Wasn't That Ad for an iPad? Display Advertising's Impact on Advertiser- and Competitor-Branded Search

Randall Lewis* Yahoo! Research

Dan Nguyen† University of Chicago

November 13, 2012

Abstract

We measure how display ads influence customers' online searches for both the advertised brand and its competitors by randomizing the delivery of display ads to visitors to Yahoo!'s front page, www.yahoo.com, and comparing subsequent activity on Yahoo! Search. In three advertising campaigns featuring car insurance, a car, and a tablet computer, display ads increased searches for the advertised brand by 30% to 45% *and* increased searches for its competitors' brands by up to 23%. Strikingly, the total number of incremental searches for the competitors' brands was 2 to 8 times more than that for the advertised brand. These spillovers from online display advertising to online search create cost-complementarities that can influence the bidding behavior of advertisers in search advertising. Further, an extension of Grossman and Shapiro (1984) suggests that positive spillovers may reduce firms' investment in advertising relative to social-welfare-maximizing levels.

Keywords: *search advertising, display advertising, advertising synergies, cross-media advertising, multimedia advertising effectiveness, strategic complements of advertising, natural experiments, spillovers, externalities*

---

*This research was performed by the authors while at Yahoo! Research during the summer of 2011. Randall Lewis is now employed by Google, Inc. and can be contacted at ralewis@google.com, 1600 Amphitheatre Parkway, Mountain View, CA 94043.

†Dan Nguyen can be contacted at dan.nguyen@chicagobooth.edu. Additionally, we would like thank Yahoo! Inc. for providing extensive support to and granting academic independence to this research. We would also like to express our thanks to David Reiley, Justin Rao, Preston McAfee, and Catherine Tucker for their constructive feedback and many other participants at the Berkeley Marketing Seminar, 2012 NBER SI of IT & Digitization Workshop, and FTC Microeconomics Conference for their questions, comments, and critiques.

1 Introduction

1.1 Overview

Online advertising encompasses a number of new advertising channels distinct from the traditional offline channels such as print media and television. These new channels include search, display, online classified ads, lead generation, online video, video games, and mobile devices. The market for these channels has been growing throughout the past decade as indicated by online advertising's growing share of total media expenditure.[1] Because so much money is invested in these channels, it is important for both advertisers and online publishers to understand their impact. This paper increases our understanding by studying the impact of display advertising on online branded searches.

We hypothesize that display advertising increases searches for the advertised brand because it increases interest in the brand. Further, display advertising also increases interest in the whole product category, resulting in more searches for the competitors' brands. These additional searches represent positive spillovers (i.e., externalities) from brand display advertising to competitors.

To test these hypotheses, we measure the search spillovers from display advertising by exploiting exogenous variation generated by a natural experiment on Yahoo!'s front page (www.yahoo.com).[2] On certain days, Yahoo! sells the main display ad space on the front page as an "ad split" in which two display ads alternate every second throughout the day. As long as visiting the front page on an even or odd second is orthogonal to a visitor's response to the display ad, the ad split creates exogenous variation to identify the causal effects of the display ad. To take advantage of this natural experiment, we recorded front page visitors' ad exposures and subsequent search behaviors on three ad split days when one

[1]From 2009 to 2010, US online advertising expenditure grew from 22.8 to 26.1 billion USD, a 14.5% increase, while traditional media grew from 124.4 to 126.9 billion USD, a 2.0% increase (Mintel Reports Database, 2010).

[2]Lewis (2010) uses this same variation to examine the impact of frequency on clicks and new account sign-ups.

of the advertisers was from a well-defined product category (the target ad) and the other advertiser was from an unrelated category (the control ad). We compared the impact of exposure to the target and control ads by examining users' search behavior immediately following visits to the front page. Focusing on the ten minutes following exposure, we find that the display ads caused a 30 to 45% increase in searches for the advertised brand and a 1 to 6% increase in searches for any competitors' brands. When analyzing the display ad's effect on searches by brands, we found as high as a 23% increase in searches for a competitor's brand. There was significant heterogeneity in the effect across competitors' brands, perhaps reflecting differences in competitive distance, but no competitor experienced a statistically significant decrease in searches from the display ad. Therefore, display advertising increases the likelihood of consumers to search for many brands in the product category.

We conducted several robustness checks of our results. We found that our results hold when we limited our sample to only the period after the first impression and when we limited our sample to users who only saw one impression. We also looked at the most statistically impacted words, queries, and domains clicked. The results from these analyses mirror our main results and also give more insight to the type of searches conducted. Specifically, we found that display ads increased searches that navigate users to a brand's website, searches for informative sites about the whole product category, and searches relating to purchase intent.

This paper makes several contributions to the advertising literature. This paper adds to the numerous papers that measure the effects of display advertising. Most of these papers concentrate on the effects of display advertising on clicks and survey response data while our paper looks at the effects on online search. There are a growing number of papers that try to measure the effect of advertising on consumers' online searches, but our paper is the first to use an experiment with clean exogenous variation to measure these effects. This paper is also the first field experiment to measure the positive spillovers of brand advertising on the whole product category. This finding is opposite to behavioral lab results and calls into question the assumption that advertising only affects the consideration of the advertised product in many economic models of informative advertising.

Besides its interest to academics, these findings will also be of interest to advertisers. Display advertisers should highlight their differential value knowing consumers will search for other products in the category. When their competitors are conducting large-scale display advertising campaigns, brands will want to develop a marketing communication strategy responding to the ad because the ad will subsequently cause consumers to search for their products. Assuming searches are meaningfully correlated with sales, display advertisers may find our methodology of tracking the effects of display ads on branded search to measure ad effectiveness useful, especially because measuring the effect of ads on sales is generally quite difficult (Lewis and Rao, 2012). Finally, we show that search advertisers should consider adjusting their bids in the search ad auction or at least have search ads when either they or their competitors run display advertisements on heavily trafficked sites such as Yahoo!.

Publishers and regulators may also find these results worth considering. The heterogeneity of search lifts across competitors may provide one way to quantify relative competition with the advertiser among competitor firms. We also find that an extension of Grossman and Shapiro (1984) suggests that positive spillovers can lead to reduced advertising investment, lower publisher revenues, and socially suboptimal advertising. However, more research is required to understand the total economic implications of the positive spillovers on both the publisher's inventory and the advertising marketplace as a whole.

The paper is organized as follows: Section 2 reviews related literature. Section 3 details the natural experiment, describes the display campaigns and search queries, and confirms the randomization. Section 4 describes the empirical models, presents the main results, and summarizes our robustness checks. Section 5 discusses the findings and concludes.

1.2 Related Literature

Our hypothesis comes from research suggesting that ads motivate consumers to learn and, hence, search for more information about the product category.[3] In a theoretical model, Mayzlin and Shin (2011) show that there exists a separating equilibrium where a high quality firm would choose an advertisement devoid of any attribute information in order to invite the consumer to search. In laboratory settings, Swasy and Rethans (1986) found that advertising for new products created more curiosity among consumers with high product category knowledge. Menon and Soman (2002) show that curiosity-inducing advertising increased time spent and attention on gathering information but did not increase the number of clicks on links for more information.

Our key findings show that some advertisers' brand advertising creates significant increases in searches for competitors' brands, but past lab research, perhaps limited by sample sizes or appropriate field outcomes, suggests that there should be little to no increase. Alba and Chattopadhyay (1985) found that cueing a brand inhibited recall of other brands in the same product category and brands in related categories. Nedungadi (1990) found that priming of a minor brand in a less accessible subcategory significantly increases the retrieval and consideration of the major brand in the same category while priming of a major brand did not significantly increase the recall or consideration of any other brands. In contrast, we find that brand advertising increases searches for competitor brands with both higher and lower market share than that of the advertiser.

