

Empirical Analysis of Search Advertising Strategies

Bhanu C. Vattikonda
University of California, San Diego
bvattikonda@cs.ucsd.edu

Saikat Guha
Microsoft Research, Bangalore, India
saikat@microsoft.com

Vacha Dave
Microsoft, Mountain View, California
vachad@microsoft.com

Alex C. Snoeren
University of California, San Diego
snoeren@cs.ucsd.edu

ABSTRACT

Top search ad placement is the coin of today's Internet services realm. An entire industry of search engine marketing companies have emerged to help advertisers optimize their ad campaigns to deliver high returns on investment, peddling a plethora of advertising strategies. Yet, very little is publicly known about the effectiveness of online search advertising, especially when trying to compare the various campaign strategies used by advertisers.

This paper presents the first large-scale measurement of the effectiveness—measured in terms of incremental conversion gains—of online search ads. We develop a simple metric called net acquisition benefit (NAB) that admits comparisons between the efficacy of different ad campaign strategies without access to advertisers' private financial information. We study three common campaign strategies used by advertisers on a large search ad network: cannibalization, poaching, and ad extensions. Considering data from a month in the last two years, we employ NAB to identify cases where these campaign strategies are justified. Advertisers and ad agencies can replicate our methodology to apply it to other strategies of interest.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques; H.3.5 [Online Information Services]: Web-based services

General Terms

Measurement

Keywords

Sponsored search; advertising effectiveness;

1. INTRODUCTION

Measuring the effectiveness of advertising is, in general, a hard problem. As John Wanamaker famously quipped, "half the money I spend on advertising is wasted; the trouble is I don't know which half." Online search advertising holds out the promise of addressing this longstanding challenge by focusing spend on the right consumers at the right time. The sheer scale of fine-grained, user-activity data that can be brought to bear (e.g., tracking every ad click and every user action on advertisers' site) allows advertisers to reach populations of particular interest. Moreover, search queries capture the intent of the user allowing more direct connections between the ad and the user action [26]. This tight relationship contrasts with traditional brand advertising where connections between ads and purchases are more nebulous. Hence, it is no surprise that search advertising accounts for over 40% of the \$42 billion spent in online ads in the US alone, and is growing at 18% annually [25].

Despite its ever-increasing prevalence, very little is publicly known about the effectiveness of online search advertising. Indeed, large advertisers have presented conflicting anecdotal evidence. As recently as in April 2013, a study claimed that the estimated \$51 million eBay spends on search ads is ineffective since they essentially cannibalize clicks from organic search results: in the absence of eBay's ad the user would have clicked the eBay page in the organic results [10]. Similarly, three Indian online apparel retailers found that poaching each other's users by advertising on their competitors' brand names was counterproductive; while they might succeed in getting their competitors' customers to click their ad, the customers would typically not convert. For the few customers who do end up making a purchase, the cost paid by the advertiser is too high [7]. Despite these reports, ad networks and ad agencies maintain that poaching and cannibalizing organic clicks both have a net positive return on investment [7]. Opinions are similarly mixed in the mobile space, where many advertisers complain about accidental clicks while ad networks defend the effectiveness of mobile search ads [4, 6].

One factor contributing to these seemingly contradictory reports is the current inability of anyone other than the advertisers themselves to determine if a particular advertising campaign is effective. Publishers and marketing firms deal in terms of metrics like click-through ratio (CTR), which reports the number of clicks on an ad as a fraction of the number of times the ad is shown, and cost per acquisition (CPA), which reports an advertiser's total spend divided by the number of resulting purchases. Unfortunately, neither of these metrics lends insight to the key question—namely whether a particular advertising campaign is *profitable* for the advertiser or not, which is captured by the metric profit per impression (PPI). However, determining an ad campaign's PPI fundamentally requires knowledge of an advertiser's cost and revenue structure—

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

IMC'15, October 28–30, 2015, Tokyo, Japan.

© 2015 ACM. ISBN 978-1-4503-3848-6/15/10\$15.00

DOI: <http://dx.doi.org/10.1145/2815675.2815694>.




Figure 1: Examples of search ad types.

information they are likely loathe to share, even with their marketing agencies.

Instead, we present a simple metric—net acquisition benefit (NAB)—that leverages information about how much an advertiser bids for ad placement (which, for rational advertisers, is a lower bound on the profit they expect) to estimate profit per impression (PPI). Using data about billions of clicks from a large search ad network we conduct a month-long measurement study of the effectiveness of three search ad campaign strategies—cannibalization, poaching, and ad extensions—that are widely employed by advertisers today.

We find that the advertising landscape is quite nuanced: a given campaign strategy may be effective for some advertisers but not others. For instance, advertisers with well-established brands may find little benefit in advertising on search queries mentioning their brand, while less well-known brands may need to do so to protect themselves from competitors looking to poach their customers. Similarly, ads on mobile devices may be more effective for businesses with brick-and-mortar storefronts than for businesses with purely online presence. We hope that NAB will allow not only advertisers but ad agencies and other third parties to replicate our methodology to measure and compare other advertising strategies.

2. BACKGROUND

This section provides a brief overview of search advertising and related concepts. Readers already familiar with search ads may prefer to skip ahead and refer back to this section as needed.

Search ads and organic results. When a user types a search query into a search engine like Google, Bing, or Yahoo!, the resulting page contains two types of listings. Organic search results—traditionally understood to be the search results—are determined by the search engine through proprietary heuristic algorithms; the listed websites do not pay the search engine for placement or click-throughs. Paid search results, on the other hand, which we refer to as search ads, are typically sold by the search engine through an auction mechanism where the destination website (i.e., advertiser) bids for inclusion. The top paid search results are typically displayed alongside organic search results and formatted to look similar except for, e.g., a slightly different background shade, or a ‘Sponsored’ or ‘Ad’ label.

Content. The content of an ad includes at a minimum a title, the advertiser’s domain name, and two short lines of descriptive text, typically rendered in blue, green, and black, respectively as illustrated in Figure 1. The advertiser also provides the URL the user should be directed to if the ad is clicked. Search engines increasingly support ad extensions that allow advertisers to include additional information or actions in the rendered ad. The call extension, for instance, allows an advertiser to provide a phone number; when the ad is shown on a mobile phone, the extension is rendered as a button that invokes the dialer as shown in Figure 1(b).

Targeting. Advertisers select individual keywords or phrases that must be present in the (normalized) user search query for the ad to be included in the ad auction. User search queries are normalized to fix misspellings, stem words, and resolve ambiguities using various heuristics. The advertiser may further target their ad by de-

vice type, geographic region, time of day, and user demographics. Targeting is often quite fine-grained in practice, with top advertisers managing tens of millions of keywords [10].

Campaigns and pricing. Ads are priced through an auction mechanism that is designed to encourage advertisers to bid the maximum amount they are willing to pay for a click. The amount the advertiser is charged if the ad is clicked is based on the next-lower-ranked bid in a form of generalized second-price auction [20]. The position of the ad on the results page is based on some function of the bid and the probability the ad will be clicked (e.g., based on historical click-through rates). Advertisers typically manage bids for a group of related keywords and targeting criteria through a logical construct called a campaign. Each bid is associated with a particular ad campaign, which can specify fine-grained bid modifications based upon device type and individual keywords.

