Search Engine Revenue from Navigational and Brand Advertising

Jeffrey Gleason[1], **Alice Koeninger**[1], **Desheng Hu**[3], **Jessica Teurn**[1], **Yakov Bart**[1], **Samsun Knight**[1], **Ronald E. Robertson**[2], **Christo Wilson**[1]

[1] Northeastern University
[2] Stanford University
[3] University of Zurich

gleason.je@northeastern.edu, koeninger.a@northeastern.edu, desheng@ifi.uzh.ch, teurn.j@northeastern.edu, rer@acm.org, cbw@ccs.neu.edu

Abstract

Keyword advertising on general web search engines is a multi-billion dollar business. Keyword advertising turns contentious, however, when businesses target ads against their competitors' brand names—a practice known as "competitive poaching." To stave off poaching, companies defensively bid on ads for their own brand names. Google, in particular, has faced lawsuits and regulatory scrutiny since it altered its policies in 2004 to allow poaching.

In this study, we investigate the sources of advertising revenue earned by Google, Bing, and DuckDuckGo by examining ad impressions, clicks, and revenue on *navigational* and *brand* searches. Using logs of searches performed by a representative panel of US residents, we estimate that ads on these searches account for 28–36% of Google's search revenue, while Bing earns even more. We also find that the effectiveness of these ads for advertisers varies. We conclude by discussing the implications of our findings for advertisers and regulators.

1 Introduction

Keyword advertising on general web search engines is a multi-billion dollar business. Alphabet reported earning $149 billion in 2021 (the majority of their revenue) from ads on Google Search (Alphabet Inc. 2022). Their largest competitor in the US, Microsoft, earned $12 billion in 2021 from ads on Bing (Microsoft Corp. 2022). Smaller search engines like DuckDuckGo also rely on ads to fund their businesses.

One contentious use of search ads occurs when businesses target ads against their competitors' brand names—a practice known as "competitive poaching" (Sayedi, Jerath, and Srinivasan 2012) in the academic literature and "conquesting" among advertisers (Stern 2017). Google altered their policies to allow advertisers to bid on trademarked keywords in 2004, and subsequently won several lawsuits that legitimized this business practice (Goldman 2008, 2010; Reuters). Bing also allows advertisers to bid on trademarked keywords, and they provide ads for other search engines, like DuckDuckGo. To stave off the threat of poaching, businesses defensively bid on ads for their own brand names, resulting in situations—like the one illustrated in Figure 1—

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

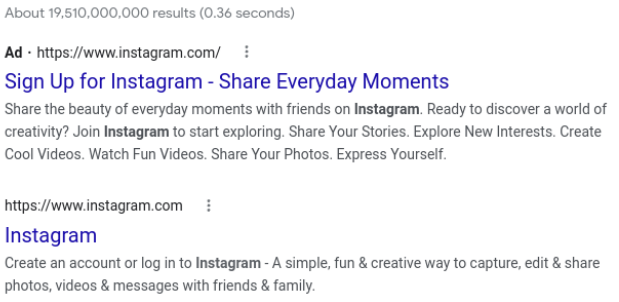


Figure 1: After searching for "instagram" on Google Search, the top two ranks are paid and organic links to Instagram.

where businesses pay to show an ad for themselves directly above the organic link to their own website.

Although competitive poaching has become normalized, businesses and regulators remain concerned about it. The US House Subcommittee on Antitrust, Commercial, and Administrative Law and the UK Competition & Markets Authority (CMA) both investigated Google Search and observed that businesses depend on it for referral traffic because it commands such a large share of the market for general web search (Subcommittee on Antitrust, Commercial and Administrative Law of the Committee on the Judiciary 2020; Competition & Markets Authority 2020). In the US House report, one business owner stated that these market conditions "force [Google's] advertising customers to pay for the ability to reach consumers who are searching specifically for the customer's brand." In 2023, the Delhi High Court went a step further and ruled that Google had to remove ads targeting trademarked keywords (Singh 2023).

While regulators frame Google's allowance of competitive poaching and defensive advertising as a form of rent seeking, these assertions are based on complaints from a small number of businesses. The reports from the US House Subcommittee and the CMA present scant empirical evidence about the scope of the issue, either in terms of advertising volume or revenue for search engine operators. The only publicly available statistic, which is from 2004 litigation, attributed 7% of Google's revenue to trademarked keywords (Golden and Horton 2021). Furthermore, Google's

choice to combine URL-based navigation and web search into a single bar in Chrome in 2008—a design emulated by all modern web browsers—may habituate the practice of typing business names, as opposed to URLs, into the browser's address bar, thus creating more opportunities for search engines to serve ads against brand names.

In this study, we investigate the sources of advertising revenue earned by Google, Bing, and DuckDuckGo by examining ad clicks and ad revenue on *navigational* and *brand* searches. Following Broder (2002), we define "navigational searches" as those with an *intent to navigate to a specific domain* (e.g., the example search shown in Figure 1). In contrast to navigational searches, we define "brand searches" as those with both an *intent to navigate to a specific domain* and a *broader, exploratory intent* (e.g., "instagram reels"). We say that the "focal brand defends" if the brand that is the target of the query advertises in the top position on the page (e.g., Instagram in Figure 1). § 3.2 describes how we operationalize these definitions. Finally, we compare the properties of ads on navigational and brand searches to the properties of ads on 18 categories of "non-brand search" (e.g., Shopping, Finance & Banking, and News & Media). We use the phrase "non-brand search" to describe queries that are neither brand nor navigational. Specifically, our research questions are:

- **RQ1:** What fraction of Google's search ad revenue comes from navigational and brand search?
- **RQ2:** How effective are competitor ads (a) across navigational, brand, and non-brand search, and (b) with and without focal brand defense?
- **RQ3:** Are navigational and brand search similarly lucrative for Google, Bing, and DuckDuckGo?

To answer these questions, we rely on logs of searches performed by a representative panel of 926 US residents. Our dataset contains the full web browsing history of participants from August–December 2020, including all searches they executed on Google, Bing, and DuckDuckGo. Additionally, our dataset contains copies of the exact search engine results pages (SERPs) that Google showed to participants. We parse the SERPs to identify ad impressions and parse the browsing logs to identify ad clicks. Additionally, we merge the ad clicks with ad cost-per-click (CPC) estimates from Google and Microsoft's advertiser APIs to quantify search engine revenue across search categories. Finally, we leverage keyword-level $\frac{mobile}{desktop}$ search ratios from the advertiser APIs to extrapolate our desktop revenue estimates to mobile, which is an especially important modality for Google Search.[1] We present the following key findings:

- **RQ1:** We estimate that navigational and brand ads account for 12.4–17.8% and 14.2–20.5% of Google Search's revenue, respectively. Google earns more from navigational and brand ads on desktop than on mobile.
- **RQ2:** We find exploratory (non-significant) evidence that navigational ads are the least effective for competitors and that defense lowers the effectiveness of navigational ads the most.
- **RQ3:** We estimate that Bing earns slightly more from navigational and brand search ads than Google.

§ 2 describes the search advertising market and prior work on navigational and brand search advertising. We introduce our dataset and measurement approaches in § 3, followed by analysis in § 4. We discuss our findings, implications, and the limitations of this study in § 5.

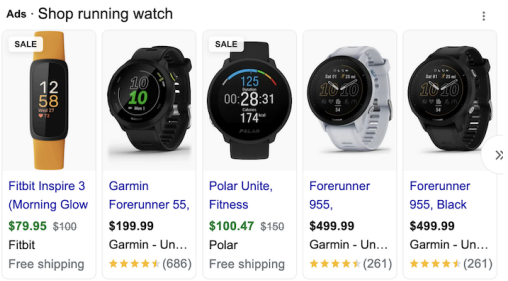


Figure 2: Shopping ads for the query "running watch" on Google. Shopping ads on Bing and DuckDuckGo have a similar appearance.

2 Background

2.1 Search Advertising

Search engines run real-time auctions that allocate ads to positions on the SERP. Advertisers submit maximum CPC bids—i.e., the maximum amount they are willing to pay for a click—on keyword phrases, e.g., "running watch." However, bids are not the only factor that determine an ad's position. For example, Google's auction also takes ad quality, ad format, and search context into account.[2] Advertisers are charged according to a generalized second price (GSP) auction mechanism, in which one pays the minimum cost required to keep the ad in the position where it was served (Edelman, Ostrovsky, and Schwarz 2007).