There is a growing academic literature studying the effectiveness of display advertising. Recent papers have mostly focused on the effects from frequency and targeting on click-through rates and survey response data. Drèze and Hussherr (2003) found results from eye tracking data that suggest users avoided looking at display ads, but they found that multiple deliveries of a display ad to the same user increased unaided brand recall. Also looking at the effects of ad frequency, Lewis (2010) found, using the same exogenous variation as this

[3]The information provided by display advertising is limited by the nature of the medium.

paper, that significantly increasing the frequency of an advertiser's display ads on Yahoo!'s front page generally only modestly diminished the click-through rate. Goldfarb and Tucker (2011b) studied how the EU regulation that limited the targeting capabilities of advertisers reduced display advertising's positive effect on surveyed purchase intent while Goldfarb and Tucker (2011a) showed that both matching an ad to the website content and increasing its obtrusiveness increased surveyed purchase intent.

The only empirical study that we know of that tries to measure the effects of advertising on online search behavior using field data is by Joo et al. (2011). They show that there is a significant correlation between television advertising for financial services brands and consumers' tendencies to search online for these brands. We add to this literature by using a natural experiment and show that display advertising causes consumers to search online for both the advertised brand and other brands in the same product category.

2 Experimental Design and Data Overview

2.1 The Experiment and Data Collection

The empirical approach leverages a natural experiment with display ads on Yahoo!'s front page, www.yahoo.com. Yahoo! participates as a publisher in both the search and display advertising markets. In 2010, approximately 34% of Yahoo's revenue came from display while 50% come from search.[4] The front page is one of its most heavily trafficked sites. Approximately 40 million unique users visit the front page on any given day. Figure 1 is a screenshot of the front page and shows a Progressive Insurance display ad from our study. We collected anonymized data on users' ad exposures on the front page and subsequent search behaviors.

On any given day, an advertiser can purchase a display ad on Yahoo!'s front page for the whole day (i.e., all impressions) or purchase a spot in an ad split (i.e., half of the day's

[4] Yahoo! 10-K FY 2010.

impressions). If an advertiser chooses to be part of a split, its ad is shown to visitors who arrive at Yahoo!'s front page on, for example, any odd second, while another advertiser's ad is shown in the same position on the webpage to those who arrive at the site on any even second. As a result, each user has a 50% chance of being exposed to either ad in the split during each page visit. In addition, the number of impressions of a given ad delivered to a user will be distributed $binomial(v, .5)$ where v is the total number visits to the front page. If there is no systematic difference between odd and even second visits then the ad shown at each impression and the percentage of the total number of impressions delivered to a user are exogenous–essentially random.[5] Therefore, front page ad splits provide a natural experiment with which we can study the effects of advertising.

Each ad impression defines a treatment or control group event. Specifically, we associate all behavior immediately following the delivery of the target ad with our test group and all behavior following the delivery of the other split ad with our control (baseline) group. A period begins upon the delivery of an impression and ends either when another impression is delivered to the same user or after ten minutes have elapsed, whichever comes first.[6] Ten minutes should be long enough to allow for users to perform enough activities after the delivery of an impression but short enough to avoid misattributing activities to the delivered impression.[7]

We recorded each anonymous user's searches on Yahoo! during each ten-minute period after an impression was delivered. We focused on searches related to the brands of the advertised products and their competitors' brands. A search was defined as related to the

[5]We refer the interested reader to Lewis (2010) who discusses this experiment and its virtues and failings in greater depth.

[6]Only 18% of impressions were followed by another impression from the study within ten minutes.

[7]Examining behavior during the ten minutes following a user's exposure to a display ad impression yielded the most statistically powerful results in independent tests. While we may underestimate the total number of searches caused by the display ads, we intentionally make this tradeoff to focus on the time window when the signal to noise ratio is sufficiently large to make meaningful statistical measurements. The ten-minute limit still captures a majority of all exposed users' searches due to the one-day nature of the campaign. Multiday campaigns can experience greater gains in statistical power from imposing the ten-minute window, motivating its use in this study.

advertiser's or a competitor's brand if it included the name of the brand or its product.[8]

Finally, we removed from our data the 0.02% of users who visited the front page more than 100 times. Such extreme behavior is more characteristic of webcrawlers (i.e., computer programs) than humans.[9] Less than one percent of users visited the front page more than 30 times.

2.2 Advertising Campaigns and Search Keywords

The three display advertising campaigns were for Progressive Car Insurance, the Acura TSX sports wagon, and the Samsung Galaxy Tab. Each campaign was part of a front page split with a control display ad unrelated to the advertised product category. The Progressive campaign was part of a split with an ad for a TV show; the Acura TSX campaign was part of a split with an ad for an electronic retail store; and the Samsung Galaxy Tab campaign was part of a split with an ad for an action movie. Table 1 shows the creatives for each display ad campaign and the corresponding control display ads in their splits.

To capture as much of the complementarities between display and search advertising, we wanted to consider all the brands in the advertised product category or market as search keywords. To identify other category brands for the Progressive Auto Insurance campaign, we used a December 2009 Mintel Marketing Intelligence research report about the Auto Insurance Industry.[10] In total, we found 14 brands other than Progressive Auto Insurance.[11] These competitors include all the top 10 underwriters, which make up over 67% of premiums written in 2008 (Mintel Reports Database, 2009).

[8]We acknowledge that this method will incorrectly categorize some searches as searches for the brand's product, but instead, they are really searches for topics related to the brand like queries for financial statements for the brand. After scanning some of the categorized searches, we believe that most if not all of the categorized searches are searches for the brands' products.

[9]The maximum number of visits to the front page by a unique user in a single day was around $4,000$, averaging once every 15 seconds during waking hours.

[10]The report contains the results from a survey asking adults 18 and over whose households own an insured vehicle about their preferred auto insurance company. We took those that were named as other brands in the market.

[11]The brands other than Progressive are State Farm, Allstate, Geico, Farmer's Insurance, Nationwide Insurance, Liberty Mutual, USAA, AIG, American Family Insurance, 21st Century Insurance, Travelers Insurance, Hartford AARP, Erie Insurance, and Safeco.

For competitor-branded search keywords for the Acura campaign, we used all makes listed on Autobytel.com, a consumer website about passenger vehicles.[12] In total, we identified 36 brands other than Acura that range from standard brands such as Ford to sports and luxury brands such as Porsche.[13]

Finding the category brands for the Samsung Galaxy Tab campaign was more difficult. During June 2011, the month of the ad split, the tablet PC market was fairly new and dominated by one product, Apple's iPad; however, new products were entering the market on a monthly basis. To summarize these options and upcoming releases, CNET, a media site about technology products, wrote an article titled "CNET looks at current and upcoming tablets" (Franklin, 2011). In this article, they listed all upcoming and currently available tablet PCs from March to the end of 2011.[14] From the article, we used the 15 products and brands listed as potential search keywords.[15]

For the three advertising campaigns, we have identified many brands in each advertised markets, but we might have ignored some close substitutes. For example, the Samsung Galaxy Tab also competes with e-readers, smartphones, netbooks, and ultra-portable laptops. By ignoring these product market substitutes, our estimates are more conservative and underestimate the magnitude of the possible search spillovers.

