From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising *

Francesco Decarolis Gabriele Rovigatti

May 8, 2020

Abstract

This paper analyzes the impact of intermediary concentration on the allocation of revenues in online platforms. We study *sponsored search* documenting how advertisers increasingly bid through a handful of specialized intermediaries. This enhances automated bidding and data pooling, but lessens competition whenever the intermediary represents competing advertisers. Using data on nearly 40 million Google keyword-auctions, we first apply machine learning algorithms to cluster keywords into thematic groups serving as relevant markets. Using an instrumental variable strategy, we estimate a decline in the platform's revenues of approximately 11 percent due to the average rise in concentration associated with intermediary M&A activity.

Keywords: Buyer power, concentration, online advertising, platforms, sponsored search

JEL Classification: C72, D44, L81.

*Francesco Decarolis: Bocconi University, Department of Economics and IGIER. Via Roentgen 1, Milan, Italy. E-mail address: *francesco.decarolis@unibocconi.it.* Gabriele Rovigatti: Bank of Italy, Research Department. Via Nazionale 91, Rome, Italy. E-mail address: *gabriele.rovigatti@bancaditalia.it.* The authors would also like to thank for their many comments colleagues at their institutions as well as seminar participants at the EU, UK and US competition authorities (DG Comp, CMA and FTC), at university departmental seminars (BU Questrom, MIT Sloan, Toulouse School of Economics, Paris School of Economics, Tilburg and McGill) and conferences (NBER, MaCCI, NET Institute, IIOC, BECCLE) where earlier drafts of this study were presented. *The use of the word "collusion" in this essay is unrelated to any assessment of the legal implications of agencies, agency networks or advertisers behavior under the competition laws of the US or other countries.*

"....*Essentially, we are investment managers for our clients, advising them how to spend around $90 billion of media. So it makes sense that WPP should offer platforms that are agnostic, and help clients plan and buy media. To that end, we are applying more and more technology to our business, along with big data. We are now Maths Men as well as Mad Men (and Women). Thus we go head to head not only with advertising and market research groups such as Omnicom, IPG, Publicis, Dentsu, Havas, Nielsen, Ipsos and GfK, but also new technology companies such as Google, Facebook, Twitter, Apple and Amazon...*"
(Sir Martin Sorrell, WPP founder and former CEO, WPP's 2012 Annual Report)

I Introduction

Advertising in the internet era is about capturing the attention of consumers browsing the web. This requires both detailed data to effectively *target* the ad at the right customers and fast algorithms to bid in the online auctions where ad space is sold. These needs have led to a major, but understudied, shift in the industry: rather than bidding individually, advertisers increasingly delegate their bidding to highly specialized intermediaries. This concentration of demand within a few large intermediaries raises the question of countervailing buyer power. Can the emergence of intermediaries counterbalance the highly concentrated supply of online ads? The supply of search advertising is indeed mostly in the hands of a single firm: Google. Google's dominance of search advertising has put it at the center of the current debate about the power and influence of the web giants.[1]

This study contributes to two key issues in the debate about concentration by studying the market for sponsored search ads. The first is quantification of demand concentration, with an emphasis on the proper use of data to identify markets. Starting from granular keyword-level data, we propose a machine learning method based on state of the art classification algorithms to cluster individual keywords into markets. The second regards the effects of concentration. We focus on the idea of countervailing power – how downstream buyers respond to the high concentration in the upstream industry. Galbraith [1952] famously remarked that "the best and established answer to economic power is the building of countervailing power: the trade union remains an equalizing force in the labor markets, and the chain store is the best answer to the market power of big food companies." Are advertising intermediaries a remedy to the dominance of Google?

The market for sponsored search ads represents about half of all internet advertising revenues, or about $40 billion dollars in 2017 [IAB, 2018]. This market is highly concentrated, with Google earning 75-80 percent of total search advertising revenue in the United States between 2016-18 [eMarketer, 2018]. Advertisers are the demand side. They seek the attention of potential customers searching the web. To do so, they compete

[1]See, among others, Autor et al. [2017], De Loecker and Eeckhout [2020], Werden and Froeb [2018], Gutierrez and Philippon [2017] and Weche and Wambach [2018], as well as the Obama administration's CEA [2016] and the press coverage by Economist [Economist, 2016*a*,*b*] and Guardian [Stiglitz, 2016].

against other advertisers to buy one of a limited number of 'slots' available on the search engine's result page for a given search term or *keyword*.[2] In the early days of this market, advertisers used to operate individually but over time, more and more ad buying is conducted through intermediaries. In our data, intermediaries are involved in about 75 percent of the slots sold.

To understand why this shift can have profound implications for the allocation of revenues, three features need to be considered simultaneously: i) search engines use auctions to sell slots; ii) advertisers run marketing campaigns through digital marketing agencies (henceforth, MAs) to which they give mandates to bid on their behalf in auctions; iii) most MAs further delegate bidding to specialized entities which, for those MAs belonging to an agency network (henceforth *network*),[3] are the network's agency trading desks. While thousands of MAs operate in the market, essentially only seven networks are responsible for collecting data and optimizing bidding algorithms for most advertisers. For the search engine, their specialized activity can be useful in terms of economies of scale and scope for the use of data, as well as for the implementation of automated bidding. However, it can also trigger major revenue losses if the network lessens competition among its clients bidding in the keyword auctions pertaining to the same market.