Two prominent formats for search ads are text and shopping ads. Figure 2 shows an example shopping ads carousel on Google Search. Aside from the ability to display a photo and price, the only difference between shopping and text ads is that advertisers bid on product attributes in the former and keywords in the latter. In both situations, ad position and payment are determined by a GSP with CPC bidding.[3]

Eye-tracking studies have consistently shown that people typically scan SERPs from top to bottom (Cutrell and Guan 2007; Pan et al. 2007). One region frequently scanned is located near the top left of the page and referred to as the "Golden Triangle" (Hotchkiss, Alston, and Edwards 2005). Though often discussed in the context of web search (Granka, Joachims, and Gay 2004; Nettleton and González-Caro 2012; Papoutsaki, Laskey, and Huang 2017), this top-to-bottom viewing pattern has also been found in

[1] Publicly available statistics purport that 70% of Google ad clicks are on mobile devices (Statista 2022).

[2] https://support.google.com/google-ads/answer/6366577

[3] https://support.google.com/google-ads/answer/2454022

eye-tracking studies examining query autocomplete interfaces (Hofmann et al. 2014), commercial website homepages (Granka, Hembrooke, and Gay 2006), and webpages more broadly (Buscher, Cutrell, and Morris 2009; Dimitrov et al. 2016). Together, these studies demonstrate the robust value, in terms of user attention, that ads receive from placement at the top of the search rankings.

2.2 Navigational Search

Broder (2002) introduced an influential taxonomy for web search, sorting searches into three classes: navigational, informational, and transactional. Navigational searches are those that a person uses to navigate to a specific known website. Early manual analyses of AltaVista query logs identified 11.7–20% of searches as navigational (Broder 2002; Rose and Levinson 2004). Jansen, Booth, and Spink (2008) proposed a rule-based classifier for search intent. They operationalized navigational search by checking whether the search query contained (a) domain suffixes or (b) company/organization names. A comparison to manual classification found that this classifier had an accuracy of 74% across the three intents. This approach identified 10% of searches on the metasearch engine Dogpile as navigational.

Teevan, Liebling, and Geetha (2011) formalized the definition of navigational searches using a query's click entropy. Intuitively, a query that produces clicks on a specific domain with high probability strongly signals navigational intent. Unfortunately, we cannot adopt this classification approach in our study because the vast majority of Google queries in our dataset (see § 3.1) do not have sufficient click volume to reliably estimate click entropy ($> 99\%$ have ≤ 5 clicks).

Finally, some work has studied ads on navigational search and users' interaction with them. Multiple studies have identified higher CTR on navigational ads compared to informational and transactional ones (Schultz 2020; Jansen and Spink 2009; Ashkan et al. 2009). Furthermore, Schultz (2020) found that navigational ads have a lower CPC because they face less competition and have a higher conversion rate.

2.3 Brand Search Advertising

Studies about ads on navigational search are closely related to literature about brand search advertising and competitive poaching. There are many empirical studies about the effectiveness of brand search advertising. Interestingly, these have yielded quite different results. Early research from Google found that 50% of ad clicks that match the top organic result are incremental, i.e., would not have occurred without an ad campaign (Chan et al. 2012). Subsequently, researchers ran experiments that halted eBay's brand search ads and found that organic traffic replaced 99.5% of paid traffic (Blake, Nosko, and Tadelis 2015). However, a similar follow-up experiment with a much smaller brand, Edmunds.com, found that organic traffic replaced $< 50\%$ of their paid traffic (Coviello, Gneezy, and Goette 2017).

Recent work has also studied the effects of competition on the effectiveness of brand search advertising. Simonov, Nosko, and Rao (2018) found that brand ads, in the absence of competition, increased focal brand clicks by 1–4%. This effect was smaller for larger brands, e.g., eBay. In a follow-up study, the authors found that competitors poached 6–15% of clicks when the keyword-owning brand was removed from the top ad position (Simonov and Hill 2021), but the quality of poached clicks was often low. On the other hand, Golden and Horton (2021) found that winning keyword auctions on a competitor's brand name had no effect on clicks to the poacher's site.[4]

Finally, there is theoretical work about advertiser strategy under the dynamics of competitive poaching. Sayedi, Jerath, and Srinivasan (2012) found that smaller firms poach more often on larger firms' keywords. Additionally, Desai, Shin, and Staelin (2014) found that bidding on competitors' brands can create a prisoner's dilemma where both firms are worse off and the search engine captures the profits.

		Participants		
		N	**%**	**US Census**
Gender	Female	415	52.5	50.4
	Male	375	47.5	49.6
Race and Ethnicity	White	583	73.9	58.9
	Black	92	11.8	13.6
	Hispanic	64	8.1	19.1
	Asian	16	2.1	6.3
	Native American	6	0.8	1.3
	2+ races	11	1.5	3.0
	Other	14	1.8	–
Age	< 18	0	0.0	21.7
	18-64	589	74.6	50.4
	65+	200	25.4	17.3

Table 1: Demographics of participants who installed our browser extension, compared to the US Census.

3 Data and Measurement

In this section, we present our datasets and data sensemaking procedures.

3.1 Data Collection

Beginning in August 2020, we engaged the survey company YouGov to recruit a panel of US residents to take a survey and install a browser extension we developed for Chrome and Firefox. YouGov reached out to a nationally-representative subsample of people who had previously taken part in the 2018 Cooperative Congressional Election Survey. Out of 2,000 respondents, 926 people completed the survey and installed the browser extension. Table 1 describes the demographics of participants who installed our browser extension. Compared to the US Census,[5] our participants were Whiter (73.9% vs. 58.9%) and older (because we only recruited participants older than 18). We collected data from participants' web browsers through December 2020. Our

[4] Another challenge that makes existing studies difficult to compare is that some define a "brand ad" as a query containing a brand name, while others only consider exact matches.

[5] https://www.census.gov/quickfacts/fact/table/US/PST045223

Search Engine	№ Users	№ Searches	% Nav.	% Brand	Ad API Coverage
Google	703	325687	6.5	10.9	84.9
Bing	188	124763	5.1	10.9	75.5
DDG	48	10553	7.9	9.8	80.7

Table 2: Descriptive statistics of our dataset, including number of users and searches, fraction of navigational and brand searches, and fraction of ad clicks covered by data from the search engines' ad APIs.

study was IRB-approved and we describe our procedures for protecting participants in § 5.2. It's important to note that our study period includes both seasonal (the 2020 US election and the winter holidays) and exceptional (the Covid-19 pandemic) events that affected digital advertising. Thus, our analysis represents online search advertising at a specific point in time.

Our browser extension collected several types of data from participants' web browsers, two of which we leverage in this study: (1) browsing activity and (2) snapshots of Google SERPs. The browsing activity data contains the timestamped sequence of URLs that participants viewed in their browser during our observation window. The snapshot data contains the complete HTML of the Google SERPs that participants saw in response to their search queries. Our extension did not collect snapshots of Bing or DuckDuckGo SERPs. However, we re-crawled participants' Bing and DuckDuckGo search queries in January 2023, from an IP address in Boston, to produce approximations of the SERPs they saw. We parse the Google SERPs using `WebSearcher` (Robertson and Wilson 2020) and the Bing and DuckDuckGo SERPs using `SearchParser`, an open source package we created for this project.[6]

We define a participant as a *user of a search engine* if they made at least one search on that search engine during our observation window. This definition permits participants to count as users of multiple search engines.

Table 2 shows the number of Google, Bing, and DuckDuckGo users in our sample and the total number of searches they made. We do not count searches made on vertical search engines, e.g., Google Shopping and Bing Images. We filter out searches that contain the string "Gamification_DailySet" in the URL query parameter, which are queries that are automatically generated when a user interacts with a Bing Rewards quiz.[7]

3.2 Navigational and Brand Search Classification

One key goal of our study is to examine ad impressions and revenue in response to navigational and brand searches. To facilitate this, we identify navigational and brand searches using methods from Jansen, Booth, and Spink (2008) and Simonov and Hill (2021).