2.3 Data Summary

Table 2 summarizes the data. There were 38-41 million unique users who visited the front page 161-171 million times each day of the splits. As expected by the natural randomization created by the ad split, the target ad was delivered on 50% of those visits while the control

[12] Autobytel (2011) retrieved June 30, 2011.

[13] The brands other than Acura are Audi, BMW, Cadillac, Chevrolet, Dodge, Ford, GMC, Honda, Hyundai, Infiniti, Jeep, Kia, Lexus, Lincoln, Mazda, Mercedes, Mini, Mitsubishi, Nissan, Scion, Subaru, Suzuki, Toyota, Volkswagen, Volvo, Buick, Chrysler, Fiat, Jaguar, Land Rover, Porsche, Rolls Royce, Saab, Smart, and Tesla.

[14] Retrieved July 29, 2011.

[15] The brands other than Samsung's Galaxy Tab are Dell Streak, Motorola Xoom, HTC, Blackberry Playbook, Asus, Acer Iconia, Apple iPad, Amazon Tablet, Sony, G-Slate, Toshiba, Vizio, Viewsonic, and Micro Cruz.

ad was delivered on the other 50%.

A small fraction of users searched for any category brands on the day of the campaign. Only 0.06% of front page visitors searched for the advertised brands and its competitors on the day of the Progressive Auto Insurance ad campaign, 0.8% searched on the day of the Acura TSX ad campaign, and 0.04% searched on the day of the Samsung's Galaxy Tab ad campaign. In spite of these small percentages, the scale of the Yahoo! front page still yielded a large number of users who search: $15,545$, $331,047$, and $24,258$ users, respectively, for the three campaigns.

Most visitors to Yahoo!'s front page in a day return at least once that same day. Figure 2 shows a histogram of the total number of users visiting for each day of the three campaigns; the histogram is heavily skewed right with the mean visits in a day exceeding the median of two. On average, a user visits the front page four times during the day and, as a result, sees four front page display ad impressions, randomly split between the target and control ads.

Figure 3 shows the distribution of the number of target ad exposures delivered to each user for the three campaigns. This distribution is also heavily skewed right. Given N, the number of front page visits by a user, the number of target ads shown to a user should be distributed binomial(N, $p = 0.5$). Figure 4 compares four histograms: three showing the empirical distribution of target ad exposures for users who visited the front page ten times during the three campaigns and, for comparison, one plotting the binomial distribution for $N = 10$ and $p = 0.5$. The similarity of the sample distributions to the theoretical distribution confirms the impression-level randomization mechanism used by the ad splits.

3 Empirical Analysis and Results

3.1 Econometric Model

We measured the impact of display advertising on branded search from the display impression and search query data we collected for each advertiser's campaign. Our analysis compares the search probability for a brand between our test and control groups. To measure the difference in probabilities, we estimated the linear probability model

$$Search_{ijt} = \alpha_j + \beta_j AD_{it} + \epsilon_{ijt} \tag{1}$$

using ordinary least squares (OLS) for each brand j on the data collected from the ad campaign relating to brand j's product category. $Search_{ijt}$ is an indicator variable equal to one if user i searched for brand j during period t and zero otherwise. AD_{it} is another indicator variable equal to one if the target ad was delivered to user i during period t. ϵ_{ijt} is the error term.

The estimated β_j measures the average causal effect of front page display advertising exposure on the tendency to search for brand j within the ten-minute time window. This estimate would be biased if AD_{it} were correlated with ϵ_{ijt}, but this would require users to be systematically more or less likely to search when they visit the front page on an even second relative to an odd second. On the other hand, we might underestimate the total impact by not considering all searches caused by the display ad. For example, if a user's first exposure continues to influence their searching behavior all day long, and not just during the limited ten-minute window, we will underestimate the impact of the ad on search behavior by ignoring later searches caused by the ads and misattributing some of those searches to the baseline search behavior .

We estimated standard errors by clustering the errors on the user level to account for random effects and any autocorrelation between periods. For example, a user who sees a headline of interest on the front page might spend an entire day researching that topic, and therefore, may be less prone to search in the advertised category (or any category).

We are interested in understanding the impact of a full day of advertising on Yahoo!'s Front Page, www.yahoo.com, not the impact of the split. We do this for expositional simplicity. All of our main estimates below present the results for the counterfactual, "What would the ad effect have been had the advertiser purchased the full day of advertising?" We compute these estimates simply by multiplying the estimates by the total number of impressions shown from both advertisers (i.e., twice the impressions shown in an ad split). We pause to point out that wearing out of the effects, which could complicate this simplified counterfactual computation, are found to be of modest importance.[16] We focus on the conceptually simpler target versus control comparison: "What did people search for after seeing one display ad rather than the other?"

3.2 Advertiser and Competitor Search Lifts

Table 3 contains the results from estimating equation (1) for each of the three ad campaigns. We have transformed the linear probability model's probability estimates to report the display ad's impact in terms of number of searches for the advertised brand and for any of its competitors' brands. We also separately report estimated effects for the competitors' brands with the four highest t-statistics. The baseline estimates, representing the total number of searches by all visitors that would have occurred if only the control display ad were delivered, are the estimated intercepts multiplied by the total number of impressions delivered on that day. The second column of estimates shows the estimated number of incremental searches over the baseline from showing the target display, computed by multiplying the estimated lifts in the propensity to search with the total number of impressions delivered. Both t-statistics based on the OLS standard errors and clustered standard errors are reported by each of the estimates.

In all three campaigns, we found a statistically significant lift in searches for the advertised brand from its display ad ($p < 0.001$ for all three campaigns). Both Acura and Galaxy Tab ad campaigns caused a 44% lift in searches for their brands while the Progressive ad campaign only caused a 28% lift in searches for its brand. If the advertised brands decided to run their campaigns for the full day then Samsung would have had 424 more searches, Acura would

[16] The frequency model's results are available in an earlier draft of this paper. They do not show significant effects of decreasing marginal returns, but the confidence intervals grow in the number of impressions.

have had $1,555$ more searches, and Progressive would have had $1,135$ more searches relative to not advertising.

Not only did we find search lifts for the advertised brand but also for its competitors. Two of the three campaigns caused a significant increase in searches for competitors. The Samsung Galaxy Tab ad campaign caused a 6.0% increase in searches for any of its competitors. Most of this increase was due to searches for Apple's iPad. If the Galaxy Tab ad campaign were run for the full day then it would have produced 857 more searches for the iPad. This number is about two times larger than that for the Galaxy Tab, 424. In essence, at least for driving online searches, the Galaxy Tab ad campaign was more of an advertisement for the iPad than it was for the Galaxy Tab. The other product that had a statistically significant increase in searches from the Galaxy Tab display ad was the Motorola Xoom. Though not as big of an increase as that of the iPad in absolute terms, the percentage lift for Xoom searches was 22.8%.

Acura's ad campaign also caused an increase in searches for its competitors, though the estimated effect was smaller at 3.0%. Note that the relative sizes of the baselines matter–with 36 competitors to account for, some of which are much larger than Acura–the total search impact for all the competitors combined, in terms of number of searches, is approximately 7.7 times larger than for the advertiser, Acura. Again, we see that advertising produced more searches for its competitors than for its own brand. The brands with the largest increase in absolute terms are Volkswagen, Hyundai, Lexus, and Volvo with effects ranging from 478 to 894 more searches.