In the 2014-2017 period covered by our data, the market share of the four largest networks was approximately 70 percent of search volumes and was growing over time. Our analysis quantifies the extent to which such increases in concentration affect Google's revenues. Our strategy is based on three ingredients: a novel dataset on keyword ads, advertisers and intermediaries; a method to define the relevant markets; and an identification strategy to assess the causal impact of demand concentration on search engine revenues.

The first ingredient is a novel dataset built by combining multiple sources. We have obtained from Redbooks – the most comprehensive database on marketing agencies – the list of the 6,000 largest US online advertisers. For these advertisers, we have the MAs they are affiliated with, together with the network each individual agency belongs to. We combine this with data on Google's sponsored search auctions from SEMrush, a major data provider for MAs. For all Redbooks advertisers, we know which keywords, if any, they bid on in Google's auctions. For each keyword and year, we know the position of the domain (and consequently of the advertiser) in the search results page, the volume of searches (i.e., the average number of search queries for the given keyword in the last 12 months) and the keyword-specific average price advertisers paid for a user's click on the ad (Cost-Per-Click, or CPC).

The second ingredient is market definition. We use natural language processing to move from the 23 industries provided by Redbooks to more granular clusters of keywords representing individual markets. The approach involves a 2-layer clustering procedure: keywords are initially split into thematic clusters on the basis of their semantic content (via the *GloVe* algorithm of Pennington, Socher and Manning [2014])

[2]In the advertising industry, all the activities involved in the process of gaining website traffic by purchasing ads on search engines are indicated as SEM (Search Engine Marketing).

[3]The seven networks are IPG, WPP, Publicis Groupe, Omnicom Group, Dentsu-Aegis, Havas and MDC.

and then each thematic cluster is further partitioned using a similarity measure based on the advertisers' co-occurrency across keywords. Although not in a strict antitrust sense, we can treat these latter groups as relevant markets. They contain keywords closely connected in terms of both consumer perceptions and advertiser competition: the consumer side is captured in the first layer, where the algorithm is trained over 840 billion documents in a way that resembles how consumers learn about products from the web, while advertiser side is captured in the second layer.

The third ingredient is an instrumental variable strategy. Instruments are needed for two reasons: measurement error in the proxy for demand concentration and potential omitted variable bias. For instance, there might be unobservable shocks to the popularity of some keywords that drive changes in both revenues and demand concentration. Similar to Dafny, Duggan and Ramanarayanan [2012], we address this problem by exploiting the variation in intermediary concentration driven by changes in network ownership of MAs. In our sample period, there were 21 acquisitions and 2 divestures, affecting 6 out of the 7 agency networks. These M&A operations, especially the larger ones involving a multiplicity of markets, are a useful source of variation in demand concentration as the revenue dynamics in each local market are too small by themselves to cause the M&A operations. We discuss in detail this strategy and evaluate its robustness.

We find with both OLS and IV estimates that greater network concentration induces lower search engine revenues. Under our baseline IV model, a change in the HHI of 245.10 points – the average HHI increase observed across the markets experiencing a merger event – leads to a decrease of 11.32 percent in revenues. This is robust for several sensitivity analyses, including different market definitions and model specifications.

We complement the baseline analysis with a number of exercises aimed at identifying the main drivers of the changes in revenues. On the one hand, the decline in revenues appears to be driven by a decline in the price of the chosen keywords: demand concentration is negatively associated with the average cost-per-click, but not with the number of keywords or their associated search volumes. This result provides direct evidence of the incentive to ease price competition, which we relate to both algorithm capabilities (in terms of coordinating bids and targeting for the same keywords and segmenting the markets across keywords) and network bargaining strength. It also complements in an interesting way the detailed descriptive evidence that we present on the differences between bidding by intermediaries and by individual advertisers, where we show that the former tends to select keywords that are more specialized, more precisely targeted and, as a result, less expensive on average. On the other hand, we also find a high degree of heterogeneity among industries, which reflects the variety of strategies followed by intermediaries in implementing ad campaigns. Overall, despite the potential efficiencies created by intermediation in certain markets, the empirical evidence indicates that increasing buyer power results in a worsening of the search engine's revenue prospects.

Our findings contribute to at least three branches of the literature. At the most general level, this paper contributes to the ongoing debate on concentration (see references in footnote 1) and its results highlight two key aspects. One is the well known problem of the inadequacy of industry-level data to analyze concentration and its effects. In our setting, this emerges as a marked difference between industry-level and market-level estimates. The other aspect, also well known, is that the profitability of even the most concentrated industry crucially depends on other forces, among which is the degree of buyer power. Our analysis illustrates how the market power of Google has been partially eroded by technological innovations and concentration among buyers. Thus, the old idea of countervailing power still matters for one of the most modern and dynamic industries.[4] In contrast to the existing buyer power literature which is mostly centered around bargaining models, the fact that auctions are the selling mechanism makes our study close in spirit to the classic work of Snyder [1996] on auctions. Specifically for online markets, competition authorities have identified concerns about the imbalance of bargaining power in favor of platforms [EU Commission, 2017], and these concerns have led to buyer concentration being identified as a possible solution to this imbalance [Mullan and Timan, 2018]. This study is the first to quantify the effectiveness of this remedy. In the conclusions, we discuss the pros and cons of buyer power relative to alternative policy interventions that are currently being debated. We also discuss a handful of recent trends in the market, from disintermediation to changes in auction reserve prices, that can be interpreted as a response from Google to the increased strength of intermediaries.