Conversion. The purpose of an ad is obviously to drive revenue to the advertiser. Conversion (or equivalently, acquisition) refers to the user performing some desirable action on the advertiser’s site after clicking the ad, e.g., signing up for a newsletter, purchasing a product, or subscribing to a service. Sophisticated search-ad networks assist advertisers in monitoring the effectiveness of their campaigns by providing support for analytics. In particular, the advertiser can inform the search-ad network of a conversion event by placing JavaScript code provided by the search-ad network on the page on which the conversion happens. The JavaScript directs the browser to contact the search-ad network’s server with a user cookie which can then be used to link the conversion event to any previous user actions performed on the search-ad network. The advertiser can track campaign performance along different types of conversions by passing an opaque tag to the conversion JavaScript; the search-ad network typically does not learn the semantics of the tag or the value of the conversion and simply tallies the counts of these opaque tags for the campaign.

3. METRIC

In this section we present NAB, a simple metric we use for measuring the effectiveness of online search ads. Our primary goal is to design a metric that can be used by ad networks, ad agencies, or any entity that manages large search advertising campaigns. We discuss the challenges in choosing the right metric to measure the performance of ad campaigns. In particular, we address the fact that details of customer transactions (e.g., profit margins) may not be available to the entities running the search engine marketing campaigns for the advertisers.

3.1 Challenges

Our choice of metric for comparing the performance of different advertising strategies is driven by the following challenges.

3.1.1 Advertiser diversity

Advertisers on our partner ad network span orders of magnitude in terms of their scale. The effective cost per acquisition (CPA) for an advertiser, computed as the total amount charged divided by the total number of conversions, is a rough proxy for the monetary utility the advertiser expects to derive from advertising [21]. Figure 2 plots the normalized CPA for advertisers reporting at least 10 conversion events in our dataset vs. the total normalized money spent by the advertiser; the values are normalized to the median value along each axis. (Because each advertiser has at least 10 conversions, there can be no points above the $CPA = TotalSpend/10$ diagonal.) As one might expect, different advertisers have vastly different budgets. Both the total money spent by advertisers (x -axis domain) and effective cost-per-acquisition (y -axis range) span




Figure 2: Large spread in the normalized advertiser ad spend and the normalized price they pay per conversion.




Figure 3: Top queries account for a small share of overall spend illustrating the query diversity.

almost six orders of magnitude. Thus, our metric must take into account an advertisers’ target cost per acquisition and allow comparison between different advertising strategies of an advertiser.

3.1.2 Statistical significance

As mentioned, ad targeting can be extremely fine-grained focusing on specific keywords, device types, geographic regions, etc. Figure 3 plots the probability density function of the share of total advertising money spent on a sampling of the most popular keywords (ordered by the amount of money spent on them). To illustrate, note that the top keyword contributes to less than 0.2% of the total spend. There is a heavy tail of keywords with the top 200 most-popular together accounting for less than 7% of the total spend in the sample. Performing analysis at keyword granularity results in poor statistical significance. The statistical significance is lower still if the data is further sliced by user, device type, advertiser and other targeting parameters.

Hence, any useful metric must be defined over traffic aggregates, which we refer to as slices of search requests. The dimensions along which the data should be aggregated depends upon the strategy being evaluated. In general, we will consider slices that capture a particular advertising campaign. For instance, aggregating data

by device type may help evaluate the effectiveness of advertising on mobile devices, while aggregating data by query classification to, e.g., consider only ads placed on competitors’ brand names, may help evaluate the impact of poaching.

3.2 Net acquisition benefit (NAB)

Intuitively, the net acquisition benefit (NAB) is the conversion probability of a traffic slice adjusted by its cost. We define NAB for a traffic slice x as follows:

$$\text{NAB}(x) = \pi_x - \frac{\nu_x}{\lambda},$$

where:

- x, n : Traffic slice x consisting of n impressions
- π_x = Conversion probability, i.e., $\frac{\#conversions}{n}$
- ν_x = Average cost, i.e., $\frac{cost}{n}$
- λ = Advertiser’s target cost-per-acquisition

λ is the maximum amount an advertiser would be willing to pay for a conversion, which is well captured, for example, by their bid in an ad auction (see Section 2). Obviously, a rational advertiser would not want to pay more for an ad than they stand to make in profit on the conversion, so we argue that λ serves as a lower bound for the profit an advertiser expects to capture from a conversion. Note that if λ precisely equals the advertiser’s profit margin, then NAB is proportional to profit per impression (PPI).

$\text{NAB} = 1$ for an optimally beneficial traffic slice — where every case results in a conversion ($\pi = 1$) and there is no cost ($\nu = 0$). $\text{NAB} = 0$ for traffic slices that have no net benefit, e.g., where the traffic slice is so expensive that the advertiser is willing to forgo every conversion and save the entire cost (i.e., effective cost-per-acquisition is λ and $\nu = \lambda\pi$). For detrimental traffic slices, e.g., there are no conversions ($\pi = 0$) and the advertiser is losing money ($\nu > 0$), NAB is negative. In practice, NAB is on the order of 10^{-2} in our real-world dataset. (Intuitively, this makes sense, as CTRs are typically of the same order.)

3.3 Incremental NAB

The incremental net acquisition benefit (INAB) measures the relative improvement in NAB of one traffic slice over another, i.e., the effectiveness of one ad campaign vs. another. Intuitively, it is the change in conversion probability ($\Delta\pi$) adjusted by the change in cost ($\Delta\nu$). We define INAB for traffic slice x over slice y as follows:

$$\text{INAB}(x | y) = \frac{\text{NAB}(x) - \text{NAB}(y)}{|\text{NAB}(y)|}$$

INAB is defined only for two comparable traffic slices. That is, INAB can be computed for two traffic slices x and y belonging to the same advertiser but not across advertisers. Traffic slice x is more beneficial than y if and only if x has more net acquisition benefit than y , i.e., $\text{INAB}(x | y)$ is positive. Slice x is less beneficial than y if $\text{INAB}(x | y)$ is negative (or equivalently, $\text{INAB}(y | x)$ is positive). Both are equivalent if $\text{INAB}(x | y)$ is zero.

3.4 Discussion

NAB approximates profit per impression (PPI) when λ is equal to profit margin on conversion. But, unlike profit per impression, NAB does not require information about revenue derived and cost of the products. This allows ad networks and ad agencies to use the NAB metric to compare effectiveness of advertising campaigns. Note that both NAB and PPI, being impression based, can be sensitive to impression counts (impressions are cheap, one may argue). However, an impression represents the most basic intervention to

user experience on behalf of advertisers, and in absence of a better denominator, impression-based metrics are still considered industry standard [2, 38].

CTR and CPA. Click-through ratio (CTR) and cost per acquisition (CPA) are the two most commonly used metrics to evaluate effectiveness of ad campaigns. CTR captures the effectiveness of a campaign in attracting users but does not capture the benefit of attracting those users. It also does not capture the price that the advertiser has to pay for the clicks. While CPA captures the quality of users being attracted, it does not capture the efficiency of the campaign in attracting users overall, something that CTR captures.

We find that NAB captures the performance of a campaign better than both CTR and CPA. In particular, NAB is robust to the diversity of advertisers and sparsity of query-level data found in our dataset. It has been observed [12, 24] that even large search engines have very sparse data when aggregated by user queries. NAB allows comparisons between different advertising strategies of a given advertiser and can be applied to any slice of traffic with sufficient data.

Target cost per acquisition. One subtlety with NAB is that advertisers could have different target costs per acquisition for different ad campaigns. We discuss our methodology for inferring target cost per acquisition in next section.