In contrast, the Progressive Auto Insurance ad campaign did not noticeably increase searches for its competitors. A possible reason for this is that auto insurers' websites also quote their competitors' prices. Both the creative in the campaign and Progressive's website suggest that Progressive will present users with prices from its competitors. Therefore, users do not need to incur the cost of searching for information about competitors if they obtain that information from the advertised brand.

These differences in search behavior are for Yahoo! Search only. We would expect users to perform any incremental searches using their preferred search engine: whether visiting a search engine's website, typing in the search box in their web browser, or using a browser toolbar. Holding the search channel constant, we would expect to find a similar relative composition of findings across search methods.

3.3 Robustness Checks

We conducted two sets of robustness checks of our analysis. In the first set, we re-estimated display advertising's effects on searches but limited the sample in one analysis to only the period after the first impression and limited another analysis to those users who only saw one impression. By limiting the analysis to a user's first impression, the issues about misattributing searches to subsequent ad impressions and correlated error terms are no longer a concern. Also, any difference between the estimates from the full sample and those from the limited samples would suggest that frequency might moderate the effect.

Table 4 contains the results from the full sample and the limited samples. The point estimates from the limited samples have the same sign and tell the same story as those from the full sample, but they are generally smaller. However, they are not significantly different at the 5% level. That said, we cannot conclude that effect of the first impression is not different from the effects of subsequent impressions because the powers of these tests are low.[17] We greatly reduce the sample size from 171 million to 40 million when limiting to first impressions and to 13 million when limiting to those who saw only one impression.

In the second set of robustness checks, we looked at the most statistically impacted words, queries, and domains clicked. The purpose of these robustness checks is to ensure that these extra searches for the brands are related to product searches and not to other type of searches such as the financial news about the brand. To find the most statistically impacted words, queries, and domains clicked, we calculated the frequency of those outcomes

[17]First impressions (no pun intended), however, do generally occur earlier in the day and users who see only one impression may not be representative of the general user base.

under the control ads and under the test ads and looked at those with differences with the highest t-statistics.[18]

The results of this robustness check corroborate our brand analysis while providing deeper insights. As expected, the display ad caused an increase in clicking on the brands' domains from both organic and sponsored links. For example, the Acura display ad increased clicking on links to Acura.com, VW.com, Hyundai.com, Lexus.com, and Volvocars.com, and the Galaxy Tab's display ad increased clicking on links to store.apple.com. The t-statistics for these increases were greater than 3. Interestingly, the display ad also statistically significantly increased clicks on search links for trade, review, and distribution websites. For example, Acura's display ad increased clicks for Motortrend.com, Edmunds.com, and autobytel.com, and Galaxy Tab's display ad increased clicks for reviews.cnet.com, besttablet2011.com, and bestbuy.com. The increase in clicks for these "market information" websites speaks to the literature that suggest advertising invites and motivates consumers to search for more information (Mayzlin and Shin, 2011). This is true not only for information about the advertised brand but also about the whole product market.

We also looked at the most positively impacted queries and words. For the Progressive ad, the most impacted query was "progressive flo." Flo is the name of the woman in the Progressive ad. This implies that a large portion of the estimated search lift for Progressive was searches intended for Flo and perhaps not for Progressive Auto Insurance. When looking at the most impacted search words and queries due to the Acura ad, we find that the Acura ad increased searches for specific car models besides just the make of the cars. For example, the Acura ad statistically significantly increased searches for Hyundai Sonata, Ford Flex, Nissan Optima, and Honda Crosstour. Overall, these checks echo the same story about positive spillovers suggested by our previous results.

[18] We use a high level of statistical significance in order to reduce the risk of false positives. For example, for the Acura ad, we analyzed 48,656 subdomains that received at least 40 clicks in our sample; only 2 of the top 25 p-value-ranked subdomains were not obviously car-related and were considered false positives. This observation combined with the fact that car-related searches account for a small share of all online search behavior lead us to expect car-related false positives to have no meaningful effect on our findings.

4 Discussion and Conclusion

4.1 Display and Search Advertising Complementarities

Even though we do not measure the revenue generated by these spillovers, the spillovers do create complementarities between display and search advertising through the cost of advertising. This is due to the generalized second price auction. Appendix A presents a simple model and derives necessary conditions for display and search advertising to be complements. Our findings mainly depend on two necessary conditions. Our empirical finding of display ads increasing search volume satisfies one condition, and another condition is that a search ad's cost-per-click (CPC) is weakly increasing in the desired click-through rate (CTR). The second condition is a property of search ad auctions. In practice, ads positioned at the top of the page, where click-through rates are higher, cost more per click than ads at the bottom of the page.

The intuition regarding the complementarities of display and search ads is simple: if display advertising increases the number of searches, then holding revenues constant by holding the number of clicks constant permits an advertiser to bid for a lower CTR. Lowering the CPC directly translates into a larger profit margin for each search click; hence, display and search ads are complements. A very similar calculation and logic applies to a firm's display ads and the competitors' search ads because the same opportunity to reduce the CPC applies as much to the competitors as to the display-advertising firm. This larger profit margin for competitors' search ads implies that display and search ads are strategic complements.

To illustrate our point, we use the bids on February 12, 2012 for the branded keywords of the three advertisers and calculated the CPC for each position. Figure 5 shows the relationship between CPC and CTR on the search result pages for the three advertised brands: Progressive, Acura, and Samsung. Due to confidentiality, CPCs are in terms of the percentage of the top position, position 1. CTRs for the four search ad positions are averages of CTRs from a sample of queries with at least four ads that are found in Reiley et al. (2010). They are also expressed in terms of percentage of the top position. For an example of the cost savings, let's consider searches for Acura. If an advertiser decided to move from position 2 to position 3, the advertiser lowers its CPC by 27%, but loses 30% clicks by moving to a lower CTR position. However, the 44% increase in searches due to Acura's display ad on the front page will make up the loss in clicks from the lower CTR. Therefore, the advertiser in position 2 can save 27% per click for the same number of clicks by moving down to position 3 when Acura has a display ad on Yahoo!'s front page.

Beyond the complementarities discussed above, there might be other complementarities and strategic complementarities about which we can only speculate. For example, if the advertiser chooses to advertise on its search result page, the display advertiser can reinforce the marketing message, provide more product information, and advertise distribution channels. In regards to strategic complements, Burke and Srull (1988) show that consumers' recall of an ad's brand information is hindered by exposure to a competitor's ad. So, if a display ad causes users to search, a prominently displayed search ad for a competitor can attenuate the recall of the display advertiser's marketing message.

Switching our focus from the advertiser to the publisher, our results suggest that there might be revenue spillovers from the display advertising market to the search advertising market. Economides and Salop (1992) show, in the context of imperfect competition, that deadweight loss is smaller in a market where a single firm sells two complementary products relative to the market where the two products are sold by two separate firms. If the value of these spillovers is economically significant, this model suggests that allowing publishers to participate in both display and search advertising markets could reduce deadweight loss. Given the data we have, we are not able to accurately measure the magnitude of these spillovers between display and search advertising.[19] Answering this question is beyond the

[19]A previous version of this paper attempted to quantify the value of these spillovers and found them to be somewhat modest–worth perhaps single-digit percentage points relative to the value of the display advertising cost. However, several weaknesses of that analysis included considering only a very restricted set of brand-related queries and not having the exact incremental revenue from the incremental search clicks.

scope of this paper and is an important area for future research.