Second, the paper contributes to the understanding of online advertising. This is a particularly complex, economically relevant and rapidly evolving market. The existing studies on online ads have mostly focused on their effectiveness (see Goldfarb [2014], Blake, Nosko and Tadelis [2015], Golden and Horton [2018] and Johnson, Lewis and Reiley [2017]), the functioning of the selling mechanisms (see Edelman, Ostrovsky and Schwarz [2007], Varian [2007], Athey and Nekipelov [2014], Borgers et al. [2013], Balseiro, Besbes and Weintraub [2015] and Celis et al. [2015]) or platform competition and consumer welfare (see Prat and Valletti [2019]). By focusing on the role of intermediaries, our study offers new insights on the firms that have practically taken over modern advertising markets, but whose role is not yet fully understood. In fact, we complement a small number of recent studies that have looked at these players (see their review in Choi et al. [2019]). These works mostly emphasize the positive roles of intermediaries in improving the use of information, to limit winners' curse risks (McAfee [2011]), and in more effectively administering client budgets, to avoid the inefficiencies associated with budget constrained bidders (Balseiro and Candogan [2017]). To these positive effects, our analysis adds an emphasis on the perils of coordinated bidding for the search engine. Several theoretical studies had highlighted the vulnerability of online ad auctions to collusive bidding through common intermediaries (Bachrach [2010], Mansour, Muthukrishnan and Nisan [2012] and

[4]There are many examples of countervailing buyer power across different industries. For instance, in the case of US healthcare, the insurers' introduction of HMOs and PPOs is credited to have dramatically rebalanced power in favor of insurers after decades of increased hospital concentration [Gaynor and Town, 2012]. See also the related literature on hospitals' consolidation [Craig, Grennan and Swanson, 2018; Dafny, Ho and Lee, 2019; Gowrisankaran, Nevo and Town, 2015; Schmitt, 2017]. For empirical applications in different industries see [Chipty and Snyder, 1999], [Villas-Boas, 2007], [Ellison and Snyder, 2010] and the UK Competition Commission's Final Report of the Grocery Market Investigation.

Decarolis, Goldmanis and Penta [2019]). While only the latter is directly applicable to search – the former two being on VCG and second price auctions – all three studies share a narrow focus on bidding equilibria within a one-shot auction. Our empirical analysis differs by looking more broadly at how the market works. This allows us to describe multiple mechanisms through which intermediaries lower auction prices involving not only bid coordination within a keyword, but also ad targeting and the choice of keywords.

Third, this study develops a novel measure of market definition based on machine learning algorithms. Our study is thus related to Hoberg and Phillips [2016] who pioneered this approach by employing a systematic text-based analysis of firm 10-K product descriptions to construct product similarity measures. Relative to that study, our clustering approach uses a different algorithm and can be implemented for all firms bidding in search auctions, regardless of whether they file 10-K forms or not.

Finally, it should be remarked that while we will at various points use the term collusion, the behaviors that we describe below are proper from a legal perspective. They would not constitute a violation of the antitrust laws in the US or the EU because intermediaries are legal entities, independent from advertisers, operating unilaterally to maximize their profits. As such, they can freely decide how to arrange pricing algorithms that bid on behalf of their customers.[5] Toward the end of the paper, we briefly discuss how some of these algorithms are built: the most advanced ones entail forms of learning, but most of the benefits from coordination can be achieved (and appear to be achieved) by simple, non-AI based algorithms, that exploit ad targeting to segment markets. Nevertheless, the shift toward AI-powered algorithms that this industry is experiencing has the potential to exacerbate the dynamics that we analyze below, for the same motives that are stressed in recent research on algorithmic collusion [OECD, 2017; Calvano et al., 2018].

II Basic Framework

Consider a monopolist search engine selling ad slots on its results page. Consider also three advertisers – q, j and k – interested in showing their ad to consumers searching for a keyword. Allocations and payments depend on how many ad slots the search engine places on its web page and on the selling mechanism adopted. For instance, with one available slot sold through a second price auction, the winner will be the advertiser with the highest bid and his payment will equal the second highest bid.

[5] A caveat to this claim is that two situations might imply an antitrust infringement. The first case involves the discipline on "hub and spoke" cartels, Harrington [2018], which would apply if it could be proved that advertisers had agreed to delegate their bidding to a common intermediary with the explicit intent of enforcing price coordination or market splits. The second case involves the discipline on Purchasing Agreements, or Group Purchasing Organizations (GPO). Although the intermediaries that we study are not GPO, under the EU law, the limits to the activity of GPO may be invoked if one could show that an intermediary controls such a large share of the market that its coordination activity could hurt Google's revenues to the point of leading to a worsening of the quality of the Google' services.

Now suppose that advertisers do not bid directly on the search auction. They submit their bid to an intermediary who internally runs a second price auction amongst its clients (we shall refer to this as the *intermediary auction*) and then bids on their behalf in the search auction. To see why this can affect the functioning of the search auction, consider the two cases illustrated in Figure 1. In panel A, each advertiser bids through a different intermediary, which we indicate as α, β and γ. In this case, intermediaries have no incentive to distort bids in the search auction. Hence, if for instance the bids placed in the intermediary auction are $b_q = 4$, $b_j = 3$ and $b_k = 1$, the same bids will enter the search auction: $b_{q,\alpha} = 4$, $b_{j,\beta} = 3$ and $b_{k,\gamma} = 1$, as indicated by the straight arrows. Advertiser q wins the slot and pays 3 to the search engine.