[6]https://github.com/jlgleason/SearchParser

[7]https://www.microsoft.com/en-us/rewards/search-and-earn

	Query	Top Organic Domain	№ Ad Clicks
Navigational	amazon	amazon	32
	walmart	walmart	17
	best buy	bestbuy	13
	yougov	yougov	13
	bed bath and beyond	bedbathandbeyond	12
Brand	citibank	citi	9
	nintendo switch	nintendo	4
	microsoft store	microsoft	4
	amazon smile	amazon	3
	usps tracking	usps	2

Table 3: Five navigational and brand searches with the highest number of ad clicks, provided ≥ 5 distinct participants searched the query.

Entity to Domain Mapping To identify navigational and brand searches we needed a mapping from entities to their domain names on the internet. Following Simonov and Hill (2021), we use the online database Curlie,[8] which "strives to be the largest human-edited dictionary of the web." We crawled Curlie from February–April 2023 and collected a dataset of 1.3 million entity:URL pairs, which we release.[9]

Navigational Search Following Broder (2002), we define navigational searches as those with an *intent to navigate to a specific domain*. To operationalize this, we check whether each query in our corpus matches the top organic domain on the corresponding SERP (or whether it maps to this domain in the Curlie directory) (Simonov and Hill 2021).[10] The top half of Table 3 shows the five navigational searches in our dataset with the highest number of ad clicks (conditional on ≥ 5 users searching the query).

Brand Search In contrast to navigational searches, we define brand searches as those with both an *intent to navigate to a specific domain* and a *broader, exploratory intent*. To operationalize this, we check if the beginning or end of each non-navigational query in our corpus matches the top organic domain on the SERP (or if either maps to this domain in the Curlie directory). Thus, "rei camping mattress" and "camping mattress rei" are both brand searches, while "rei" is a navigational search. The bottom half of Table 3 shows the five brand searches in our dataset with the most ad clicks.

Table 2 shows that our point estimates for navigational search volume range from 5.1% of searches on Bing to 7.9% of searches on DuckDuckGo. Our point estimates for brand search volume range from 9.8% of searches on DuckDuckGo to 10.9% of searches on Google and Bing. Thus, navigational and brand search volume are individually closer to the 10% from Jansen, Booth, and Spink (2008), while their sum is closer to the 21% from Teevan, Liebling, and Geetha (2011).

[8]https://curlie.org/

[9]https://github.com/jlgleason/navigational-ads

[10]Specifically, we check if the Jaro-Winkler similarity is ≥ 0.95 to catch typos and abbreviations (Cohen et al. 2003).

3.3 SERP Topic Classification

Another key goal of our study is to compare ad clicks and revenue on navigational and brand searches to ad clicks and revenue on searches for other topics. To achieve this, we assign a category label to each Google query in our dataset based on the domains that appear in links on the SERP, using the following approach.

First, we assign each domain that appears in our sample of SERPs to one of 91 categories using the FortiGuard domain classification service. FortiGuard is a cybersecurity service that enables organizations to filter web traffic based on the category of the destination, such as Social Networking, Streaming Media, or Pornography. Prior work found that FortiGuard had greater coverage of domains and more accurate category labels than other similar services (Vallina et al. 2020).

Second, we label each SERP as a weighted distribution over categories, with weights taken from the empirical distribution of clicks over vertical ranks in our dataset. This accounts for decaying attention (Papoutsaki, Laskey, and Huang 2017) as users move down the SERP. For example, if a SERP contains three total links, the first to a shopping site (e.g., `amazon.com`), the second to a health site (e.g., `walgreens.com`), and the third to a social networking site (e.g., `facebook.com`), our method would assign the categorical distribution $[0.50, 0.28, 0.22]$. We are able to parse ≥ 1 domain and generate a categorical distribution for 97% of SERPs in our sample using this method. For each non-brand search (i.e., not navigational nor brand), we randomly sample a category from its categorical distribution.

3.4 Ad Measurement

In this section, we explain how we identified ad impressions and ad clicks, and how we obtained CPC data.

Ad Impressions We count the number of text and shopping ads on each Google SERP using `WebSearcher` (Robertson and Wilson 2020).[11] One limitation of `WebSearcher` is that it cannot parse ads in the right-hand column of SERPs. Additionally, our data cannot tell us which ads, if any, participants actually viewed. Therefore, we only count ad impressions that appear among the top four vertical ranks of the SERP and among the top four horizontal slots of shopping carousels. We adopt this approach because prior work and our own dataset indicate that people pay the majority of their attention to the top four vertical ranks on the SERP (69.4% of clicks in our dataset and an even higher fraction of examination time in Papoutsaki, Laskey, and Huang (2017)).

Ad Clicks Both Google and Microsoft use URL parameters for conversion tracking and attribution (Ursu, Simonov, and An 2021). Google allows advertisers to auto-tag URLs using the Google Click ID (`gclid`)[12] or manually tag URLs using UTM parameters.[13] Similarly, Microsoft allows advertisers to tag URLs using the Microsoft Click ID (`msclkid`)[14] or UTM parameters.[15] Thus, we record a Google ad click if the URL immediately following a Google search contains `gclid` or `utm_source=google`. We record a Bing or DuckDuckGo ad click if the URL immediately following a Bing or DuckDuckGo search contains the `msclkid` URL parameter or `utm_source=bing`. Microsoft tracking parameters allow us to identify ad clicks on both Bing and DuckDuckGo, because DuckDuckGo uses Microsoft's advertising network.[16] Finally, we also identified an Amazon-specific URL tracking parameter (`hvnetw`) that is appended to some Amazon ads instead of the Google or Microsoft tracking parameters. We also record an ad click if the URL immediately following any search contains this parameter.

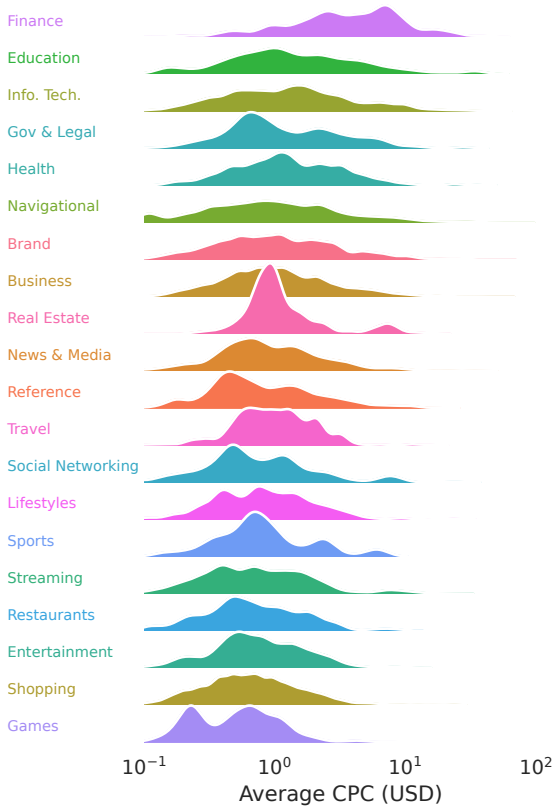


Figure 3: Google CPC distributions across search categories.

Ad CPC We collect CPC data using the Google Ads[17] and Bing Ads[18] APIs. These APIs allow advertisers to research and monitor keyword performance and manage active

[11] Version 0.2.15.

[12] https://support.google.com/google-ads/answer/9744275

[13] https://support.google.com/analytics/answer/1733663

[14] https://about.ads.microsoft.com/en-us/blog/post/january-2018/conversion-tracking-update-on-bing-ads

[15] https://help.ads.microsoft.com/apex/index/3/en/56762

[16] https://help.duckduckgo.com/duckduckgo-help-pages/company/ads-by-microsoft-on-duckduckgo-private-search/

[17] https://developers.google.com/google-ads/api/

[18] https://learn.microsoft.com/en-us/advertising/guides

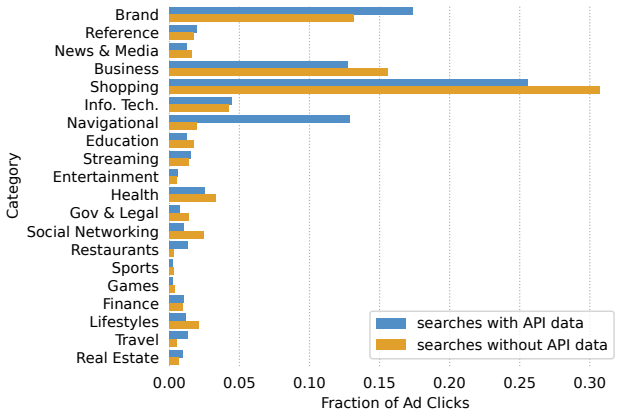


Figure 4: Evaluating selection bias in the coverage of Google searches, grouped by search category, for which API data was and was not available.