4.2 Advertising Investment and Competitive Spillovers

Our results show that there are positive spillovers of advertising to other competitors, especially in awareness. Most theoretical models on informative advertising usually assume that advertising only increases the awareness of the advertised product (e.g. Butters (1977) and Grossman and Shapiro (1984)), and our results conflict with this assumption. Using the Grossman and Shapiro (1984) model and allowing for advertising spillovers, we show in Appendix B that introducing symmetric positive spillovers in advertising decreases investment in advertising (see Figure 6). Intuitively, every additional unit of advertising creates more price competition in the product market when there are spillovers than when there are not.

In regards to the socially optimal level of advertising, the original Grossman and Shapiro (1984) model predicts that there can be either an underinvestment or overinvestment in advertising depending on the parameters of the model. In our extension, if firms are over-investing, positive spillovers shift advertising investment levels toward the socially optimal level. In contrast, spillovers make underinvestment worse.

Kaiser et al. (2005) document several marketing programs that attempted to reduce such advertising underinvestment for agricultural commodities such as "California Raisins." In discussing their work, Varian (2006) states,

> [Profitable commodity advertising] programs can be difficult to maintain. The problem lies in aligning incentives. The producer of a branded product pays all the costs and reaps all the benefits of its advertising spending, leading it to a carefully considered decision about how much to spend.
>
> By contrast, the benefits and costs of commodity advertising are spread unequally among many producers, making it tough to reach a collective decision about marketing levels.

Our measurements and model of positive spillovers suggest that brand advertisers may *not*, in fact, reap all of the benefits of their advertising. As a result, even the significant returns to the commodity advertising programs documented in Kaiser et al. (2005) may still be socially suboptimal because advertisements for branded "California Raisins" may have resulted in positive spillovers to non-California raisin producers or even other related products. Future research should explore the consequences of spillovers in awareness on advertising investment more broadly to understand its adverse welfare consequences and seek to develop mechanisms to mitigate this market failure.

4.3 Conclusion

Measuring the spillovers of online display advertising on online search is becoming increasingly important as online advertising continues to grow. We believe these spillover results offer a fascinating view into the effects of advertising on customer behavior. Will researchers find proportional effects on sales and profits? Are there similar spillovers from all other forms of advertising media? Are the spillovers illustrated in online search a proxy for customer search behavior more generally? Are online search queries a proportional representation of causal attention induced by the ad? While these questions may sound esoteric, an increasing number of day-to-day activities are assisted by online search engines paired with other digital technologies such as smart phones and tablet computers. Online search is now in your pocket to help you obtain information about who you know, where you want to go, and what you want to buy. We hope to see future research explore these and other related questions that will help advertisers and publishers improve the effectiveness of advertising and the efficiency of advertising marketplaces.

References

Alba, J. and Chattopadhyay, A. (1985). Effects of context and part-category cues on recall of competing brands. *Journal of Marketing Research*, pages 340–349.

Autobytel (2011). Autobytel - new cars. Retrieved Jun. 30, 2011: `http://www.autobytel.com/new-cars`.

Burke, R. and Srull, T. (1988). Competitive interference and consumer memory for advertising. *Journal of Consumer Research*, pages 55–68.

Butters, G. (1977). Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies*, pages 465–491.

Drèze, X. and Hussherr, F. (2003). Internet advertising: Is anybody watching? *Journal of Interactive Marketing*, 17(4):8–23.

Economides, N. and Salop, S. (1992). Competition and integration among complements, and network market structure. *The Journal of Industrial Economics*, pages 105–123.

Franklin, E. (2011). Cnet looks at current and upcoming tablets. Retrieved Jul. 29, 2011: `http://news.cnet.com/8301-17938_105-20037960-1/cnet-looks-at-current-and-upcoming-tablets/`.

Goldfarb, A. and Tucker, C. (2011a). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404.

Goldfarb, A. and Tucker, C. (2011b). Privacy regulation and online advertising. *Management Science*, 57(1):57–71.

Grossman, G. and Shapiro, C. (1984). Informative advertising with differentiated products. *The Review of Economic Studies*, 51(1):63–81.

Joo, M., Wilbur, K., and Zhu, Y. (2011). Does television advertising influence online search? (september 24, 2012). Available at SSRN 1720713: `http://ssrn.com/abstract=1720713`.

Kaiser, H. M., Alston, J. M., Crespi, J. M., and Sexton, R. J. (2005). *The economics of commodity promotion programs: Lessons from California*. Peter Lang Pub Incorporated.

Lewis, R. (2010). *Where's the "Wear-Out?": Online Display Ads and the Impact of Frequency*. PhD thesis, MIT Dept of Economics.

Lewis, R. and Rao, J. (2012). On the near impossibility of measuring advertising effectiveness. *Working paper*.

Lewis, R., Rao, J., and Reiley, D. (2011). 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*, pages 157–166. ACM.

Mayzlin, D. and Shin, J. (2011). Uninformative advertising as an invitation to search. *Marketing Science*, 30(4):666–685.

Menon, S. and Soman, D. (2002). Managing the power of curiosity for effective web advertising strategies. *Journal of Advertising*, pages 1–14.

Mintel Reports Database (2009). Auto insurance. Retrieved Aug. 1, 2011: http://academic.mintel.proxy.uchicago.edu/.

Mintel Reports Database (2010). Attitudes to internet and new media marketing. Retrieved Sept. 8, 2011: http://academic.mintel.proxy.uchicago.edu/.

Nedungadi, P. (1990). Recall and consumer consideration sets: Influencing choice without altering brand evaluations. *Journal of consumer research*, pages 263–276.

Reiley, D., Li, S., and Lewis, R. (2010). Northern exposure: a field experiment measuring externalities between search advertisements. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 297–304. ACM.

Swasy, J. and Rethans, A. (1986). Knowledge effects on curiosity and new product advertising. *Journal of Advertising*, pages 28–34.

Varian, H. R. (2006). Advertising commodities can be tricky, but it does pay off. Retrieved Nov. 12, 2012: http://people.ischool.berkeley.edu/~hal/people/hal/NYTimes/2006-06-01.html.

Appendix A: Complements and Strategic Complements

Consider the following model for a firm's profits as a function of the quantity of advertising:

$$\Pi\left(A_d, A_s\right) = A_d \cdot v_d + A_s \cdot v_s - A_d \cdot P_d - A_s \cdot P_s\left(A_s/Q_s\left(A_d\right)\right)$$

where A_d and A_s are the quantity of display and search advertising denoted in impressions and clicks, respectively, and Q_s is the total number of searches. v and P refer to the expected value and price of the advertising, A. In the case of display advertising, P_d would be denoted in cost per mille (CPM), meaning the price of 1,000 display impressions. For search advertising, the generalized second-price auction is used to sell search ad placements on search result pages. In the auction, an advertiser determines a desired click-through rate (CTR) and bids sufficiently high to achieve the ranking that yields that CTR. Here, we write P_s is the average price paid to obtain a CTR of $A_s/Q_s\left(A_d\right)$. A necessary condition for the existence of an equilibrium in the auction requires that the price paid is increasing in the CTR, $P_s' > 0$. In order for search and display advertising to be complements, marginal profits from search advertising must be increasing in the volume of display advertising. We derive the cross partial derivative of the profit function below.

$$\begin{aligned}
\frac{\partial \Pi}{\partial A_s} &= v_s - P_s\left(\frac{A_s}{Q_s\left(A_d\right)}\right) - A_s \cdot P_s'\left(\frac{A_s}{Q_s\left(A_d\right)}\right) \cdot \frac{1}{Q_s\left(A_d\right)} \\
\frac{\partial^2 \Pi\left(A_d, A_s\right)}{\partial A_s \partial A_d} &= \frac{Q_s'\left(A_d\right)}{Q_s\left(A_d\right)^2}\left(\left(1+A_s\right) \cdot P_s'\left(\frac{A_s}{Q_s\left(A_d\right)}\right) + \frac{A_s}{Q_s\left(A_d\right)} \cdot P_s''\left(\frac{A_s}{Q_s\left(A_d\right)}\right)\right) > 0
\end{aligned}$$

A key finding from our empirical work reveals that $Q_s' > 0$: we find that the number of searches, $Q_s\left(A_d\right)$, is an increasing function of display advertising, A_d. The other necessary conditions for search and display advertising to be complements are that $P_s' > 0$, which follows from the auction's structure, and $P_s'' > -\frac{1+A_s}{A_s} \cdot Q_s\left(A_d\right) \cdot P_s'\left(\frac{A_s}{Q_s(A_d)}\right)$, which bounds the degree of concavity for the average price paid as a function of the CTR.