Figure 1: An Example of Bidding through Intermediaries

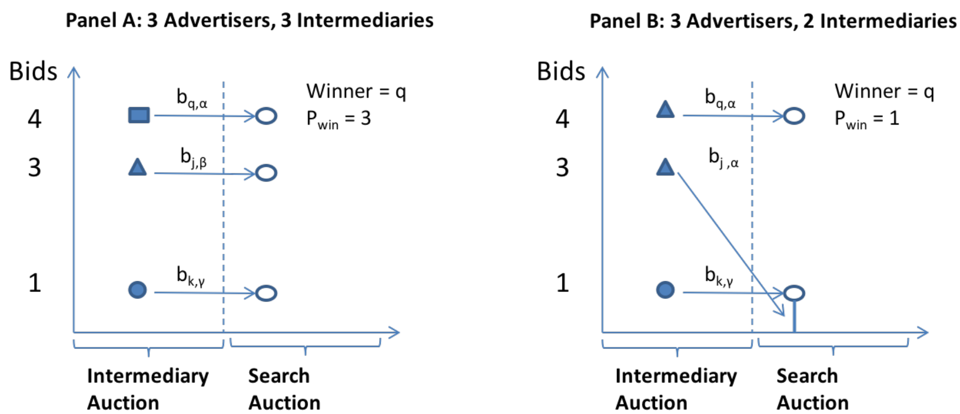

Notes: There are three advertisers (q, j and k) submitting arbitrary bids ($b_q = 4$, $b_j = 3$ and $b_k = 1$) to a second price auction held by the intermediary (α, β or γ) to which they are affiliated. Intermediaries then bid in the search auction. In panel A, each advertiser has a different intermediary. In panel B, q and j share intermediary α. The arrows indicate how the intermediary translates the bids in its own auction into the bids placed on the search auction. In panel A, bids are passed without distortions; in panel B, j's bid is reduced. q wins in both cases, paying the second highest bid which is either 3 (panel A) or 1 (panel B).

In panel B, we plot the same game, but with 2 intermediaries: both q and j bid through intermediary α. This intermediary can now alter the search auction outcomes by retaining or amending the bids it places on behalf of its two clients: it can report just the highest bid among the two, $b_{q,\alpha} = 4$, or both bids, but setting $b_{j,\alpha} \in [0, 1]$. In all these cases, q wins the slot and pays only 1 instead of 3, thanks to the reduction in $b_{j,\alpha}$ implemented by the intermediary α. This example provides intuition on how intermediary concentration may lower the CPC in an auction, and consequently the search engine's revenues. Implementing this in practice would not be so simple for an intermediary handling many advertisers active over thousands of keywords, each with its own competitive structure dynamically evolving as Google updates the quality scores used to reweigh bids as the set of rivals changes. Although algorithms for bid coordination in search have been proposed,[6]

[6]We will mention below an algorithm performing this type of within-auction bid coordination developed by engineers at a leading MA, iProspect. Decarolis, Goldmanis and Penta [2019] also propose three different

keyword multiplicity allows for a simpler form of coordination: market split by keywords. Consider a modification to the example above so that there are three branded keywords and the three advertisers can bid only on their own branded keyword (i.e., the keyword including their own brand name) or also on those of their rivals. As in a typical prisoner's dilemma, all advertisers might be better off bidding only on their own brand, but – absent coordination – they bid on rival brands too. Explicit coordination by advertisers would be outright collusion and, so, illegal.[7] Delegation to a common intermediary that autonomously implements the market split is a solution to the dilemma that would not be in breach of the law.

Given the nearly infinite combination of keywords, there is ample scope for this type of strategy. Furthermore, targeting allows forms of market split to occur even for the same keyword. Consider a simple algorithm exploiting targeting on two dimensions: Google Ads allows geographical targeting (down to the zip code level) and schedule targeting (down to 15 minutes past the hour intervals). For most keywords, however, the traffic volume is not so finely differentiated. This means that an algorithm that rotates bids between two advertisers so that they never meet in these zip code / quarter of hour intervals would expose these advertisers to the same audience, but without making them compete. Considering that other targeting parameters are usually feasible, the number of possible market segmentations is nearly infinite.

What all of the above strategies have in common is that they would induce a lower CPC and, through that, lower revenues for the search engine. This negative effect, however, need not be the final outcome of intermediary concentration. Indeed, intermediaries can foster entry, by bringing to these auctions advertisers who would otherwise not enter. Moreover, thanks to their superior bidding technology, they can also increase the number of keywords on which advertisers bid. Furthermore, the literature has identified the presence of externalities that agencies could help to internalize, thus improving their appeal for advertisers: for a given keyword-advertiser-slot, the number of clicks that the advertiser receives under different configurations of the set of rivals displayed might be very different ([Jeziorski and Segal, 2015; Gomes, Immorlica and Markakis, 2009]). In the closely related context of ad exchanges, the literature has identified further problems related to limited information leading to winners' curse [McAfee, 2011] and budget constraints leading to inefficiencies [Balseiro and Candogan, 2017] that a common intermediary might help to solve. Next we illustrate industry facts that are helpful in understanding this nuanced relationship between intermediaries and search engines.

III Industry Background

The premise to what follows is acknowledging the impossibility of fully describing a heterogeneous and rapidly evolving industry. The lack of public regulation also means that both search engines and networks are essen-

algorithms of this kind, depending on the exact objective function of the MA.

[7]In 2019, the FTC charged 1-800 Contacts inc. for having entered into bidding agreements with at least 14 competing online contact lens retailers that eliminated competition on branded keywords search advertising.

tially free to arrange their contracts and methods of working in an unconstrained and often non-transparent way. Nevertheless, below we offer a schematic account incorporating our best knowledge, developed from interactions with market participants and direct involvement as both bidders and intermediary clients.