4. DATA AND METHODOLOGY

This section describes the dataset we use and the methodology we follow for measuring effectiveness of various ad campaign strategies in the subsequent sections.

4.1 Dataset

Our dataset sample is several terabytes in size, consisting of billions of search and ad clicks on a large search ad network in the United States English language market. We report upon clicks capturing all actions taken by hundreds of millions of users who issue hundreds of millions of unique English-language queries over a period of four contiguous weeks within the last two years. (Analysis of a different four-week time period obtains qualitatively similar results.) The dataset contains a representative sample of clicks from desktops, tablets and phones. Our dataset sample covers many millions of dollars in advertising spend¹ by hundreds of thousands of advertisers. Our dataset does not cover specialized search verticals like image, video and map, or product listings.

For each click our dataset includes the following anonymized information: *i*) the normalized search query issued by the user and the search ad network’s internal query classification; *ii*) information about the browser including version and operating system, device form factor *iii*) the list of organic search results and paid search ads that were presented to the user; *iv*) the details of the associated ad campaign including bid amounts, keywords targeted, and ad extensions; *v*) the organic search results or paid search ads on which the user actually clicked (if any) including clicks on multiple results and ads; and, lastly, for ads clicked, *vi*) the second-price bid charged to the advertiser along with any advertiser-reported conversion event(s) for that click along with the URL of the pages for which the user conversion was reported and the (opaque) conversion types.

Along with click data, the analytics system collects user activity data on advertiser websites to track the performance of their ad campaigns. Whenever a user performs an action that the advertiser

¹We are obliged to report only the magnitude or normalized values for some sensitive quantities when doing so does not compromise the scientific value of our results.

wants to track, JavaScript embedded in the page sends information about the action—along with a user cookie allowing the data to be connected to the user’s search behavior—to the ad network. This system allows advertisers to declare which user actions constitute conversions.

4.2 Methodology

In order to compute NAB we infer target cost per conversion (λ) and conversion rate of traffic acquired organically from the data. We aggregate data over queries that identify traffic representing different advertising strategies.

Conversions. In order to identify the conversions that an advertiser obtains from a slice of traffic we have to attribute the conversion to a specific prior search. For this, we identify the user actions on the search engine prior to the conversion event on the advertiser website. We then attribute the conversion to the latest user click (regardless of whether the click was on an organic result or a search ad) that led the user to the advertiser’s website—as long as the click happened in the 24 hours prior to the conversion event.

Not all advertisers report conversion signals, or not in significant numbers. Unless otherwise mentioned, we omit advertisers for whom we have less than 30 conversion reports in our dataset sample.

Inferring target cost per acquisition. Recall from Section 2 that advertisers place a bid which is the maximum amount they are willing to pay for a click. We infer the maximum amount the advertiser is willing to pay for a conversion (λ) by dividing their total bid amount for the ads clicked by the number of conversions they received. Since the bid values are always more than the actual cost of advertising, overall NAB for any advertiser is always positive. Note that by making this choice we consider all conversions that the advertiser receives in US English market—irrespective of campaign—equivalent.

Aggregating queries. NAB must be computed over a significant aggregation of traffic. As mentioned earlier, individual search queries are too granular. We follow the search-ad network’s internal classification scheme [41] to aggregate queries into the following four classes of particular interest: navigational, local, commercial, and other, which includes informational queries.

Since the ad network’s internal query classifier uses heuristics, we verify the correctness of classification by manually investigating a representative sample. In all we manually verify 200 queries and find that in the large majority of cases (94%) our manual label matches the classifier’s; in the remaining 6% of the cases we believe the classifier misclassified the query. We compute the sensitivity and specificity measures for classification of navigational queries and find that that 77% of the time, a navigational query is classified as navigational, whereas 4% of the time, a non-navigational query is classified as navigational. Overall, the query classification—while not perfect—seems sufficiently accurate for the purposes of our study.

Traffic slices. In subsequent sections we compare the effectiveness of various campaign strategies used by advertisers by comparing the benefits across different traffic slices. Each slice of traffic is defined by the query classification, device type, the position of the advertiser in the organic results (if at all), whether the advertiser’s ad is shown or not, and whether the call button was present for the ad or organic result (in the case of mobile devices). Table 1 lists these features and describes the values they take. Table 2 labels the various combinations of these features with a name that we use to refer to that traffic slice in subsequent sections. We discuss our choice of the traffic slices in Section 8.

Query	
nav	Navigational query; user seeks specific site
com	Commercial query; user has purchase intent
all	All queries
Device	
phone	Mobile smartphones
pc	Desktops and laptops
all	All devices
Organic	
yes	Present in first page of results
no	Not in first page of results
top	Top-most organic search result
poor	Ranked 2 or worse, or not on first page
n	Ranked n
$n+$	Ranked n or worse
all	All cases whether present or not
Ad	
yes	Ad present
no	Ad not present
Ext. (set)	
ad:call	Ad has call button
org:call	Organic result had call button
ad:comp	Competitor has an ad

Table 1: Traffic features used to define traffic slices

Slice	Query	Device	Org.	Ad	Ext.
<i>Section 5: Cannibalizing Organic</i>					
org-n-noad	all	pc	n	no	
org-n-ad	all	pc	n	yes	
nav-noad	nav	pc	top	no	
nav-ad	nav	pc	top	yes	
nav-comp-noad	nav	pc	top	no	ad:comp
nav-comp-ad	nav	pc	top	yes	ad:comp
noorg-ad	all	pc	no	yes	
<i>Section 6: Poaching</i>					
poach-ad	nav	pc	poor	yes	
poach-noad	nav	pc	poor	no	
commerce-ad	com	pc	all	yes	
<i>Section 7: Ad Extensions</i>					
phone-orgcall	all	phone	yes	no	org:call
phone-org	all	phone	yes	no	
phone-ad	all	phone	all	yes	
phone-adcall	all	phone	all	yes	ad:call
phone-noorg-ad	all	phone	no	yes	
phone-noorg-adcall	all	phone	no	yes	ad:call
phone-orgcall-adcall	all	phone	all	yes	org:call, ad:call

Table 2: Traffic slices used in this paper

5. CANNIBALIZING ORGANIC TRAFFIC

We refer to the scenario where advertisers show ads for queries where they have an organic presence as cannibalization: in the absence of ads, users could have navigated to the advertiser using organic results. We consider three particular scenarios of cannibalization and use NAB to measure the benefits of each. We find that, *i*) for queries where the advertiser is the top result 56% of the advertisers derive no benefit from advertising, *ii*) the incremental benefit of advertising increases as the organic rank of advertiser decreases and *iii*) over 61% of the advertisers achieve limited benefit by advertising on navigational queries.

5.1 Improving visibility

In this section we consider the general case of searches where the user does not already have a particular destination website (advertiser) in mind. In such an instance, even though an advertiser’s




Figure 4: Most advertisers see little to no benefit in advertising for queries where they are the top result.




Figure 5: 56% advertisers achieve zero or negative INAB by advertising on queries where advertiser is the top result.

site appears in the organic search results, the advertiser may wish to increase its visibility to induce the user to visit its site as opposed to a competitor’s. We consider an advertiser to be attempting to improve their visibility if the advertiser chooses to display an ad despite already being included in the organic results likely to be considered by the user (i.e., ranked in the top ten organic results which are chosen based on their relevance to the user query).