ad campaigns. They also offer several parameters when requesting bid prices and CPCs. When collecting data, we controlled for six parameters—keyword, date, ad match type, ad position, location, and language—that we aligned across both APIs:

1. **Keyword**: For each query in our corpus, we removed invalid symbols[19] then submitted (a) the full query string, and (b) noun phrases extracted from the query using `spaCy`'s default pipeline.[20] If multiple phrases returned results, we used the phrase that returned the highest CPC, in accord with the search engines' auction structure.
2. **Date**: We crawled both APIs in February 2023 from an IP address in Boston. Both APIs only return data corresponding to a 30 day lookback window. Thus, an important limitation is that the time periods for our CPC data and our browsing activity data do not match.
3. **Ad Match Type**: When submitting full query strings, we use "exact match," which targets searches that *match* a keyword's meaning.[21] When submitting noun phrases, we use "phrase match," which targets searches that *include* a keyword's meaning.
4. **Ad Position**: "Top of the first page," corresponding to our operationalization of ad impressions.
5. **Location**: United States, corresponding to the location of our participants.
6. **Language**: English, corresponding to the most used language in the United States.

We focus on three measures—bids, CPC, and device type—that the APIs return:

1. **Bids**: Microsoft returns the mean bid price for a keyword, while Google returns the 20$^{\text{th}}$ and 80$^{\text{th}}$ percentile bid prices.

[19]https://support.google.com/google-ads/answer/7476658
[20]https://spacy.io/models/en#en_core_web_sm
[21]https://support.google.com/google-ads/answer/7478529

Variable	Description
n	number of search engine users
s_i	number of searches made by user i
u_i	YouGov survey weight for user i
$\mathbb{I}_{ij}^{g}$	search j from user i is part of category g
a_{ij}^{t}	text ad impressions on search j
a_{ij}^{s}	shopping ad impressions on search j
a_{ij}	total ad impressions on search j ($a_{ij}^{t} + a_{ij}^{s}$)
c_{ij}	ad clicks on search j
p_{ij}	average CPC for search j
w_{ij}	$\frac{\text{mobile}}{\text{desktop}}$ search ratio for search j

Table 4: Notation for search, ad, and CPC variables.

2. **CPC**: Microsoft returns the average CPC for a keyword, while Google returns a "forecast curve" that maps bid prices to CPCs.[22] Thus, we use the average of the 20$^{\text{th}}$ and 80$^{\text{th}}$ percentile bid prices to look up an "average" CPC for a keyword on Google.
3. **Device Type**: Microsoft returns ad impression volumes that are disaggregated across desktop and mobile, while Google returns search volumes that are disaggregated across desktop and mobile.

Figure 3 shows empirical CPC distributions on Google for the top 20 search categories in our dataset, which cover 91% of all searches. The distributions include all SERPs with at least one ad impression and a non-zero CPC. `Finance & Banking` searches have the highest mean CPC ($5.10). `Navigational` ($1.86) and `Brand` ($1.84) have mean CPCs that are slightly higher than the overall mean ($1.66). `Shopping` searches have the second lowest mean CPC ($0.96). Within each category, there is significant variation in CPCs across keywords.

Figure 4 assesses one potential threat to the validity of our revenue estimates: selection bias in ad API coverage. Specifically, Figure 4 compares the ad click distribution in the subset of searches with API coverage to the ad click distribution in the subset without API coverage. The fraction of ad clicks on navigational searches is substantially higher in the subset with API coverage (12.8% vs. 2.0%, $P \approx 0$ using two-sided z-test). Overall, we expect that this selection bias toward navigational ad clicks is offset by our low estimate of navigational search volume (4.9–7.8% vs. 10–21% in previous studies).

4 Results

We now present the results of analysis on our datasets. We begin by examining Google Search on the desktop, where our data is richest, and then expand our frame to compare Google against other search engines across desktop and mobile modalities. Table 4 contains notation for the variables that we use throughout this section.

[22]https://developers.google.com/google-ads/api/docs/keyword-planning/generate-forecast-metrics#generate_a_forecast_curve

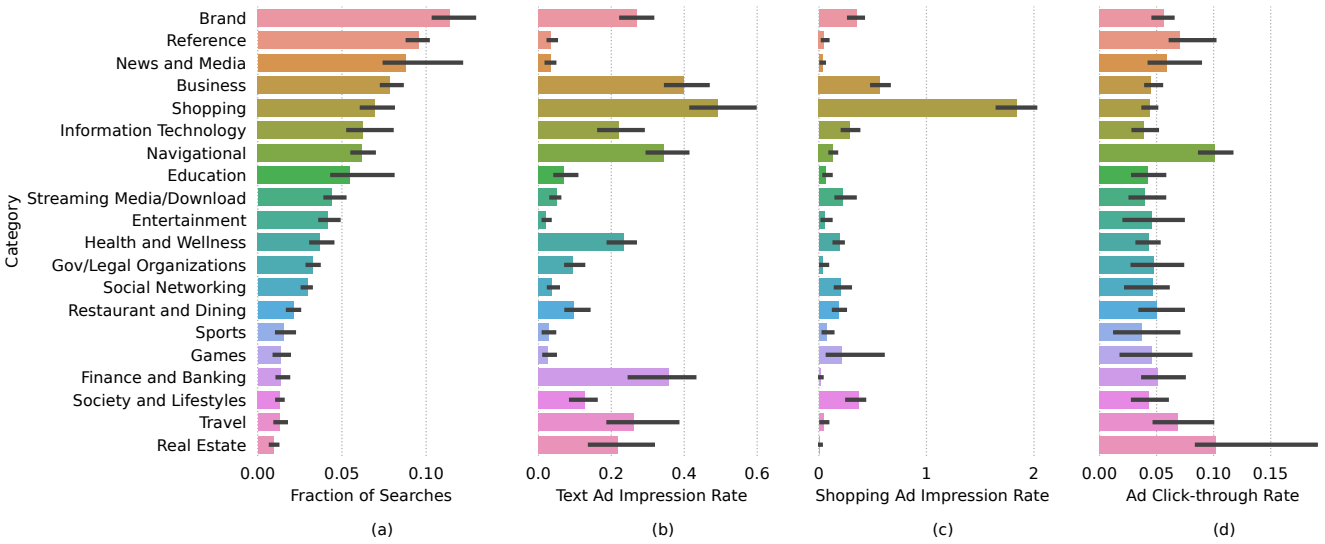


Figure 5: Fraction of searches, ad impression rate, and ad CTR across search categories for desktop Google Search.

4.1 Google Search Ads on the Desktop

We begin by providing a descriptive analysis of ad impressions and clicks on the desktop version of Google Search. Figure 5 shows the (1) fraction of searches, (2) text ad impression rate, (3) shopping ad impression rate, and (4) ad CTR across the top 20 search categories. The text ad impression rate in category g is defined as the number of text ad impressions divided by the number of searches (where users are weighted according to their survey weight):

$$\frac{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} a_{ij}^t \mathbb{I}_{ij}^g}{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} \mathbb{I}_{ij}^g}. \tag{1}$$

The shopping ad impression rate is defined analogously, swapping a_{ij}^s in for a_{ij}^t.

The ad CTR in category g is defined as the number of ad clicks divided by the number of ad impressions:

$$\frac{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} c_{ij} \mathbb{I}_{ij}^g}{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} a_{ij} \mathbb{I}_{ij}^g}. \tag{2}$$

We compute 95% confidence intervals using the BCa bootstrap at the user level with 1000 replications (Efron 1987). In each bootstrap iteration, we sample categories for searches that are not labeled as navigational or brand.

Figure 5a shows that Brand is the largest category on the desktop, accounting for 10.5–12.8% of searches. Navigational is the 7th largest category, accounting for 5.6–6.9% of desktop searches. Figures 5b and 5c show that Reference and News & Media, the 2nd and 3rd largest categories, have very low ad impression rates. Nine categories—Shopping, Business, Finance & Banking, Navigational, Brand, Travel, Health & Wellness, Information Technology, and Real Estate—have high text ad impression rates (0.2–0.5 text ads per SERP). Shopping SERPs have far more shopping ads—almost two per SERP—than any other category. These patterns align with our expectations and lend a strong degree of face validity to our ad impression measurement.