Internet advertising is mostly split between sponsored search and display advertising. Our study focuses on the former. In essence, an advertiser opens an account on the platform auctioning off 'slots' on the search engine results page (for instance, *Google Ads*, formerly *AdWords*) and enters a bid amount, a budget and a brief ad for all the keywords of interest. Each time a user queries the search engine for one of these keywords, an auction is run to allocate the available slots (typically up to seven) among the interested advertisers. The slot order reflects the ranking emerging from the bids (reweighted by ad quality in the case of Google and of some other platforms like Bing), and payment occurs only if the user clicks on the ad.

In recent years, the market has become more complex. On the supply side, even though search advertising remains firmly dominated by Google, new players have emerged to challenge Google (mostly Amazon). For the other types of internet ads, video (mostly YouTube) and social media (mostly Facebook) have grown more relevant and display advertising has seen the emergence of new ad exchanges. On the demand side, a major evolution is the advent of intermediaries: advertisers are supported by intermediaries in bidding across the many platforms where online ad space is traded.[8] There are many types of intermediaries. In search, advertisers typically contract out marketing campaigns to a digital marketing agency (MA) that designs and optimizes the campaign. Thousands of MAs operate in the US market, but most of them belong to one of the seven agency networks, the MA holding companies. Within networks, bidding typically happens through centralized entities, the so called "agency trading desks" (ATDs). ATDs represent the demand-side technological response to the incentive to improve bidding performance through better data and faster algorithms. Compared to the independent MAs, they represent a heterogenous pool with some more involved with the creative part of ads and using third parties to manage the bidding, while others develop their own bidding systems internally. An interesting feature of the intermediary sector relevant for our analysis is the acquisition of these independent MAs by the networks. As we shall see in section VI, over the sample period analyzed, 21 MAs were acquired, with 6 different networks involved.

Finally, regarding the bidding algorithms, automated systems for bidding is what crucially characterizes network bidding: they allow faster bidding and the implementation of data-intensive targeting strategies with limited human intervention. Bidding algorithms that perform these tasks often use machine learning, while the diffusion of AI algorithms with learning abilities is still in its infancy. Nevertheless, all the industry experts that we consulted (see appendix) agree that AI-based bidding is rapidly penetrating this sector.[9] The algorithms used differ depending on the environment in which they work, such as search vs. display.

[8] In some platforms, *only* qualified bidders and intermediaries can bid – like in the financial exchanges.

[9] An example of an AI system that fully automatizes the handling of marketing campaigns is Koa, `https://www.thetradedesk.com/products/koa-artificial-intelligence`. Google itself offers a suite of automated bidding options which are based on learning algorithms.

Regardless of algorithm specificities, their common goal is to allow networks to effectively use large amounts of consumer profiling data, both internal and purchased from third parties (*data exchanges*). Pooling all advertiser data into a single network allows immediate access to relevant data and can both save on the costs charged by the *data exchanges* and improve speed. For our analysis, the presence of ATDs implies that the most relevant level at which we should analyze intermediary concentration is that of agency networks.

IV Data

The dataset that we use allows us to observe a large set of keyword auctions, the advertisers bidding on them and their intermediaries, both at the MA and at the network level (when applicable). Indeed, the minimal data requirements to test the effects of intermediary concentration on the search engine's revenues are information on: *i)* the advertisers' affiliation to intermediaries, *ii)* the set of keywords on which they bid and *iii)* the associated average CPC and search volume of these keywords. Our new dataset contains all this information, and more. Figure 2 shows the hierarchical structure of the data: the highest level (the networks, for non-independent agencies) group the individual MAs (layer 2). These, in turn, cluster the advertisers (layer 3), each bidding on a different set of keywords (layer 4). Solid lines indicate the cases of coalitions: in Figure 2, for example, MA 2 participates in the auction for *key5* on behalf of both Advertiser 3 and Advertiser 4. But we also consider *key2* as having a coalition because Advertiser 1 and Advertiser 4 both bid through Network 1, although via different MAs.

Figure 2: Redbooks-SEMrush Data Structure

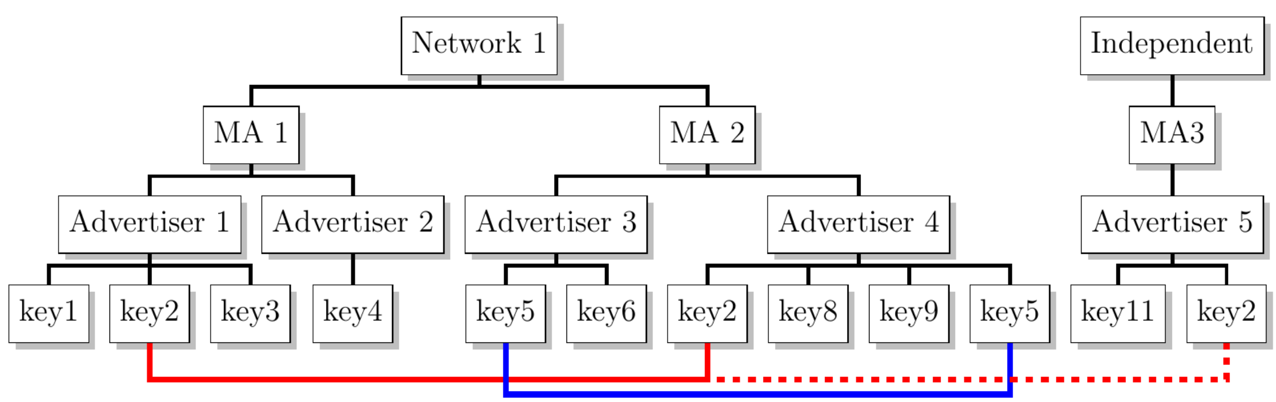

Notes: Hierarchical structure of the data. From bottom to top: keywords (SEMrush), advertisers (Redbooks/SEMrush), agencies and networks (Redbooks). Solid lines represent examples of coalitions: within MA (blue) and network (red).