5.1.1 Same-query visibility

The main challenge in measuring the impact of advertising is obtaining comparable search impressions. In the ideal scenario we would compare user actions in the presence and absence of an ad while everything else remains the same. Since we do not have such data, we compare the performance of ads by comparing benefits of advertising over two sets of impressions for the same query, one with ads and one without ads. Comparing user actions over impressions for the same query ensures that the most significant variable—the user query—remains the same in both sample sets.

We begin by considering the effectiveness of advertising when an advertiser is already the top organic search result. For each (advertiser, query) pair, where advertiser is the top organic result for that




Figure 6: The benefits of cannibalizing organic clicks increase with decreasing organic search ranking.

query, we identify impressions with and without ads by the advertiser. I.e., a single advertiser may be considered multiple times if they are the top search result for more than one query. Figure 4 plots the NAB of advertising (query-org-1-ad) vs. not advertising (query-org-1-noad) for all (advertiser, query) pairs where we have more than 30 conversions (as discussed in Section 4). In our dataset, we have 824 (advertiser, query) pairs covering 345 distinct advertisers. The diagonal line represents an INAB of zero, i.e., the effectiveness of advertising is equivalent to the effectiveness of not advertising for that advertiser. Advertising is more effective than not for advertisers above the diagonal. A slight majority are below the line where the inverse is true—i.e., it is not worth advertising for that query. The dashed red curves above and below the diagonal are intended to aide comparison as the graph spans four orders of magnitude on each axis: they represent a ΔNAB ($\text{NAB}(\text{ad}) - \text{NAB}(\text{noad})$) of $+0.01$ and -0.01 respectively.

As an alternative representation of the same data, Figure 5 plots the CDF of the corresponding INAB metric. Recall that INAB is normalized to the NAB of the base strategy, so a value of 1.0 represents a $2\times$ improvement. For 56% of the (advertiser, query) pairs the corresponding advertiser gets zero or negative incremental benefit by advertising on the query (i.e., fall below the diagonal in Figure 4). Again, this means that these advertisers should reconsider advertising on those queries.

5.1.2 Over all queries

The constraint of comparing the benefits of advertising by holding the user query a constant limits our analysis to a few advertisers who have significant number of conversions for the same query. Even for those advertisers, these few queries contribute to a small proportion of their entire ad spend. Hence, there is reason to believe the same-query results may not be representative. By design, we can apply our metric to measure effectiveness of ads on an arbitrary slice of traffic. Here we expand our analysis to compare traffic slices where the advertiser has a particular rank in organic results, irrespective of user query.

Figure 6 considers only non-navigational queries (i.e., those where the user likely does not have a particular destination website in mind) and plots NAB of advertising (org-n-ad) vs. not advertising (org-n-noad) aggregated across all queries for each advertiser (i.e., each advertiser is a single data point) where we have sufficient data as described in Section 4. We separate campaigns based upon their rank in organic results. Figure 6(a) considers the aggregate performance of all non-navigational queries for which an advertiser


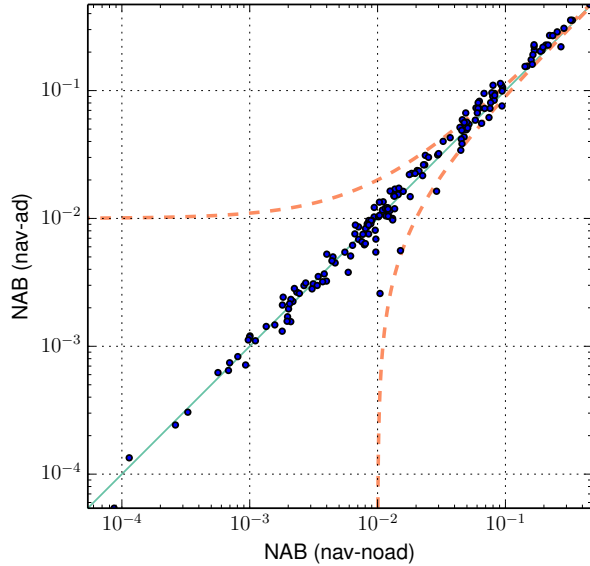


Figure 7: As rank increases in organic results benefits of cannibalistic advertising turn positive, albeit only slightly.


anyway obtains the top organic result (c.f. Figure 4 which shows a distinct point for each unique top-result-producing query an advertiser may obtain). Figures 6(b) and 6(c) show the same comparison for campaigns where the advertiser appears third or fifth in organic search results, respectively.

Figure 7 plots the CDFs of the corresponding INAB metrics for each of the three classes of campaign considered in Figure 6. In contrast to the same-query results above, most of the advertisers gain by advertising for non-navigational queries despite being present in the organic results. This is likely because non-navigational queries tend to be competed for more aggressively and users are flexible with choosing any business that meets their needs. Also, while the relative benefit (i.e., INAB) of advertising increase as the organic rank drops, the absolute benefits (i.e., NAB) of advertising are very low.

Interestingly, for 32% of the top-ranked websites, there appear to be significant benefits to advertising ($\text{INAB} > 50\%$). Manual investigations find that these advertisers belong to two categories. One set of advertisers have less-well-recognized brands. If a more well recognized brand advertises for the query, and is thus placed above the top-ranked organic result, the latter lose out. The other category of advertisers who gain are those who share their brand with other advertisers. For example, car manufacturers lose potential converting users to competing ads placed by car dealers trying to attract the same users. For 19% of the top-ranked websites, advertising is a drain on their ad spend ($\text{INAB} < 0$) since they end up competing on hotly contested queries resulting in higher cost per acquisition than their overall average.



(a) Advertising for navigational queries.



(b) INAB of navigational queries.

Figure 8: The incremental benefit of advertising for navigational queries is very small for most advertisers.

5.2 Navigational queries

A subset of queries (explicitly excluded above) for which an advertiser is the top organic result is known as navigational queries—i.e., the user is using the search engine as a way to navigate to the advertiser’s site. We consider a given query navigational for an advertiser if: *i*) the search query is classified by the search ad network as navigational, i.e., the user query includes branded terms or keywords (e.g., Amazon or Facebook) that suggests the user is seeking a specific website, and *ii*) the advertiser is the top-most organic result for that query. Here, we consider whether it is effective for a website to advertise for queries where the user is explicitly seeking to navigate to the advertiser website.

Figure 8(a) plots the NAB of advertising (nav-ad) versus not advertising (nav-noad) for all advertisers where we have at least 30 conversions attributed to clicks following queries matching the above criteria. Figure 8(b) plots the same data as a CDF of the INAB of advertising over not advertising on these queries. As is evident from points clustering along the diagonal in the scatter plot, the majority of advertisers (61%) receive very little incremental benefit (less than 10%) from advertising on navigational queries where they are the top organic result. Moreover, for 32% of ad-




Figure 9: INAB has no correlation with CTR and most advertisers receive a high click through rate on navigational queries.

vertisers, advertising on navigational queries for which they are the top organic result is a net loss when compared to not advertising (i.e., INAB is < 0). In 9% of the cases, however, such advertising bears significant fruit (INAB is $> 25\%$).