Figure 5d shows that ad CTRs often have higher variance than ad impression rates. This is because there are far fewer ad clicks than impressions. Navigational has the 2nd highest ad CTR—10.1% (8.8–11.6%)—which aligns with previous work (Schultz 2020; Jansen and Spink 2009; Ashkan et al. 2009). Brand has the 6th highest ad CTR: 5.6% (4.7–6.4%).

The combination of high search volumes, ad impression rates, and ad CTRs provides our first piece of evidence that navigational and brand ads may be systemically important for Google's business.

4.2 Estimating Google Search Ad Revenue

Next, we estimate Google Search's ad revenue across categories on desktop and mobile using the keyword CPC and $\frac{mobile}{desktop}$ search ratios from the Google Ad API (§ 3.4).

Figure 6 shows the (1) desktop revenue share, (2) mobile revenue share, and (3) combined revenue share across the top 20 categories. We define the desktop revenue share in category g as the fraction of CPC-weighted ad clicks:

$$\frac{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij} \mathbb{I}_{ij}^g}{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij}}. \tag{3}$$

We define the mobile revenue share in category g as the fraction of CPC-weighted ad clicks, adjusted for $\frac{mobile}{desktop}$ search ratios:

$$\frac{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij} w_{ij} \mathbb{I}_{ij}^g}{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij} w_{ij}}. \tag{4}$$

Finally, the total revenue share in category g combines desktop and mobile revenue:

$$\frac{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij} (1 + w_{ij}) \mathbb{I}_{ij}^g}{\sum_{i=1}^{n} u_i \sum_{j=1}^{s_i} p_{ij} c_{ij} (1 + w_{ij})}. \tag{5}$$

As before, we compute 95% confidence intervals using the BCa bootstrap at the user level with 1000 replications.

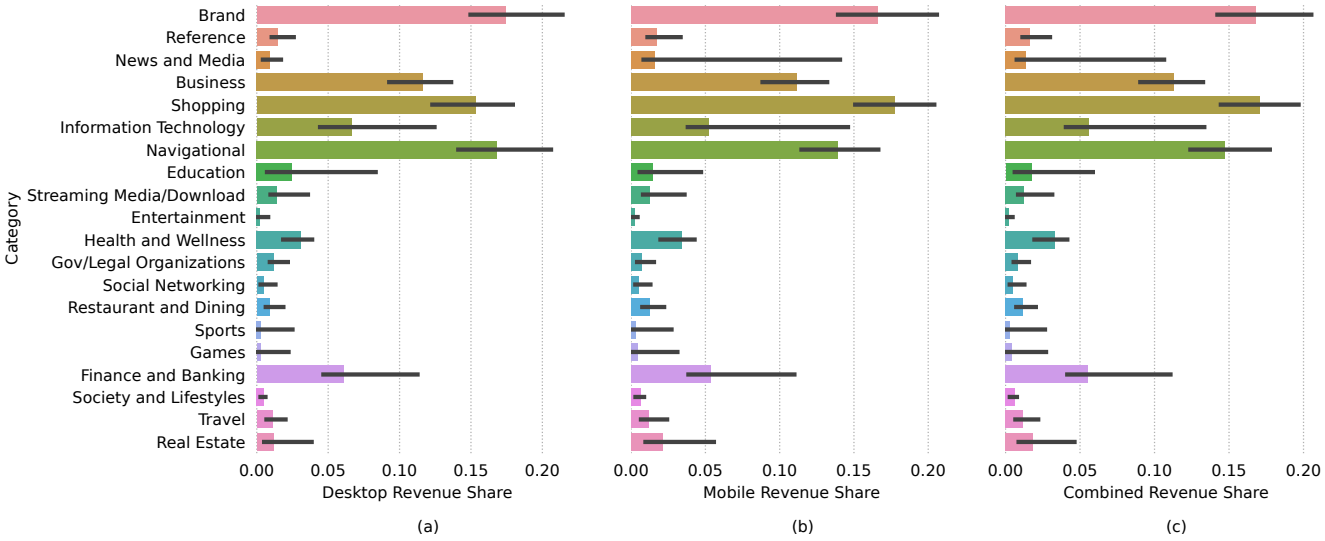


Figure 6: Desktop, mobile, and combined revenue shares across search categories for Google Search.

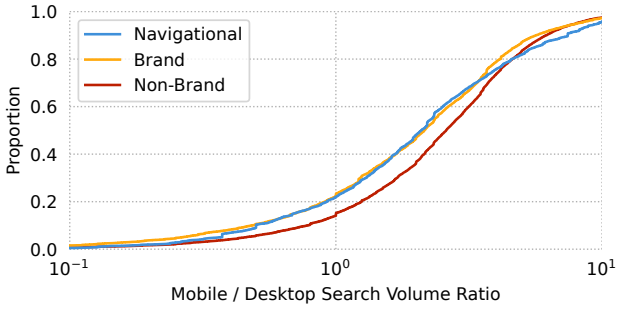


Figure 7: Comparing CDFs of $\frac{mobile}{desktop}$ search volume across navigational, brand, and non-brand searches.

Figure 6a shows that Brand is the largest revenue category on the desktop, accounting for 17.4% (95% CI 15.0–21.4%) of revenue. Navigational is the 2nd largest revenue category, accounting for 16.8% (14.1–20.6%). Shopping is the 3rd largest revenue category, accounting for 15.3% (12.3–17.9%) of revenue. Three other categories with non-negligible revenue shares are Business, Information Technology, and Finance & Banking.

Figure 6b, however, shows that the story changes on mobile devices. Shopping's revenue share increases to 17.7% (15.1–20.4%). Meanwhile, Brand and Navigational's revenue shares decrease to 16.6% (13.9–20.6%) and 13.9% (11.5–16.7%), respectively. Business, Information Technology, and Finance & Banking retain non-negligible revenue shares on mobile.

Finally, Figure 6c shows that Shopping is the largest revenue category overall, accounting for 17.1% (14.4–19.7%) of revenue. This reflects the importance of mobile search to Google's ad business (Statista 2022). Brand accounts for 16.8% (14.2–20.5%) of revenue overall, while Navigational accounts for 14.7% (12.4–17.7%) of revenue overall.

Figure 7 illustrates the reason for the decreases in Navigational and Brand revenue share on mobile. While 84% of non-brand searches have more mobile than desktop search volume ($w_{ij} > 1$), only 78% of navigational searches and 75% of brand searches have more mobile volume. This finding aligns with recent research from the SEO industry, which also found that navigational searches on Google are less common on mobile devices (Beus 2021).

4.3 Estimating Google Search Ad Effectiveness

The large fraction of ad revenue from navigational and brand search raises the question: how effective are ads on these types of searches compared to ads on non-brand searches?

Following Simonov and Hill (2021), we operationalize the "focal brand" as the top organic domain (operationalized as effective second-level domain) on the SERP. We use the term "competitors" to describe advertisers who are not the focal brand. The focal brand "defends" if it advertises in the top position on the page, i.e., above any competitors. Competitors can occupy up to the top four positions on the SERP (i.e., the "mainline"). However, a competitor can only occupy the top position if the focal brand does not defend. We analyze effectiveness from the perspective of competitors using two metrics: (1) competitor ad CTR and (2) "high-quality" competitor ad CTR. A high-quality click occurs when a user does not return to Google Search within 30 seconds of the click (Simonov and Hill 2021).