The same derivation used for complements now applies to strategic complements. To simplify notation, we write the profit function with only the relevant strategic variables:

$$\Pi^j\left(A_d^i, A_s^j\right) = A_s^j \cdot v_s^j - A_s^j \cdot P_s^j\left(A_s^j / Q_s^j\left(A_d^i\right)\right).$$

We write i for the display advertiser and j for competitors. As the strategic variables, we consider the quantity of display advertising, A_d^i, and the competitor's search advertising, A_s^j. Again, P_s^j is the average price paid per click to achieve the CTR of $A_s^j/Q_s^j\left(A_d^i\right)$ and is increasing for all market participants: $P_s' > 0$. From the empirical results we know that the number of competitors' searches is increasing in the firm's display advertising, $Q_s^{j\prime}\left(A_d^i\right) > 0$. Given these identical comparative statics, the strategic complementarity of display and search advertising follows from applying a superscript j to all variables other than A_d, to which one applies an i superscript:

$$\frac{\partial^2 \Pi^j\left(A_d^i, A_s^j\right)}{\partial A_s^j \partial A_d^i} > 0.$$

Appendix B: Ad Awareness Investment and Spillovers

In Grossman and Shapiro (1984), a unit mass of consumers are uniformly distributed over a line of unit length. Each consumer has a reservation price R for the product but suffers a transportation cost τ for each unit distance away from his ideal point. There are two firms in the market and position themselves at opposite ends of the line. A consumer knows the existence of a firm's product and its price if he is aware of it.

In the original Grossman and Shapiro (1984) model, consumers are only aware of the firm if it receives an ad from it, but we depart from this model by assuming that a firm's ad will also increase the awareness of the other firm in some consumers. Formally, firm i chooses the level of advertising by choosing percentage of the unit mass that would receive its ad, ϕ_i. However, every unit of advertising will increase the awareness level of the other firm by δ. This δ represents the values in the "Competitor/Own" column in Table 3. If there are no spillovers from advertising, $\delta = 0$. The resulting demand curve for firm i's product is

$$D_i(P_i, P_{i'}, \phi_i, \phi_{i'}) = (\phi_i + \delta\phi_{i'}) \left((1 - (\phi_{i'} + \delta\phi_i)) + (\phi_{i'} + \delta\phi_i) \frac{P_{i'} - P_i + t}{2t} \right). \tag{2}$$

In the game, both firms simultaneously choose prices and advertising levels that maximize their profits, $D(P_i, P_{i'}, \phi_i, \phi_{i'})(P_i - c) - \frac{a}{2}\phi_i^2$ where c is the marginal cost of production, $\frac{a}{2}\phi_i^2$ is the cost of advertising, and $a > \frac{t(1+\delta)^2}{2}$. Because costs, demand, and spillovers are symmetric, equilibrium prices and advertising levels should also be symmetric. That is $\phi_i^e = \phi_{i'}^e = \phi^e$ and $P_i^e = P_{i'}^e = P^e$. The results of the equilibrium profits and prices are

$$\begin{aligned} P^e &= c + \tau \frac{2 - (1+\delta)\phi^e}{(1+\delta)\phi^e} \\ \Pi^e &= \tau \frac{(2 - (1+\delta)\phi^e)^2}{2} - \frac{a}{2}(\phi^e)^2 \end{aligned} \tag{3}$$

where ϕ^e is the equilibrium advertising level. This level is

$$\phi^e = \frac{(2+\delta) - \sqrt{(2+\delta)^2 - 4\left[\frac{(1+\delta)^2 - 2\frac{a}{\tau}}{1+\delta}\right]}}{(1+\delta)^2 - 2\frac{a}{\tau}}. \tag{4}$$

In our empirical analysis, we found δ to be less than one for spillovers to most competitive brands, and as a result, we plot the equilibrium advertising, prices, and profits for $\delta \in [0, 1]$ respectively in Figure 6, Figure 7, and Figure 8. We find advertising levels fall as spillovers increase because for each unit advertising, advertising spillovers create more price competition. As for the equilibrium price, spillovers decrease it initially because the effect from the increase of price competition outweigh the effect from a decrease in advertisement, but eventually, the effect from the decrease in advertisement will dominate and cause the equilibrium price to rise.

Appendix C: Tables and Figures

Figure 1: Snapshot of the Frontpage

Table 1: Creatives for Ad Campaigns

<table>
<tr><th>Date of Ad Split</th><th>Target Ad</th><th>Control Ad</th></tr>
<tr><td>11 January 2011</td><td>
</td><td>
</td></tr>
<tr><td>10 February 2011</td><td>
</td><td>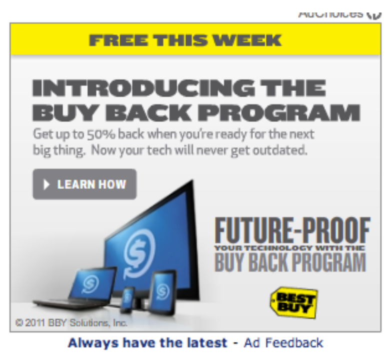
</td></tr>
<tr><td>29 June 2011</td><td>
</td><td>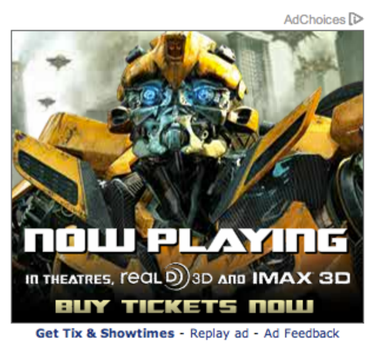
</td></tr>
</table>

Table 2: Summary Statistics

Variables and Statistics	Progressive Auto Insurance	Acura TSX	Samsung's Galaxy Tab
Date of Ad Split	2011/01/11	2011/02/10	2011/06/29
Sample Sizes			
Total Number of Unique Visitors	40,673,687	41,313,836	37,620,318
Total Number of Visits	171,953,331	171,593,781	161,460,200
Total Number of Exposures to the Target Ad	86,152,779	85,684,914	80,866,903
Percentage of Users Who Searched for Relevant Keywords	0.06%	0.80%	0.04%
Total Number of Visits per User			
Mean	4.23	4.15	4.29
Standard Deviation	6.59	6.53	6.64
Minimum	1.00	1.00	1.00
25th Percentile	1.00	1.00	1.00
50th Percentile	2.00	2.00	2.00
75th Percentile	5.00	5.00	5.00
Maximum	4,902.00	4,896.00	4,897.00
Total Number of Exposure to the Target Ad per User			
Mean	2.12	2.07	2.15
Standard Deviation	3.45	3.43	3.50
Minimum	0.00	0.00	0.00
25th Percentile	1.00	1.00	1.00
50th Percentile	1.00	1.00	1.00
75th Percentile	3.00	3.00	3.00
Maximum	2,448.00	2,453.00	2,447.00