Our finding that 61% of advertisers receive limited benefit (INAB approximately zero) from advertising on navigational queries vs. not advertising squarely contradicts reports from other ad networks that suggest 89% of ad clicks are incremental, and would be lost without advertising [13]. We reconcile these results by observing that the previous study [13] does not consider conversions and focuses solely on clicks. Indeed, the study’s authors explicitly state that advertisers should consider conversions since relying on clicks alone may be misleading. We show below that advertising on navigational queries does, in fact, inflate click counts without increasing conversions for the vast majority of advertisers. Another reason for the divergent result is that the previous study considers campaigns that are paused due to budget shortage. This would bias the choice of campaigns towards very small advertisers who would not have strong presence for navigational queries.

For advertisers that value conversions, however, our data discredits conventional wisdom that promotes advertising on navigational queries for which the advertiser is the top organic result. Finally, while our general finding is consistent with experimental evidence from large advertisers including eBay [10], we nevertheless encourage advertisers to conduct their own experiments and track conversions to determine if they belong to the small (9%) set of advertisers for whom advertising on navigational queries brings significant benefits. These advertisers, as we discussed in Section 5.1.2, either have a weaker brand or are competing against other advertisers who can legitimately advertise on their brand.

5.3 Click count inflation

Many advertisers appear to optimize for clicks rather than conversions. While we can only speculate as to why they choose to do so, anecdotal evidence ranges from naïveté, e.g., unawareness of metrics other than the click-through rate, to financial, e.g., ad agencies that collect commission per-click and advertisers that specify CTRs and minimum click counts (rather than conversions) in contracts with such agencies [3]. One of the easiest ways for advertisers to inflate ad click counts is to cannibalize the advertiser’s own navigational queries.

Figure 9 plots achieved click-through rate of campaigns that advertise on an advertiser’s own navigational queries as a function of their INAB. While the INAB of such a strategy is close to zero for the vast majority of campaigns as one would expect, the click-through rate for these campaigns is extremely high as compared to typical search ad CTRs of around 1-2% [5]. Advertisers should run

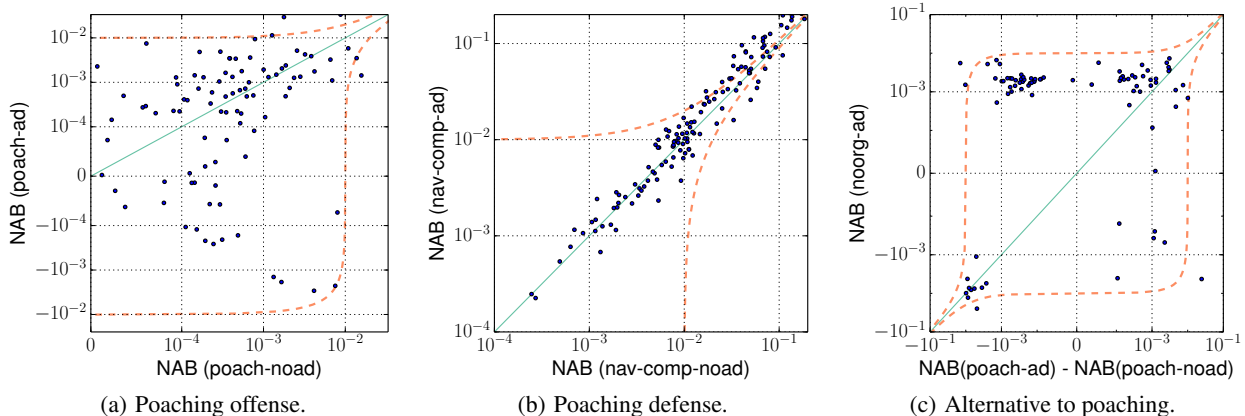


Figure 10: Poaching often leads to wasteful spend and defending against it is not worthwhile either. A better alternative is to try to gain users on commercial queries where there is no organic presence.

carefully calibrated A/B tests to determine the incremental benefit of advertising for queries where the user is already looking to navigate to the advertiser.

6. POACHING COMPETITOR TRAFFIC

Poaching refers to advertising strategies that involve bidding on navigational queries specifically seeking a competitor’s website (e.g., queries with competitor’s trademarked terms or brand names). Initially, Google’s trademark usage policy prevented advertisers from bidding on trademarked keywords they did not own or from using them in their ad content; the restriction of bidding on competitors’ trademarked keywords was relaxed in 2004—even if the trademark owner explicitly objected—and restrictions on using them in ad content were relaxed in 2009 [1, 2]. The relaxed policy effectively increased the cost per click of trademarked keywords by allowing competition from third parties, thereby increasing trademark owners’ costs by compelling them to bid defensively to protect their trademark.

We find that poaching may help smaller advertisers get users who would not navigate to the advertiser in the absence of ads. But, the costs of attempting to gain such users may be too high for some advertisers when compared to the cost they pay for an average conversion. Moreover, given the varied effectiveness of poaching, defending against poaching has uneven results and can, in fact, occasionally lead to negative results due to the high cost of defense.

6.1 Offense

Here, we analyze the benefit of trying to poach users seeking to navigate to a competitor. Figure 10(a) compares the NAB for poaching ads (poach-ad, i.e., ads on navigational queries where the advertiser is not the top-most organic search result) against the NAB of not advertising on the same queries (poach-noad, i.e., navigational queries where the advertiser has an organic presence but is *not* the top-most organic result, and does not display an ad). Each data point represents an advertiser where we have sufficient conversion data. About half the advertisers fall below the diagonal showing that poaching is actually detrimental. For the half that do see relative gains over not advertising, the absolute benefits are negligible (note the absolute value of NAB is generally below 0.01).

Figure 11 plots the CDF of the INAB for these advertisers. For 50% of advertisers, poaching is of negative value—likely due to few conversions and high costs associated with such ads. There




Figure 11: For most advertisers, poaching competitor’s navigational queries yields little benefit, while some see extreme gains.

are however, a minority that see significant gains relative to their performance when ads are not shown. This is due to the very poor rate at which they obtain traffic when they do not show ads; said another way, while the relative gains are substantial, in absolute terms they still do not receive many conversions.

6.2 Defense

Regardless of how effective poaching is for the advertiser, competitors may still be harmed because they value lost conversions more dearly. Here, we consider whether it is useful for an advertiser to defend against potential poaching by out-bidding competitors for ad space despite being the top organic result for a navigational query. This data represents a specific scenario of advertising for navigational queries (discussed in Section 5.2) where a competitor ad is also present. We consider an advertiser a potential victim of poaching for a given query if: *i*) the search query is classified as navigational, i.e., the user query includes branded terms (e.g., Amazon or Macys) that suggests the user is seeking a specific website, *ii*) the advertiser is the top-most organic result for that query, and *iii*) another advertiser advertises on the query.

Figure 10(b) compares the NAB of defensive ads (nav-comp-ad, i.e., ads on navigational queries where the advertiser is the top-most organic search result) against the NAB of not advertising on the same queries (nav-comp-noad). Figure 12 plots the CDF of the INAB for these advertisers. The results are mixed: 20% of advertisers realize significant (INAB > 25%) benefits, while a 32% are negatively impacted by attempting to defend these queries.