First, we benchmark our estimates of effectiveness on navigational search against those in Simonov and Hill (2021). This is important because Simonov and Hill (2021) leverage search data with randomized ad allocation, while our search data is purely observational. Despite representing a different population (Bing vs. Google) and year (2017 vs. 2020), these estimates represent causal effects and thus inform us about the magnitude of bias in our data. Following

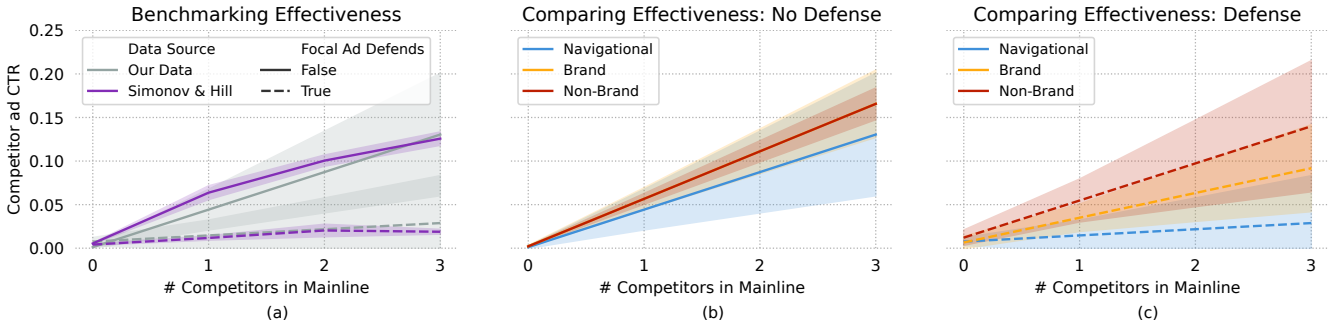


Figure 8: Ad effectiveness measured using competitor ad CTR. (a) benchmarks estimates on navigational search against Simonov and Hill (2021), (b) compares effectiveness across navigational, brand, and non-brand search without defense, and (c) compares effectiveness with defense. Shaded regions are 95% CIs for the predicted mean.

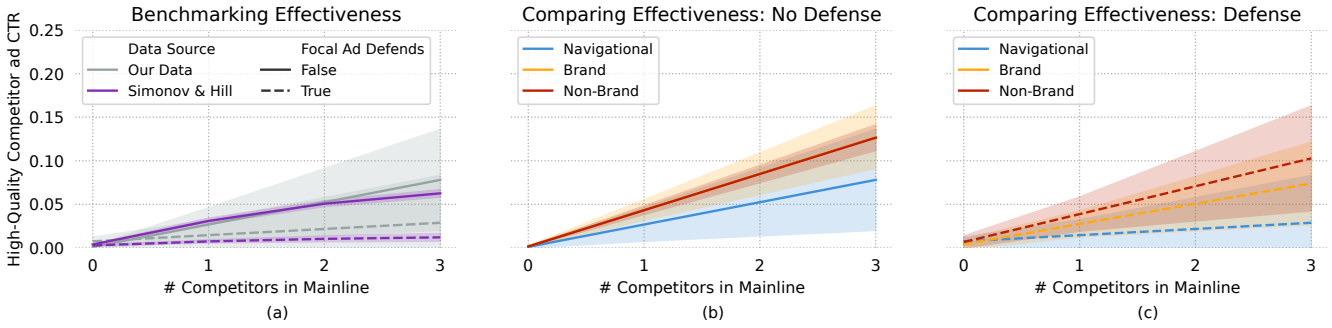


Figure 9: Ad effectiveness measured using *high-quality* competitor ad CTR. (a) benchmarks estimates on navigational search against Simonov and Hill (2021), (b) compares effectiveness across navigational, brand, and non-brand search without defense, and (c) compares effectiveness with defense. Shaded regions are 95% CIs for the predicted mean.

Simonov and Hill (2021), we restrict our sample to searches with zero shopping ads, ≤ 1 click, and one of the top 1,500 focal brands (ordered according to our data). Thus, these effects represent competitors targeting text ads against navigational searches for well-known brands.

Figures 8a and 9a compare the two sets of estimates. Specifically, we plot competitor ad CTRs as we vary (1) whether the focal brand defends and (2) the number of competitors in the mainline. Simonov and Hill (2021) estimate a separate effect for each combination of defense and competitors, adjusted for the exact slate of focal and competitor ads. Given our smaller sample, we ignore the specific ad slate and smooth over the number of competitors using a linear probability model. Let i index searches, y represent a click on any competitor ad, x represent the number of competitors, and d represent whether the focal brand defends. Our model is:

$$y_i = \beta_0 + \beta_1 x_i + \beta_2 d_i + \beta_3 x_i d_i + \epsilon_i \tag{6}$$

where β_3 and β_1 represent competitor effectiveness with and without focal brand defense, respectively. We cluster standard errors at the user level. Figures 8a and 9a demonstrate that our model produces similar estimates of both competitor ad CTR and high-quality competitor ad CTR to Simonov and Hill (2021), though our estimates are more uncertain.

This comparison gives us confidence to explore effectiveness across navigational, brand, and non-brand searches. To do so, we add a term representing search type (q) to Equation 6 and interact it with the number of competitors (x) and focal brand defense (d). Specifically, let q^b and q^{nb} represent brand and non-brand searches, respectively:

$$\begin{aligned} y_i = \beta_0 &+ \beta_1 x_i + \beta_2 d_i + \beta_3 q_i^b + \beta_4 q_i^{nb} \\ &+ \beta_5 x_i d_i + \beta_6 x_i q_i^b + \beta_7 x_i q_i^{nb} + \beta_8 d_i q_i^b \\ &+ \beta_9 d_i q_i^{nb} + \beta_{10} x_i d_i q_i^b + \beta_{11} x_i d_i q_i^{nb} + \epsilon_i \end{aligned} \tag{7}$$

Figures 8b and 8c contrast competitor ad CTR, while Figures 9b and 9c contrast high-quality competitor ad CTR. We make three exploratory observations, which Tables 5 and 6 express as contrasts between regression coefficients. Note that we cannot reject any null hypotheses after adjustment for multiple comparisons.

1. Navigational ads are the least effective for competitors.
2. Focal brand defense lowers the effectiveness of navigational ads the most and lowers the effectiveness of brand ads more than non-brand ads.
3. Undefended brand and non-brand ads have similar effectiveness.

Obs.	Contrast	Coef	Std. Err	Adj. Pval
1	$\beta_6 = 0$	0.011	0.014	0.562
1	$\beta_7 = 0$	0.011	0.012	0.562
1	$\beta_6 + \beta_{10} = 0$	0.021	0.013	0.407
1	$\beta_7 + \beta_{11} = 0$	0.036	0.016	0.224
2	$\beta_{10} = 0$	0.010	0.019	0.690
2	$\beta_{11} = 0$	0.024	0.020	0.562
2	$\beta_{11} = \beta_{10}$	0.014	0.016	0.562
3	$\beta_6 = \beta_7$	-0.000	0.007	0.999

Table 5: Contrasts of competitor ad CTR across search scenarios (see Equation 7). P-values are adjusted using Benjamini-Hochberg.

Obs.	Contrast	Coef	Std. Err	Adj. Pval
1	$\beta_6 = 0$	0.016	0.011	0.411
1	$\beta_7 = 0$	0.016	0.010	0.411
1	$\beta_6 + \beta_{10} = 0$	0.016	0.013	0.411
1	$\beta_7 + \beta_{11} = 0$	0.025	0.014	0.411
2	$\beta_{10} = 0$	-0.000	0.017	0.999
2	$\beta_{11} = 0$	0.009	0.017	0.806
2	$\beta_{11} = \beta_{10}$	0.009	0.014	0.806
3	$\beta_6 = \beta_7$	0.000	0.006	0.999

Table 6: Contrasts of *high-quality* competitor ad CTR across search scenarios (see Equation 7). P-values are adjusted using Benjamini-Hochberg.

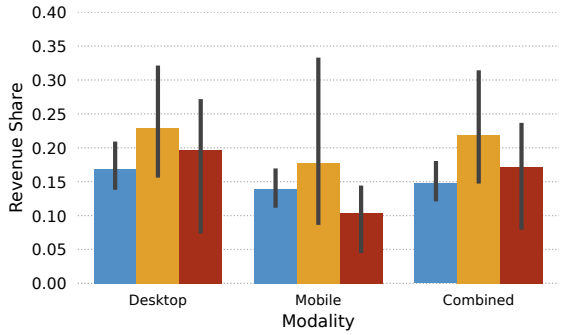
4.4 Ad Revenue Across Search Engines

Finally, Figure 10 compares the desktop, mobile, and combined revenue shares from navigational and brand search across Google, Bing, and DuckDuckGo. Bing and DuckDuckGo revenue shares are also calculated using Equations 3, 4, and 5, with two caveats—the Bing Ads API returns: (1) average CPCs (p_{ij}) that are specific to desktop/mobile, and (2) device ratios (w_{ij}) based on ad impression volume instead of search volume.