Figure 2: Distribution of the Total Number of Visits

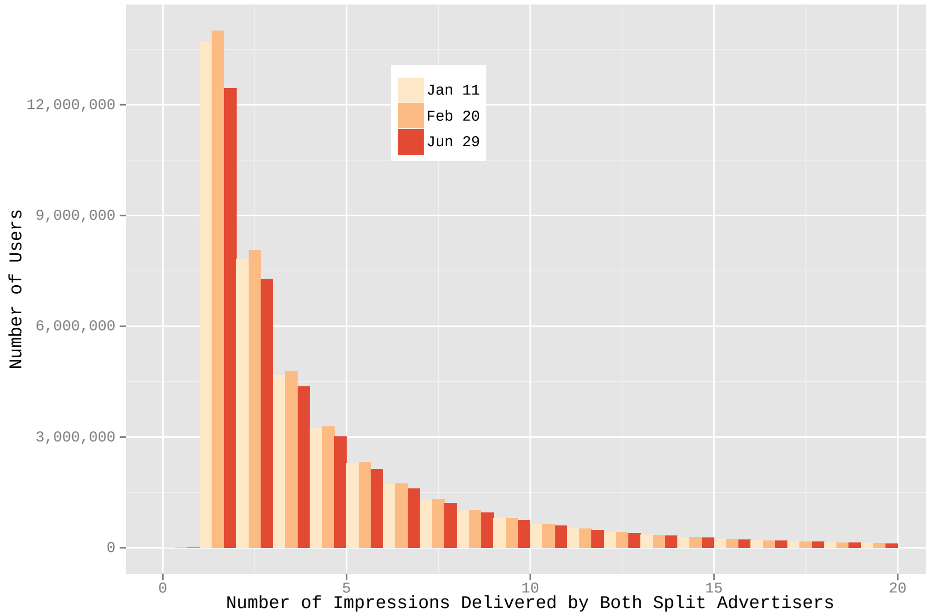

Figure 3: Distribution of the Total Number of Exposures to the Target Ad

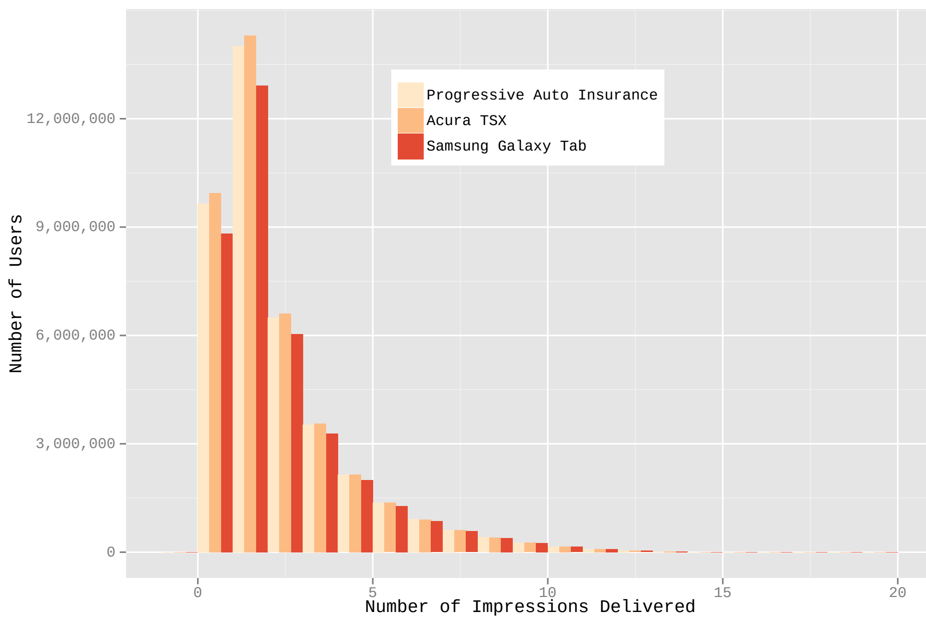

Figure 4: Distribution of the Total Number of Exposures to the Target Ad for Users Who Only Visited the Front Page 10 times