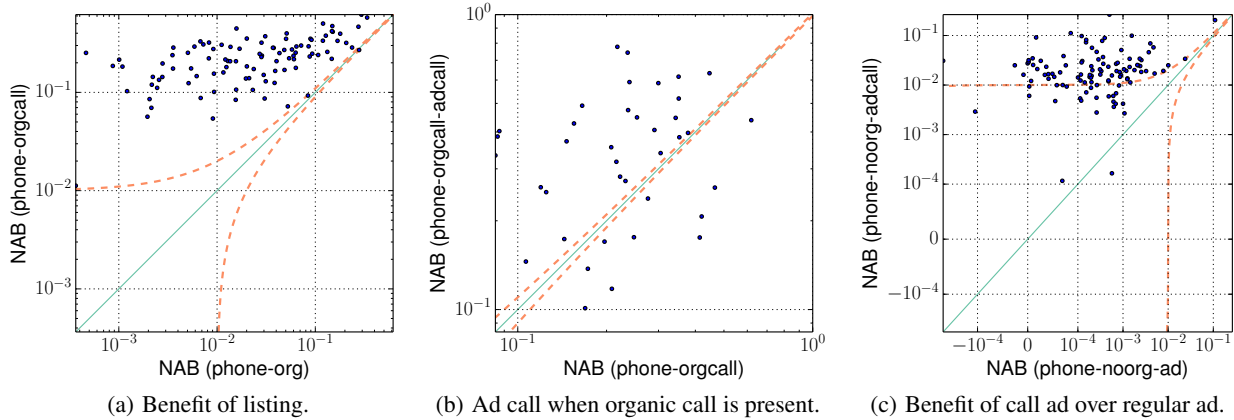


Figure 13: Listing business details benefits advertisers and in the presence of organic call options, ad call extensions yield mixed results. But, showing an ad with extension is more beneficial than a regular ad.




Figure 12: The benefits of defending against poaching vary dramatically across advertisers.

6.3 Spending smarter

As an alternative to poaching a competitor’s customers (i.e., users who have issued a navigational search query for a competing website), an advertiser might try instead to recruit customers who are likely to convert somewhere, but have not yet decided on a particular vendor. Here we consider an advertiser deciding between spending money on poaching ads vs. spending that money to compete on commercial queries (i.e., those that are likely to lead to conversions) where they may not be prominently ranked in the organic results, Figure 10(c) plots the additional benefit of advertising on commercial queries where the advertiser does not have an organic presence i.e. $NAB(\text{noorg-ad})$ vs. the benefit of poaching ads over not poaching ($NAB(\text{poach-ad}) - NAB(\text{poach-noad})$). For 78% of the advertisers, advertising on commercial queries where the advertiser has no organic presence dominates poaching by a wide margin (i.e., they are significantly above the diagonal).

7. DISPLAYING AD EXTENSIONS

Call extensions allow advertisers to explicitly add a call button to their ads when rendered on mobile phones. Call buttons are also shown for business listings in the organic search results. These features are provided by our partner ad network for no additional cost to the advertisers. We have limited data for different advertising strategies on mobile devices. However, our preliminary results show that creating a business listing, which would allow the ad network to show a call button in organic results, is beneficial to the




Figure 14: Listing a business in local results brings clear benefit to advertisers.

advertisers. The effect of adding a call button to an ad when an organic result with call button is already present is mixed.

7.1 Organic business listings

Businesses can create a (free) listing that includes their location, phone number, store hours, parking information, and payment methods accepted through the network’s “Places for Business” portal or third-parties like Yelp. The ad network uses this structured information to enhance the presentation of organic results including showing the call button, map directions, and so on.

We first look at the effectiveness of the call button for organic business listings vs. plain search results. It is challenging, however, to define conversion rates in these scenarios. We obviously have no way of knowing from the search logs how many users that call the business end up converting in a way that is equivalent to the conversions in previous sections; here, we consider the simple act of a user contacting the business by clicking the call button as a conversion event for the purposes of computing the NAB. For organic listings without call buttons, however, no similar data is available. Instead, we substitute the conversion rate the advertisers obtain when they advertise on local queries. While conversions are potentially an undercount when compared to calls, we hope that using the conversion rate for ads as opposed to organic clicks restores some of the balance, offering a reasonable baseline of user engagement.

Figure 13(a) plots the NAB of having an organic call button (phone-orgcall) vs NAB of plain search results (phone-org) for searches performed on mobile phones. Note that since both NABs




Figure 15: When organic call is already listed, advertising with a call ad has mixed results.




Figure 16: When advertiser does not have organic presence, advertising with a call button is more beneficial.

are for organic results—for which businesses do not pay—the cost term in the NAB computation is zero, and NAB reduces to conversions per impression. Figure 14 plots the CDF of the INAB of call button over plain results. As evident from the figures, median conversions per impression increases by a factor of 10 when a organic call result is present.

7.2 Ads with call extensions

When an advertiser chooses to place a call button in an ad, it is possible the ad results in a conversion off-line; i.e., rather than navigating to the website and converting, the user may instead call the advertiser and “convert” by making a purchase or similar analogous activity without further web interaction. Hence, when considering conversions for ads with call extensions, we define a conversion event to be either a call or an advertiser-reported conversion. For the very few impressions that result in both, we report only one conversion.

Figure 13(b) compares the NAB of mobile search ads with call extensions (phone-orgcall-adcall) vs. the NAB of organic business listings with call buttons (phone-orgcall) for the very few businesses in our dataset that both advertise call extensions as well as created an organic listing, as well as reported a sufficient number of conversion events. Figure 15 plots the CDF of the corresponding INAB. We find that even if organic search results contain a call button extension, ads employing the same are effective for 74% of advertisers, but ineffective for the remaining 26%. The orders-of-magnitude larger NAB (and INAB) values in this section can be attributed to the relative paucity of data for results using the call extension.

For businesses that do not have a rank high enough to be listed in the organic results, Figure 13(c) compares the NAB of mobile search ads with call extensions (phone-noorg-adcall) vs. the NAB

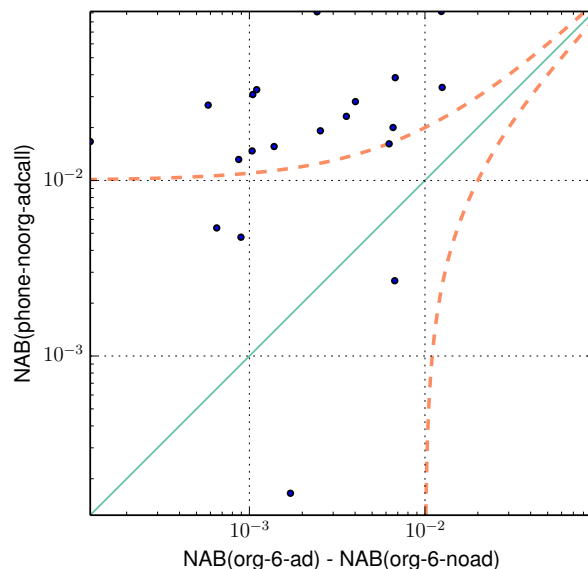


Figure 17: Call ads on mobile are better than regular desktop ads in the absence of organic presence.

of ads without the call extension (phone-noorg-ad). Figure 16 plots the CDF of the corresponding INAB. We find that ads with a call extension are universally more effective.

Preliminary data suggests that there are mixed benefits to using call extensions for businesses where organic listings include call button. For businesses that do not have a high-enough organic rank for their listing to appear in the first page of ads, however, there is a consistent boost from call extensions in search ads. That being said, since call extensions and business listings are both recent features and very few advertisers have opted-in to both, our results are preliminary and we encourage advertisers to conduct their own experiments and compute their respective INABs to assess the effectiveness of the call extension in their specific case.