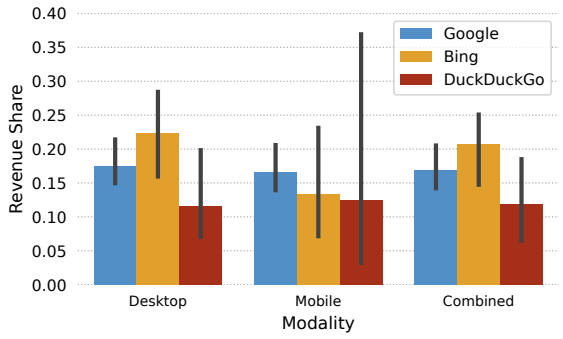
Figure 10a demonstrates that Bing has the highest navigational revenue share: 22.0% (95% CI 15.0–31.1%). DuckDuckGo has a similar navigational revenue share to Google (17.2% vs. 14.7%), but a much wider confidence interval (8.2–23.4%). Bing and DuckDuckGo's wider confidence intervals reflect their smaller sample of users. We also observe that navigational revenue shares are higher on desktop than on mobile for all three search engines.

Figure 10b shows that Bing also has the highest brand revenue share: 20.7% (14.7–25.1%). However, Google's brand revenue share is quite similar: 16.8% (14.2–20.5%). DuckDuckGo's brand revenue share is slightly smaller, but again has a wide confidence interval: 11.8% (6.5–18.5%). As before, brand revenue shares are generally higher on desktop than on mobile.

Overall, this section demonstrates that it is not only the dominant search engine, Google, that benefits from large navigational and brand ad revenue shares. These are lucrative markets for any general search engine.



(a) Navigational Search



(b) Brand Search

Figure 10: Navigational and brand revenue share across modalities and search engines.

5 Discussion

The primary findings from our study are accurate estimates of keyword ad revenues for major search engines, stratified across different categories of ads. In § 4.2 we find that brand and navigational ads are the 2nd and 3rd largest revenue categories, respectively, on Google Search, collectively accounting for 27.7–36.4% of Google's search ad revenue. This is on the order of $50 billion annually. In § 4.4 we find that Bing earns an even greater share of its revenue from navigational and brand ads than Google.

These findings are vital because these forms of ads—poaching competitor's trademarked brands—are contentious. Google successfully settled several lawsuits in the US that attempted to halt this practice (Goldman 2008, 2010; Reuters) and our results present a strong motivation to do so: these ads are very lucrative. However, case law in the US surrounding these business practices is far from settled, and other jurisdictions are just beginning to grapple with them (Singh 2023). Our results demonstrate the stakes of this conflict, in terms of both potential revenues for search engines and expenditures by advertisers.

Irrespective of how law and policy decisions play out, a more fundamental question is about the effectiveness of these ads for advertisers. In § 4.3, we find exploratory (non-significant) evidence that navigational ads are less effective in terms of competitor CTR than brand and non-brand ads. Brand ads, on the other hand, are more complex: without de-

fense, brand and non-brand ads are similarly effective, but defense might be slightly more effective on brand ads than non-brand ads.

While these effectiveness differences are interesting, on their own, they are not sufficient guidance for advertisers. Simonov and Hill (2021) emphasize that the metric advertisers should use to make bidding decisions is a keyword's quality-adjusted cost per incremental click. Our study comments on part of this metric (quality-adjusted incremental clicks), but does not comment on cost. We cannot comment on cost because the CPC data from Google and Microsoft's ad APIs omits a key axis of variation: the difference, for a given keyword, between focal brand and competitor CPC. This difference can be large for navigational and brand keywords, e.g., Simonov and Hill (2021) found that the most relevant competitor must increase their bid 7–8x to dislodge the focal brand from the top ad slot. Thus, future work that combines quality-adjusted incremental clicks with competitor-specific CPCs could demonstrate that competitors are overspending on navigational ads.

5.1 Limitations

Our study has several limitations. First and foremost, the time period for our CPC data (February 2023) does not match the time period for our browsing activity data (August–December 2020). In order to evaluate the impact of this discrepancy, we acquired Google CPC estimates corresponding to August–December 2020 from a marketing analytics firm named Semrush. In order to protect participants' anonymity, we only shared 5,790 Google queries that $\geq$ 2 participants searched. We reproduced our analysis on this query subset and found that revenue shares were similar whether we used our 2023 CPC data (53.0–72.4% navigational, 10.8–29.1% brand) or 2020 Semrush CPC data (39.6–73.9% navigational, 9.7–49.3% brand).[23]

Second, we relied on Bing and DuckDuckGo SERPs that we crawled in January 2023, while our Bing and DuckDuckGo queries come from late 2020. We use both the query and the SERP to label navigational and brand searches, so this introduces measurement error that is not captured in our confidence intervals. Third, we classified searches as exclusively navigational, brand, or non-brand, even though there is explicit overlap between these categories. Future work could explore a more probabilistic treatment. Relatedly, we treat shopping and text ads equivalently in our revenue calculations, even though we know advertisers bid on product attributes instead of keywords in shopping ad auctions. We made this simplification because the ad APIs only provide data about keywords. Our results are also limited to the US market and English queries; it is not clear how ad CTRs and CPCs vary across geography and language. Finally, we scoped our study to general search, thus ignoring ad clicks on vertical search engines (e.g., Google Shopping and Images), which account for 12.5%, 7.4%, and 2.7% of Google, Bing, and DuckDuckGo ad clicks, respectively.

[23]The absolute revenue share for navigational search is higher in this subset because we only shared repeated queries with Semrush, which are more likely to be navigational.

5.2 Ethics Statement

We followed standard ethical protocols throughout this IRB-approved study.[24] We informed potential participants about the data our browser extension would collect and asked for consent before collecting any data (see § 6.2). Participants were compensated and could exit the study at any time. Our extension uninstalled itself after the study concluded. We did not share any data collected by our extension with YouGov, and are not permitted to publicly release participant-level data, given its sensitivity. The total cost of this study was $78,000, which included fees to YouGov and compensation for participants. We do not anticipate any harms to participants—either as a result of being a study member or due to the publication of this study—or negative societal impacts in general.

Acknowledgements

The collection of data used in this study was funded in part by the Anti-Defamation League, the Russell Sage Foundation, and the Democracy Fund. This research was supported in part by NSF grant IIS-1910064. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders.

References

Alphabet Inc. 2022. 2021 Annual Report (Form 10-K).

Ashkan, A.; Clarke, C. L.; Agichtein, E.; and Guo, Q. 2009. Classifying and characterizing query intent. In *European conference on information retrieval*, 578–586. Springer.

Beus, J. 2021. The proportion of mobile searches is more than you think – What you need to know. Sistrix.

Blake, T.; Nosko, C.; and Tadelis, S. 2015. Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1): 155–174.

Broder, A. 2002. A taxonomy of web search. In *ACM Sigir forum*, volume 36, 3–10.

Buscher, G.; Cutrell, E.; and Morris, M. R. 2009. What Do You See When You're Surfing?: Using Eye Tracking to Predict Salient Regions of Web Pages. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, 21.

Chan, D.; Kumar, D.; Ma, S.; and Koehler, J. 2012. Impact Of Ranking Of Organic Search Results On The Incrementality Of Search Ads. Technical report, Google Inc.

Cohen, W. W.; Ravikumar, P.; Fienberg, S. E.; et al. 2003. A Comparison of String Distance Metrics for Name-Matching Tasks. In *IIWeb*, volume 3, 73–78.

Competition & Markets Authority. 2020. Online Platforms and digital advertising.

Coviello, L.; Gneezy, U.; and Goette, L. 2017. A large-scale field experiment to evaluate the effectiveness of paid search advertising.

[24]Northeastern IRB protocol #20-03-04.

Cutrell, E.; and Guan, Z. 2007. What Are You Looking for? An Eye-Tracking Study of Information Usage in Web Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 407–416.

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

Dimitrov, D.; Singer, P.; Lemmerich, F.; and Strohmaier, M. 2016. Visual Positions of Links and Clicks on Wikipedia. In *Proceedings of the 25th International Conference Companion on World Wide Web - WWW '16 Companion*, 27–28.

Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1): 242–259.

Efron, B. 1987. Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397): 171–185.

Golden, J.; and Horton, J. J. 2021. The effects of search advertising on competitors: An experiment before a merger. *Management Science*, 67(1): 342–362.