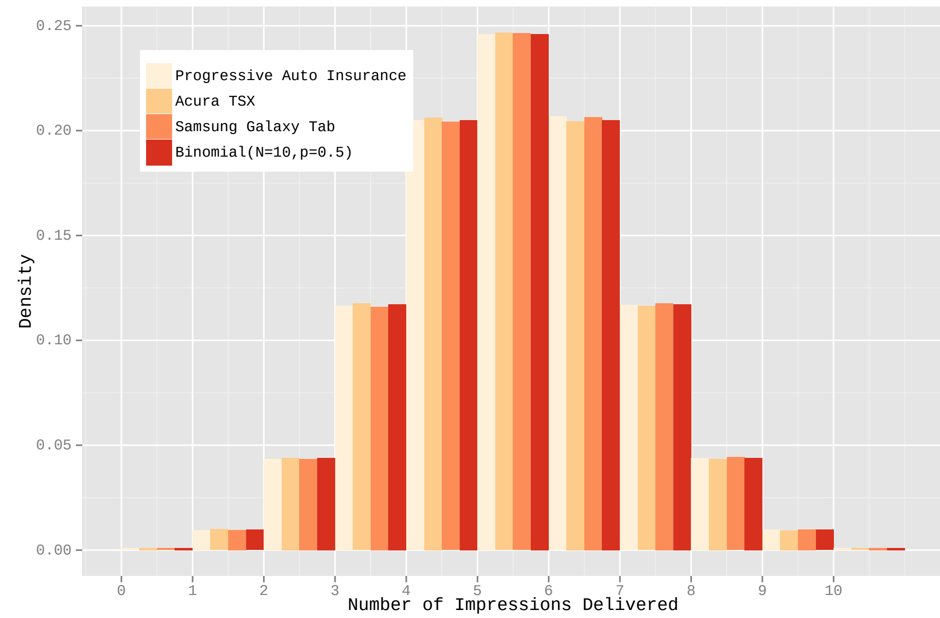

Figure 5: The Relationship between CPC and CTR on a Search Result Page

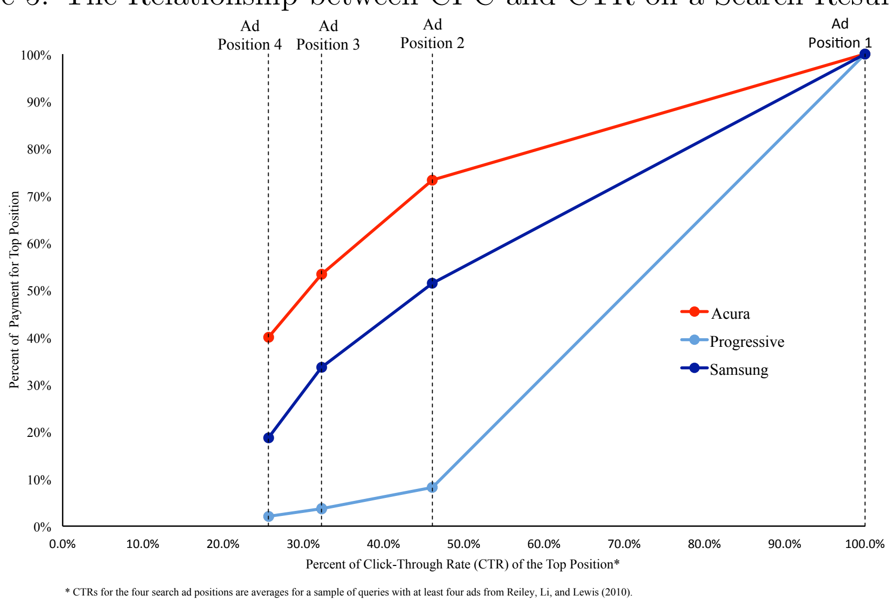

Table 3: Main Results of the Effects of Advertising on Search

Searches	Control: Estimate	Control: OLS T-stat	Control: Cluster T-stat	Search Lift from Advertising: Estimate	Search Lift from Advertising: OLS T-stat	Search Lift from Advertising: Cluster T-stat	Percentage Lift	Competitor/ Own
Samsung Galaxy Tab Advertising Campaign								
Samsung Galaxy Tab	958	19.78	20.57	424	6.20	6.32	44.3%	1.00
All Competitors	16,662	89.87	82.42	994	3.79	3.81	6.0%	2.34
Apple Ipad	9,851	68.64	63.21	857	4.23	4.25	8.7%	2.02
Motorola Xoom	663	17.23	16.74	151	2.79	2.79	22.8%	0.36
Blackberry Playbook	317	11.92	11.34	71	1.89	1.90	22.4%	0.17
Viewsonic	18	2.55	3.00	14	1.39	1.39	77.2%	0.03
Acura Advertising Campaign								
Acura	3,539	38.12	38.34	1,555	11.84	11.78	43.9%	1.00
All Competitors	401,927	445.80	389.84	12,035	9.43	9.44	3.0%	7.74
Volkswagen	5,840	52.12	48.24	894	5.64	5.62	15.3%	0.58
Hyundai	5,399	50.05	46.94	853	5.59	5.55	15.8%	0.55
Lexus	3,907	42.54	39.37	631	4.86	4.85	16.2%	0.41
Volvo	2,183	31.39	29.31	478	4.86	4.75	21.9%	0.31
Progressive Auto Insurance Advertising Campaign								
Progressive	4,104	42.41	42.76	1,135	8.30	8.34	27.6%	1.00
All Competitors	23,035	106.84	99.34	327	1.07	1.08	1.4%	0.29
Allstate	2,968	38.09	36.52	124	1.12	1.13	4.2%	0.11
USAA	7,870	62.30	56.97	187	1.05	1.06	2.4%	0.17
Safeco	214	10.01	9.73	29	0.96	0.97	13.6%	0.03
Nationwide Insurance	880	20.64	19.81	54	0.90	0.91	6.2%	0.05

Table 4: Robustness Checks Limiting the Sample to the First Impressions

	Full Sample	Limited to the first impression			Limited to users who were delivered one impression		
	Daily Total Search Lift	Daily Total Search Lift	Lower Bound 95% CI	Upper Bound 95% CI	Daily Total Search Lift	Lower Bound 95% CI	Upper Bound 95% CI
Samsung Galaxy Tab Advertising Campaign							
Samsung Galaxy Tab	424	503	190	817	-41	-603	521
All Competitors	994	257	-964	1,478	228	-1,893	2,348
Acura Advertising Campaign							
Acura	1,555	1,037	466	1,607	250	-724	1,224
All Competitors	12,035	10,161	4,259	16,062	7,437	-2,516	17,389
Progressive Auto Insurance Advertising Campaign							
Progressive	1,135	433	-234	1,100	607	-426	1,639
All Competitors	327	326	-1,225	1,877	237	-2,184	2,659

Figure 6: Magnitude of Spillovers vs. Equilibrium Advertising Levels

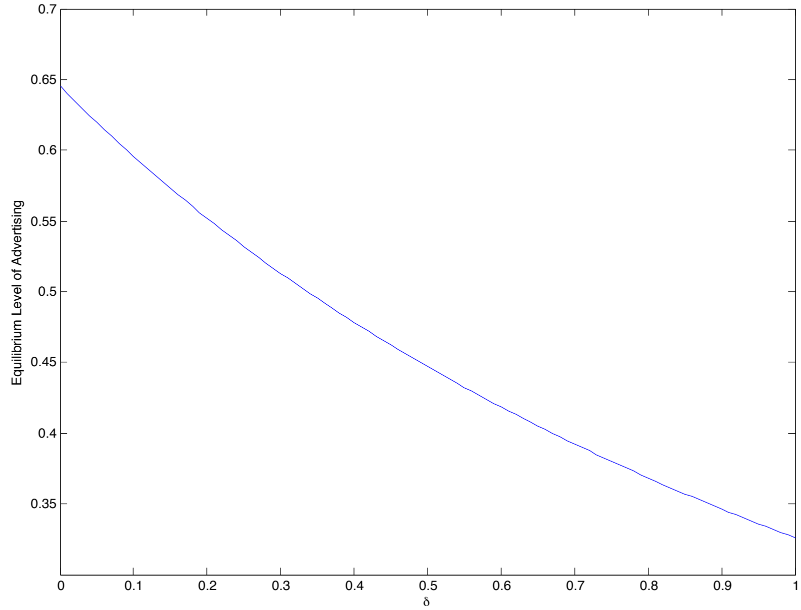

Figure 7: Magnitude of Spillovers vs. Equilibrium Prices

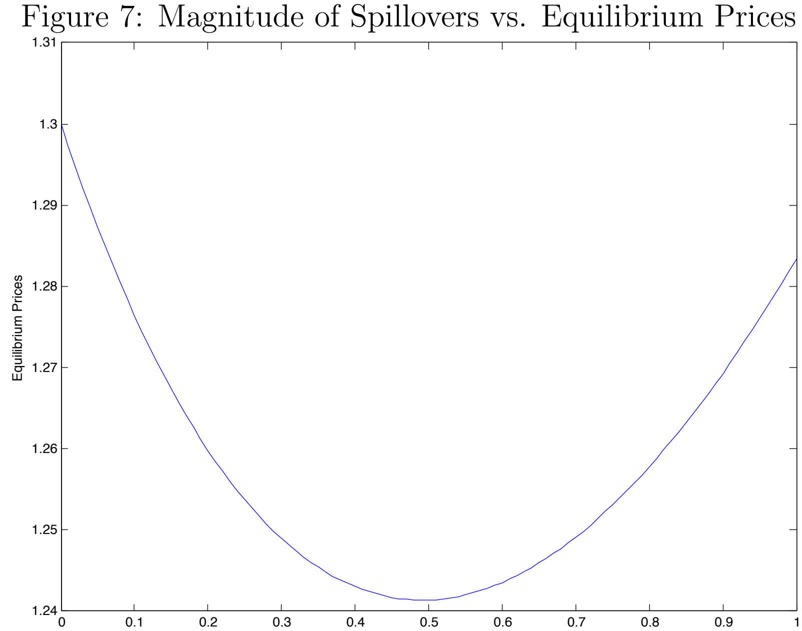

Figure 8: Magnitude of Spillovers vs. Equilibrium Profits