7.3 Spending smarter

In closing, we compare the marginal benefit of call extensions in mobile ads to the benefit of traditional advertising on computers. As we discussed previously, the benefit of the latter depends tremendously on the position of the advertiser in the organic search results. As an optimistic estimate, we focus on an advertiser who appears sixth in the organic search results. Specifically, Figure 17 plots the marginal benefit of running mobile ads with call extension over organic business listings with call button (phone-noorg-adcall) vs. advertising on desktops and laptops. For 75% of the advertisers in our dataset, it is (modestly) more beneficial to focus on mobile call ads vs. desktop ads.

While we have limited data in our dataset, our initial assessment indicates greater benefits can be obtained by adding call buttons to ads on mobile compared to traditional advertising on desktop when organic presence is poor.

8. DISCUSSION

In this section, we discuss our choice of using aggregated traffic slices and the underlying systemic bias that exists in the way the traffic slices are chosen for comparison.

8.1 Choosing traffic slices

The techniques that we use in this paper can be applied to estimate the profitability achieved by an advertiser over an arbitrary slice of traffic that is representative of a particular advertising strategy. To measure the effectiveness of advertising, we identify similar search impressions with and without ads. For this, an ideal comparison would require impressions where the presence of an ad is the only distinguishing attribute. However, even the largest search engines have only sparse data when aggregated at the query level [24]. In order to have reasonable samples of impressions, we aggregate impressions by the category of user query. Note that the techniques we discuss in this paper can however be applied to traffic slices aggregated on other attributes like geography or user characteristics.

8.2 Systemic biases

A source of bias in our analyses is that queries for which ads are shown are often more representative of the advertiser than queries where ads are not shown for comparable traffic slices. This is because whenever ads are shown, both the ad network and advertiser find the query relevant. But, when ads are not shown, the ad network finds query relevant to advertiser but advertiser does not. So, whenever results returned by the ad network are poor, users are less likely to choose the advertiser from organic results hence lowering NAB for the traffic slice without ads.

An exception to this rule is when the advertiser tries to poach users looking for a competitor (Section 6.1). This is because in case of poaching, while the impressions with advertiser ad are likely to be more attractive to the advertiser, they also cost more.

Finally, in Section 7 queries for which the ad network presents a result with a call button are often more relevant to a mobile user than queries where a web result is delivered; hence, the NAB for organic search results with call buttons may be overstated.

9. RELATED WORK

Measuring ad effectiveness. There have been a number of studies measuring the effectiveness of display advertising ad campaigns [27, 29]. These studies measure the incremental impact of advertising by creating a treatment and a control group. We take a similar approach by identifying comparable traffic slices and then measuring the incremental impact of advertising.

Reiley, Lewis and Rao have also highlighted the challenges in conducting such studies [30, 31] and how the benefits of advertising can be overestimated. They highlight this by studying the impact of display ads on user activity. They show that presence of display ads is not necessarily correlated with subsequent user activity. A key reason is that display ads are expected to perform by their mere presence in the absence of explicit expression of intent by the user. In case of web search, however, advertiser pays the search engine only when there is an ad click and user intent is captured through the query made by the user. So, a subsequent conversion can be attributed to a click performed by the user on search engine results — sponsored or organic.

Absolute metrics. Advertising effectiveness has traditionally been measured in absolute terms, e.g., click through rates [14, 22, 37, 40] or conversion rates [11, 36]. Other work focuses on prevalence [9] or effectiveness of different advertising strategies like mobile advertising [18, 39] or the interaction between ads and organic search results [10, 44]. Since different metrics focus on different dimensions of the campaign, comparing them is not straightforward. NAB is also an absolute metric, however it explicitly addresses the multi-dimensional comparison problem by converting other dimensions (i.e., cost) into the primary dimension (conversion probabil-

ity) through a conversion factor (λ) that explicitly tracks the advertiser’s relative utility across dimensions.

Incremental metrics. The Incremental Ad Clicks metric [13] and the related one for mobiles [32] take a novel approach by proposing a way to compare two traffic slices. We were inspired by IAC but ultimately identified the severe limitations in the general approach relating to: i) ignoring conversions, ii) being ill-conditioned to noise; and iii) not scaling to the orders of magnitude difference in advertiser scale. NAB and INAB are designed to address these challenges. We were also inspired by work in cost-effectiveness analysis in the health sector [43].

Other use cases. Our work complements the large body of work in the space of ad optimizations. Researchers have described a methodology for A/B testing in context of controlled experiments online for fairer comparisons [19]. INAB complements this work by providing a metric that enables direct comparison of the relative effectiveness of the two tests. [8] develops strategies that advertisers may use to modify bid amounts. Incorporating NAB in the bidding strategy may allow for joint optimizations in cost and conversions. A large body of work has looked at identifying fraud in online ads [15, 16, 17, 23, 28, 33, 34, 35, 42]. NAB of various suspicious traffic slices may be used as an additional signal in these techniques.

10. FUTURE WORK

In this section we discuss areas where the NAB metric can be refined and outline future work.

Multiple types of conversions. In this paper we treat all the conversion events of an advertiser equivalently. If an advertiser has multiple types of conversions she is interested in, she would need to compute the NAB for each type of conversion, and those NABs would not be mutually comparable. One approach would be to allow the advertiser to specify the relative value of each conversion and extend NAB to allow different types of conversions to be traded off against one another to arrive at the net NAB across all conversion types.

Probabilistic NAB and brand advertising. As defined NAB can be applied most readily to search ads since every impression, click and conversion can be tracked precisely and with certainty. In traditional brand advertising, simply estimating impression counts is a hard problem (e.g., TV viewers going to the restroom during commercial breaks), let alone tracking conversions. It may, however, be possible to arrive at probabilistic estimates of impressions, and conversions. One approach could be to extend NAB to incorporate probabilistic inputs and output the expected NAB with confidence intervals.

What-if analysis. In our design of NAB we consider only point-in-time values and do not leverage the longitudinal perspective our historical data provides. While one could compute NAB at different points in the past to extrapolate into the future, a better approach may be to directly incorporate time-series data as input and output a time-parameterized NAB function.

11. CONCLUSION

In this paper we develop the net acquisition benefit (NAB) metric to approximate profit per impression in order to measure the effectiveness of an ad campaign, and the incremental net acquisition benefit (INAB) to measure the marginal benefit of one ad strategy over another. Using these metrics and extensive search and ad click data from a major search-ad provider, we find that cannibalizing organic traffic and poaching a competitor’s traffic are frequently ineffective while call extensions on mobile phones show promise.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. This work was funded in part by the National Science Foundation through grant CNS-1237264.