From Redbooks, a comprehensive database on marketing agencies, we obtained a list of advertisers representing all the major US firms active in online marketing (see Dai [2014] for an application of these data to the pharmaceutical sector and for a review of other studies using Redbooks data). For each of these advertisers, the Redbooks data give us the full list of MAs. The data are yearly for the period 2014-2017 and covers around 6,000 advertisers (i.e., web domains) per year active in all sectors of the economy. Each advertiser is associated with one of the 23 industries in which Redbooks classifies advertisers. We also have access to a linkage variable that relates each individual MA to its agency network, if any. Overall, there are seven networks and about a thousand independent agencies.[10]

We combine the data on intermediaries with sponsored search data from the most comprehensive provider of digital ad data, *www.semrush.com* (SEMrush henceforth). For keywords searched on Google, it collects the identity and website of advertisers appearing in the sponsored ad slots. Moreover, it gathers information on the keyword-specific average CPC, the position of the ad in the search outcome page, the volume of searches associated with the keyword; the visible URL of the ad; the traffic (that is, the share of the advertiser's traffic associated with the specific keyword); and the organic links. Thanks to the visible URL and the advertiser name, we are able to link Redbooks and SEMrush data. Although the SEMrush data is available at a relatively high frequency (up to daily for certain keywords), we use the yearly average to match the frequency in the Redbooks data. CPC, volume and traffic are monthly averages, calculated over the past 12 months.[11] Although these averages are calculated through proprietary algorithms that we could not inspect, they are considered reliable (and widely used) by MAs and individual advertisers (see the appendix for a more extensive discussion of the data). Whilst the use of yearly averages implies foregoing some of the richness in the geographic and time dynamics in keyword bidding, this is necessary to match the two data sources and it is adequate to address our research question involving aggregate impacts at the level of markets (i.e., groups of keywords, as discussed below).

Table 1 presents summary statistics, by keyword and advertiser type (Panel A) and by network (Panel B). In the left columns of Panel A, we report the statistics for keywords with at least one network advertiser; in the right columns, those for keywords with at least one independent advertiser (i.e., an advertiser bidding either autonomously without any MA or through a MA not affiliated to any network). The two groups are thus not mutually exclusive. For both of them, we see a similar CPC. In terms of volume, for both groups the substantially lower value of the median relative to the mean indicates a tendency to bid on keywords that are infrequently searched. The lower value of *Traffic* (1 percent) observed for the network advertisers relative to the 6 percent for the non network advertisers is compatible with the former placing ads over more keywords. *Coalition* measures the cases of keywords where more than one of the ads shown belongs to different advertisers represented by the same agency network. Within this subset of cases, *Coalition Size* shows that the average number of advertisers bidding in a coalition is 2.38 and, indeed, the vast majority

[10]Some advertisers are affiliated to multiple MAs. With very few exceptions – that we drop from the analysis sample – these do not represent an issue, since all the involved MAs belong to the same network.

[11]Since the Redbooks data are updated each year around mid January, we downloaded the SEMrush yearly data using as reference day January 15^{th} (or the closest day on record for that keyword).

Table 1: Summary Statistics: Keywords, Networks and Markets

Panel A. Statistics by Keywords and Advertiser Type

	Keywords with at Least 1 Network Years 2014-2017				Keywords with at Least 1 Independent Years 2012-2017			
	Mean	Median	SD	Obs	Mean	Median	SD	Obs
Cost-per-click	2.34	0.90	5.79	15,383,769	2.39	0.89	6.11	21,525,056
Volume (000)	497	40	34,916	15,383,769	362	40	99,845	21,525,056
Traffic	0.01	0.00	0.53	15,383,769	0.06	0.00	1.27	21,525,054
Num of Advertisers	1.30	1.00	0.68	15,383,769	1.21	1.00	0.52	21,525,056
Organic Results	4.70	0.18	26	15,383,769	3.8	0.16	19	21,525,056
# Characters	22.79	22.00	7.74	15,383,769	22.86	22.00	7.59	21,525,056
# Words	3.71	4.00	1.35	15,383,769	3.66	3.00	1.30	21,525,056
Long Tail	0.50	1.00	0.50	15,383,769	0.48	0.00	0.50	21,525,056
Branded	0.10	0.00	0.29	15,383,769	0.07	0.00	0.25	21,525,056
Coalition	0.15	0.00	0.36	15,383,769	0.00	0.00	0.00	21,525,056
Coalition Size	2.38	2.00	0.69	332,017	-	-	-	-

