 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 term.

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 strategy, 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.

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

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

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.

Coviello, Lorenzo, Uri Gneezy, and Lorenz Goette, "A large-scale field experiment to evaluate the effectiveness of paid search adversiting," *Working Paper*, 2017.

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

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.

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.

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.

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.

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.

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.

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.