12. REFERENCES

- [1] Google Settles With Rosetta Stone Over Trademark Infringement Within Search And Ads Lawsuit. <http://techcrunch.com/2012/10/31/google-settles-with-rosetta-stone-over-trademark-infringement-within-search-and-ads-lawsuit/>.
- [2] Google:Update to U.S. ad text trademark policy. <http://adwords.blogspot.com/2009/05/update-to-us-ad-text-trademark-policy.html>.
- [3] Search Marketing Agency Price Models. <http://searchengineland.com/search-marketing-agency-pricing-models-11039>.
- [4] Smartphone user study shows mobile movement under way. <http://googlemobileads.blogspot.com/2011/04/smartphone-user-study-shows-mobile.html>.
- [5] What average Ads click through rate (CTR) is good for text ad? <https://www.en.adwords-community.com/t5/Tracking-Reporting/What-average-Ads-click-through-rate-CTR-is-good-for-text-ad/t5-p/150054>.
- [6] Will ‘Useless’ Clicks Really Cost Mobile Advertisers \$6 Billion+ by 2016? <http://searchenginewatch.com/article/2207164/Will-Useless-Clicks-Really-Cost-Mobile-Advertisers-6-Billion-by-2016>.
- [7] Is Google Gobbling Up the Indian Internet Space? <http://forbesindia.com/article/real-issue/is-google-gobbling-up-the-indian-internet-space/35641/1>, July 2013.
- [8] M. Babaioff, N. Nisan, and R. Paes Leme. Price Competition in Online Combinatorial Markets. In *Proceedings of the 23rd International Conference on World Wide Web*, 2014.
- [9] P. Barford, I. Canadi, D. Krushevskaja, Q. Ma, and S. Muthukrishnan. Adscape: Harvesting and Analyzing Online Display Ads. In *Proceedings of the 23rd International Conference on World Wide Web*, 2014.
- [10] T. Blake, C. Nosko, and S. Tadelis. Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment. *NBER Working Paper*, pages 1–26, 2013.
- [11] L. Blumrosen, J. Hartline, and S. Nong. Position auctions and non-uniform conversion rates. In *ACM EC Workshop on Advertisement Auctions*, 2008.
- [12] A. Z. Broder, P. Ciccolo, M. Fontoura, E. Gabrilovich, V. Josifovski, and L. Riedel. Search Advertising Using Web Relevance Feedback. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM ’08*, pages 1013–1022, New York, NY, USA, 2008. ACM.
- [13] D. Chan, Y. Yuan, J. Koehler, and D. Kumar. Incremental Clicks Impact Of Search Advertising. Technical report, Google, Inc., 2011.
- [14] Y. Chen and T. W. Yan. Position-normalized Click Prediction in Search Advertising. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2012.
- [15] N. Daswani and M. Stoppelman. The Anatomy of Clickbot.A. In *Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets*, 2007.
- [16] V. Dave, S. Guha, and Y. Zhang. Measuring and Fingerprinting Click-Spam in Ad Networks. *SIGCOMM Comput. Commun. Rev.*, 42(4):175–186, August 2012.
- [17] V. Dave, S. Guha, and Y. Zhang. ViceROI: Catching Click-spam in Search Ad Networks. In *Proceedings of the ACM SIGSAC Conference on Computer & Communications Security*, 2013.
- [18] M. de Sa, V. Navalpakkam, and E. F. Churchill. Mobile Advertising: Evaluating the Effects of Animation, User and Content Relevance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2013.
- [19] A. Deng, T. Li, and Y. Guo. Statistical Inference in Two-stage Online Controlled Experiments with Treatment Selection and Validation. In *Proceedings of the 23rd International Conference on World Wide Web*, 2014.
- [20] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.
- [21] Google. AdWords Help – Measuring return on investment (ROI). <https://support.google.com/adwords/answer/1722066?hl=en>.
- [22] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-Scale Bayesian Click-Through Rate Prediction for Sponsored Search Advertising in Microsoft’s Bing Search Engine. In *Proceedings of the 27th International Conference on Machine Learning ICML 2010, Invited Applications Track (unreviewed, to appear)*, June 2010. Invited Applications Track.
- [23] H. Haddadi. Fighting Online Click-Fraud Using Bluff Ads. *SIGCOMM Comput. Commun. Rev.*, 40(2):21–25, Apr. 2010.
- [24] D. Hillard, S. Schroedl, E. Manavoglu, H. Raghavan, and C. Leggetter. Improving Ad Relevance in Sponsored Search. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM ’10*, pages 361–370, New York, NY, USA, 2010. ACM.
- [25] IAB. Internet advertising revenue report. http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2013.pdf.
- [26] B. J. Jansen and S. Schuster. BIDDING ON THE BUYING FUNNEL FOR SPONSORED SEARCH AND KEYWORD ADVERTISING. *Journal of Electronic Commerce Research*, 12(1):1–18, 2011.
- [27] G. A. Johnson, R. A. Lewis, and D. H. Reiley. Location, Location, Location: Repetition and Proximity Increase Advertising Effectiveness. Available at SSRN 2268215, 2014.
- [28] A. Juels, S. Stamm, and M. Jakobsson. Combating Click Fraud via Premium Clicks. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*, 2007.
- [29] R. Lewis and D. Reiley. Online Ads and Offline Sales: Measuring the Effect of Retail Advertising via a Controlled Experiment on Yahoo! *Quantitative Marketing and Economics*, 12(3):235–266, 2014.
- [30] R. A. Lewis and J. M. Rao. On the Near Impossibility of Measuring the Returns to Advertising. 2013.
- [31] R. A. Lewis, J. M. Rao, and D. 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, WWW ’11*, pages 157–166, New York, NY, USA, 2011. ACM.
- [32] S. Lysen. Incremental Clicks Impact of Mobile Search Advertising. Technical report, Google, Inc., 2013.

- [33] A. Metwally, D. Agrawal, and A. El Abbadi. DETECTIVES: DETECTing Coalition hiT Inflation attacks in adVERTISING nEtworks Streams. In *Proceedings of the 16th International Conference on World Wide Web, WWW '07*, pages 241–250, New York, NY, USA, 2007. ACM.
- [34] A. Metwally, F. Emekçi, D. Agrawal, and A. El Abbadi. SLEUTH: Single-publisher attack dETECTION Using correlaTION Hunting. *PVLDB*, 1(2):1217–1228, Aug. 2008.
- [35] B. Miller, P. Pearce, C. Grier, C. Kreibich, and V. Paxson. What's Clicking What? Techniques and Innovations of Today's Clickbots. DIMVA, pages 164–183, 2011.
- [36] B. Rey and A. Kannan. Conversion Rate Based Bid Adjustment for Sponsored Search. In *Proceedings of the 19th International Conference on World Wide Web*, 2010.
- [37] M. Richardson, E. Dominowska, and R. Ragno. Predicting Clicks: Estimating the Click-through Rate for New Ads. In *Proceedings of the 16th International Conference on World Wide Web*, 2007.
- [38] D. Sculley, R. G. Malkin, S. Basu, and R. J. Bayardo. Predicting Bounce Rates in Sponsored Search Advertisements. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '09*, pages 1325–1334, New York, NY, USA, 2009. ACM.
- [39] V. Shankar and S. Balasubramanian. Mobile Marketing: A Synthesis and Prognosis. *Journal of Interactive Marketing*, 23(2):118 – 129, 2009. Anniversary Issue.
- [40] E. Sodomka, S. Lahaie, and D. Hillard. A Predictive Model for Advertiser Value-Per-Click in Sponsored Search. In *Proceedings of the 22nd international conference on World Wide Web*, 2013.
- [41] Y. Song, H. Ma, H. Wang, and K. Wang. Exploring and Exploiting User Search Behavior on Mobile and Tablet Devices to Improve Search Relevance. In *Proceedings of the 22Nd International Conference on World Wide Web*, 2013.
- [42] K. Springborn and P. Barford. Impression Fraud in Online Advertising via Pay-per-view Networks. In *Proceedings of the 22Nd USENIX Conference on Security*, 2013.
- [43] A. A. Stinnett and J. Mullahy. Net health benefits a new framework for the analysis of uncertainty in cost-effectiveness analysis. *Medical Decision Making*, 18(2):S68–S80, 1998.
- [44] S. Yang and A. Ghose. Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence? *Marketing Science*, 29(4):602–623, 2010.