Panel B. Statistics by Network

	Market Share (Search Volume Share)				Presence Across Keywords			
	2014	2015	2016	2017	2014	2015	2016	2017
IPG	0.21	0.19	0.21	0.19	0.26	0.32	0.33	0.38
WPP	0.17	0.20	0.16	0.23	0.29	0.29	0.33	0.43
Omnicom	0.17	0.16	0.17	0.14	0.39	0.38	0.37	0.38
Publicis	0.14	0.13	0.13	0.18	0.30	0.30	0.29	0.30
MDC	0.09	0.09	0.08	0.09	0.17	0.17	0.17	0.24
Havas	0.05	0.07	0.06	0.02	0.12	0.14	0.12	0.06
Dentsu-Aegis	0.05	0.08	0.10	0.09	0.14	0.17	0.19	0.25
Indep. Agency	0.13	0.09	0.08	0.06	0.42	0.38	0.35	0.22

Notes: Panel A: statistics at the keyword level, separately for keywords where at least one ad comes from either a network bid (columns 1 to 4) or a non-network MA bid (columns 5 to 8). The variables included are: *Cost-per-click*, reported in USD; *Coalition*, an indicator function for the presence of multiple advertisers affiliated with the same network participating to the keyword auction; *Coalition size*, which is populated for keywords with coalitions only. Both *Long Tail* and *Branded* refer to the type of keyword: the first indicates those composed by at least four terms and the latter those including the name of a brand. *Organic results* report the number of non-sponsored search results returned by the search engine (rescaled by 10 million). Panel B: on the left (columns 1 to 4) we report for the years 2014-2017 the market share (in terms of *Search Volume*) of the seven network and non-network MAs; on the right (columns 5 to 8), we report the presence of the networks across all keywords in our data – the sum of these values within columns does not add up to one since the same keyword can display ad from multiple network and non-network bidders.

of cases involve coalitions of size 2. In essentially all cases where there is a coalition, there is exactly one coalition, suggesting that different networks are specialized in different segments of the keyword markets.

Panel B shows the relative size of each of the seven networks, both in terms of the volume of searches covered and in terms of their presence across keywords. If we consider just the largest four networks – the "big four" as they are often referred to (WPP, Omnicom, Publicis Groupe, and Interpublic Group of

Companies) – their combined market share (in terms of search volume) reaches 74 percent of the total volume in 2017. The situation is similar across years and concentration tends to increase over time. The situation is also similar if we look at the network presence across keywords.

The sheer prominence of networks in the data, together with bidding centralization at the network ATD level confirms the usefulness of considering networks as the key players in our analysis. However, measuring network concentration at keyword level might be misleading: if the intermediary uses an algorithm that splits the market by dividing the keywords among the clients, then an increase in the set of competing advertisers concentrated under the same intermediary would reduce the keyword-level of both buyer concentration and CPC. To deal with this problem, we propose a method to measure concentration at the market-level, where markets are (non-overlapping) sets of keywords. We then map the keyword-level data summarized in Table 1 into market-level proxies of demand concentration and search engine revenue. Before doing that, however, we explore in the next section some keyword-level descriptive evidence.

V Keyword-Level Descriptive Evidence

In this section, we use keyword-level data to analyze whether intermediaries' choices are indicative of strategies that could reduce or increase search engine revenues. Since the market is characterized by a large number of heterogeneous MAs, each with a broad set of possible actions, we focus primarily on the seven large networks and analyze only a few of their possible actions. In particular, we begin by describing what happens when an advertiser who used to bid autonomously, joins an intermediary. Then, we show evidence suggestive of the presence of either market split or within-keyword coordination.

Individual Advertisers Joining MAs. Figure 3 reports the evolution of six variables as advertisers transition from bidding individually to bidding through a MA. We indicate $t = 0$ to be the first year after the advertiser joins a MA. Hence, to the left of the red, vertical line we report the linear fit – with the confidence intervals – of the yearly average of the variables across all the advertisers in the periods in which they bid autonomously and to the right of this line the averages under delegated bidding to MAs. The plot in the top-left corner displays a clear tendency for the number of keywords to increase under MA bidding. Indeed, the average number nearly doubles, from about 250 keywords to nearly 500 keywords. The top-right plot indicates that the average price of keywords declines as the average CPC goes from about \$4 per click to \$3 per click. The advertisers' position, instead, does not experience any significant jump, as shown by the middle-left plot. Middle-right and bottom-left plots refer to the type of keywords. *Long tail keywords* are longer, more specific keyword variations containing at least four terms. By being more specific they are both exposed to less competition and more likely to be searched by users closer to the bottom of the purchasing funnel.[12] They typically guarantee less competition (lower cost) and more clicks. *Branded* are those keywords that include as one of their terms any specific brand (see Golden and Horton [2018]). No significant change is evident for this variable. The bottom-right plot reports the number of markets entered. Although we will explain the details of how markets are constructed only in the next section, in essence these are groups of closely related keywords. Since the number of keywords (especially the *long tail*) grows, while the number of markets declines, this suggests MAs narrowing the focus of the keywords selected.