Goldman, E. 2008. American Airlines and Google Settle Keyword Advertising Lawsuit. Technology and Marketing Law Blog.

Goldman, E. 2010. Rescuecom Abandons Its Litigation Against Google. Technology and Marketing Law Blog.

Granka, L.; Hembrooke, H.; and Gay, G. 2006. Location Location Location: Viewing Patterns on WWW Pages. In *Proceedings of the 2006 Symposium on Eye Tracking Research & Applications*, 43.

Granka, L. A.; Joachims, T.; and Gay, G. 2004. Eye-Tracking Analysis of User Behavior in WWW Search. In *Proceedings of the 27th Annual International Conference on Research and Development in Information Retrieval*, 478.

Hofmann, K.; Mitra, B.; Radlinski, F.; and Shokouhi, M. 2014. An Eye-tracking Study of User Interactions with Query Auto Completion. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, 549–558.

Hotchkiss, G.; Alston, S.; and Edwards, G. 2005. Eye Tracking Study: An in Depth Look at Interactions with Google Using Eye Tracking Methodology. Technical report, Enquiro.

Jansen, B. J.; Booth, D. L.; and Spink, A. 2008. Determining the informational, navigational, and transactional intent of Web queries. *Information Processing & Management*, 44(3): 1251–1266.

Jansen, B. J.; and Spink, A. 2009. Investigating customer click through behaviour with integrated sponsored and non-sponsored results. *International Journal of Internet marketing and advertising*, 5(1): 74.

Microsoft Corp. 2022. 2021 Annual Report (Form 10-K).

Nettleton, D.; and González-Caro, C. 2012. Successful Web Searches: What Makes the Difference? An Eye-Tracking Study. *Revista Colombiana de Computación*, 13(2): 53–72.

Pan, B.; Hembrooke, H.; Joachims, T.; Lorigo, L.; Gay, G.; and Granka, L. 2007. In Google We Trust: Users' Decisions on Rank, Position, and Relevance. *Journal of Computer-Mediated Communication*, 12(3): 801–823.

Papoutsaki, A.; Laskey, J.; and Huang, J. 2017. Searchgazer: Webcam eye tracking for remote studies of web search. In *Proceedings of the Conference on Conference Human Information Interaction and Retrieval*.

Reuters. 2012. Rosetta Stone and Google settle trademark lawsuit. Reuters.

Robertson, R. E.; and Wilson, C. 2020. WebSearcher: Tools for Auditing Web Search. In *Proceedings of Computation + Journalism Symposium*.

Rose, D. E.; and Levinson, D. 2004. Understanding user goals in web search. In *Proceedings of the 13th international conference on World Wide Web*, 13–19.

Sayedi, A.; Jerath, K.; and Srinivasan, K. 2012. Competitive Poaching in Sponsored Search Advertising and Its Strategic Impact on Traditional Advertising. *Marketing Science*, 33(4).

Schultz, C. D. 2020. Informational, transactional, and navigational need of information: relevance of search intention in search engine advertising. *Information Retrieval Journal*, 23(2): 117–135.

Simonov, A.; and Hill, S. 2021. Competitive Advertising on Brand Search: Traffic Stealing and Click Quality. *Marketing Science*, 40(5): 923–945.

Simonov, A.; Nosko, C.; and Rao, J. M. 2018. Competition and crowd-out for brand keywords in sponsored search. *Marketing Science*, 37(2): 200–215.

Singh, M. 2023. Google required to remove ads that violate trademarks, Indian court rules.

Statista. 2022. Paid mobile search click share of Google and Microsoft Advertising in the United States as of 1st quarter 2021, by device.

Stern, T. 2017. Refreshing competitive search strategies in 2017. Search Engine Land.

Subcommittee on Antitrust, Commercial and Administrative Law of the Committee on the Judiciary. 2020. Investigation of Competition in Digital Markets.

Teevan, J.; Liebling, D. J.; and Geetha, G. R. 2011. Understanding and predicting personal navigation. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 85–94.

Ursu, R.; Simonov, A.; and An, E. 2021. Online advertising as passive search. *NYU Stern School of Business Forthcoming, Columbia Business School Research Paper*.

Vallina, P.; Le Pochat, V.; Feal, Á.; Paraschiv, M.; Gamba, J.; Burke, T.; Hohlfeld, O.; Tapiador, J.; and Vallina-Rodriguez, N. 2020. Mis-shapes, mistakes, misfits: An analysis of domain classification services. In *Proceedings of the ACM Internet Measurement Conference*, 598–618.

6 Appendix

6.1 Ethics Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes, see § 3.
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes, see § 3.1.
   (e) Did you describe the limitations of your work? Yes, see § 5.1.
   (f) Did you discuss any potential negative societal impacts of your work? Yes, see § 5.2.
   (g) Did you discuss any potential misuse of your work? NA
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, see § 5.2.
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? Yes, see § 4.3.
   (b) Have you provided justifications for all theoretical results? Yes, see § 4.3.
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? NA
   (e) Did you address potential biases or limitations in your theoretical framework? NA
   (f) Have you related your theoretical results to the existing literature in social science? Yes, to Simonov and Hill (2021), see § 4.3.
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? Yes, see § 5.
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? NA
   (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   (a) If your work uses existing assets, did you cite the creators? NA
   (b) Did you mention the license of the assets? NA
   (c) Did you include any new assets in the supplemental material or as a URL? NA
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes, see § 5.2.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, see § 5.2.
   (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA
   (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
   (a) Did you include the full text of instructions given to participants and screenshots? Yes, see § 5.2 and § 6.2.
   (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? Yes, see § 5.2.
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? Yes, see § 5.2. Note that YouGov handles participant compensation, thus we do not have hourly wage estimates.
   (d) Did you discuss how data is stored, shared, and deidentified? Yes, see § 5.2.

6.2 Browser Extension Informed Consent

Welcome to the study!

This extension implements a user study being conducted by researchers at Northeastern University, Dartmouth, Princeton, and University of Exeter. If you choose to participate, this browser extension will confidentially collect four types of data from your browser.

1. Metadata for web browsing (e.g., URL visited with time of visit), exposure to embedded URLs on websites (e.g., YouTube videos), and interactions with websites (e.g., clicks and video viewing time). This data is collected until the study is completed.

2. Copies of the HTML seen on specific sites: Google Search, Google News, YouTube, Facebook Newsfeed, and Twitter Feed. We remove all identifying information before it leaves the browser. This confidential data is collected until the study is completed.

3. Browsing history, Google and YouTube account histories (e.g., searches, comments, clicks), and online advertising preferences (Google, Bluekai, Facebook). This data is initially collected for the year prior to the installation of our browser extension, and we then check these sources once every two weeks to collect updates until the study is completed.

4. Snapshots of selected URLs from your browser. For each URL, the extension saves a copy of the HTML that renders, effectively capturing what you would have seen had you visited that website yourself. Once per week we conduct searches on Google Search, Google News, YouTube, and Twitter, and collect the current frontpage of Google News, YouTube, and Twitter. These web page visits will occur in the background and will not affect the normal functioning of your browser. There is a theoretical risk of "profile pollution" – that this extension will impact your online profiles, i.e., "pollute" them with actions that you did not take. To mitigate this risk, the extension will only visit content that is benign and will only execute searches for general terms. Our previous work has found that historical information of this kind has minimal impact on online services.

Additionally, if you choose to participate, you will be asked to take a survey in which we ask you several questions about your demographics, web usage, and media preferences. These data, as well as those mentioned above, will be used to analyze the correlations between your online behavior and your interest profiles.

After the study is complete on December 31, 2020, the extension will uninstall itself. All data collected will be kept strictly confidential and used for research purposes only. We will not share your responses with anyone who is not involved in this research.

You must be at least 18 years old to take part in this study. The decision to participate in this research project is voluntary. You do not have to participant and you can refuse to participate. Even if you begin our experiment, you can stop at any time. You may request that we delete all data collected from your web browser at any time.

We have minimized the risks. We are collecting basic demographic information, information about your internet habits, and copies of web pages that you visit. To the greatest extent possible, information that identifies you will be removed from all collected web data.

Your role in this study is confidential. However, because of the nature of electronic systems, it is possible, though unlikely, that respondents could be identified by some electronic record associated with the response. Neither the researchers nor anyone involved with this study will be collecting those data. Any reports or publications based on this research will use only aggregate data and will not identify you or any individual as being affiliated with this project.