Figure 3: Individual Advertisers Joining MAs

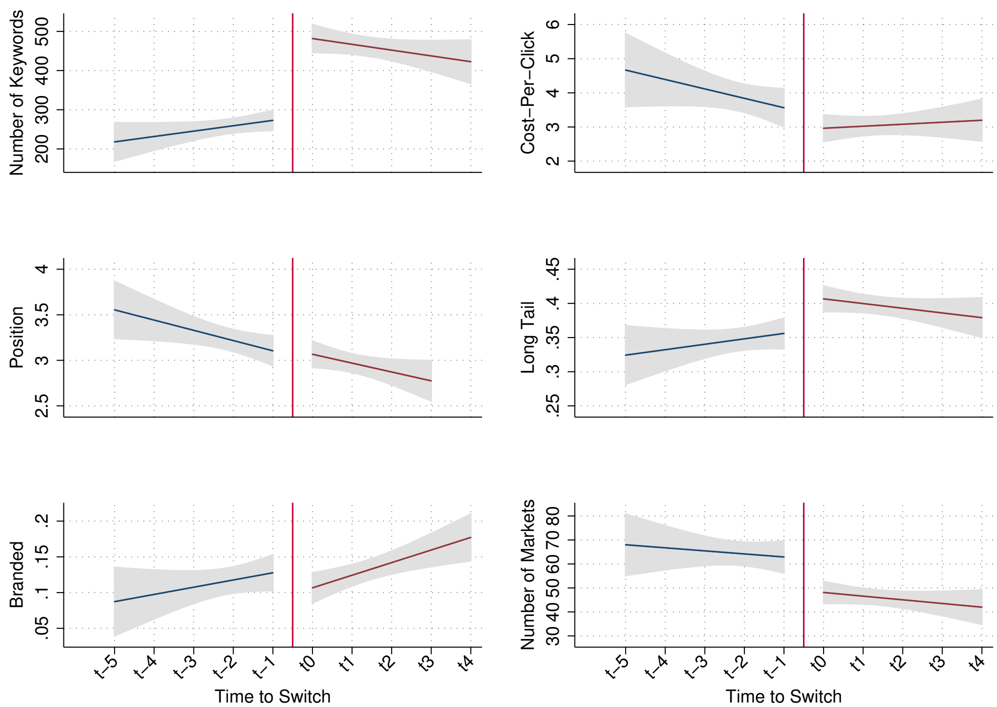

Notes: blue (maroon) lines are linear fits of average values before (after) joining a MA at *t0* (red vertical line). The reported variables are (left to right and top to bottom): *Number of Keywords*, *Cost-per-Click*, *Position*, *Long-tail Keywords*, *Branded* and *Number of Markets*. *Cost-per-Click* value is reported in USD, the shaded area corresponds to the standard deviation of the mean.

However, it is risky to analyze the effects of intermediary concentration by looking at variations in concentration driven by the incorporation of formerly independent bidders joining MAs. Some advertisers might join a MA due to their inability to optimize bidding. But then the lower CPC after joining might be explained by excessively high bids in the previous period, rather than by bid coordination by the intermediary.

[12]For instance, "charity donations furniture pickup" is an example of a long tail modification of "charity donations." The former is said to be closer to the bottom of the (purchasing) funnel to illustrate the idea that queries for keywords like "charity donations" describe the situation of a consumer still forming an initial opinion about the purchase, while queries like "charity donations furniture pickup" are more likely to be associated with a consumer closer to be finalizing the purchase decision.

In the analysis below, we therefore rely on a different type of variation in demand concentration: the one produced by ad networks incorporating previously independent MAs. In these situations, it is reasonable to assume that bids are already optimized from an individual bidder's perspective and that any strategy change is driven by the intermediaries' incentives to coordinate their advertisers' actions, as described earlier.

Network Expansions via MA Acquisitions. Next we start looking at the effects of increased demand concentration driven by network expansion via M&As. As discussed earlier, there is a distinction between what an intermediary can do within a single keyword auction and across multiple auctions. It is therefore informative to look at whether, after the acquisition of a MA by a network, there is any change in the overlap in the sets of keywords of the clients of either the network or the acquired MA. If the overlap declines, it might indicate that the intermediary splits the market by keywords, while if it stays identical (or grows) it might indicate that most of what the intermediary does takes place within-auctions. The evidence is suggestive that both strategies are adopted, although to different extents across the seven networks. As an illustration of keyword splitting, consider what happened when the MDC network acquired the Forsman MA: most of the keywords that used to be shared by clients of both MDC and Forsman before the merger stopped being shared afterwards (and only a few new common keywords were introduced). For the opposite situation consider, instead, what occurred following WPP's acquisition of the Shift MA. This was associated with an expansion of the shared keywords: most of the keywords that were shared before the merger continued to be shared after it (and new shared keywords were also introduced). In the appendix, Figure H.1 presents the exact details of these two cases, along with 4 other cases involving different networks. Overall, the great variety of possible strategies and the heterogeneity across networks in their usage make it difficult to quantify their impacts. Thus, in the next section we propose an empirical strategy using market-level data to quantify the effects of intermediary concentration on Google's revenues.

VI Market-Level Empirical Strategy

The relationship we seek to uncover is between the concentration of bidding by intermediaries and changes in Google's revenues. In particular, we assume the following linear relationship:

$$\log(R)_{mt} = \beta HHI_{mt} + \phi X_{mt} + \tau_t + \gamma_z + \epsilon_{mt}, \tag{1}$$

where the subscripts t and m indicate year and market respectively and $\log(R)$ and HHI are proxy for the search engine's (log) revenues and demand concentration respectively. As specified below, τ_t and γ_z are fixed effects for time and "thematic clusters", while X_{mt} are characteristics of the market-time (we will consider the number of organic links, plus a series of keyword-related and market-related controls). The coefficient of interest is β. A positive coefficient supports the hypothesis that greater concentration (proxied by HHI) benefits the search engine's (log) revenues, whilst a negative one would indicate that the negative effects prevail. In an ideal environment, we would like to observe different levels of HHI assigned randomly to otherwise identical markets m. But the actual data differs from this ideal in several ways. The proposed empirical strategy aims at correcting such issues in three main steps: the definition of what are the relevant markets, the construction of proxy measures for revenues and demand concentration and the formulation of an IV to deal with both measurement error and omitted variable bias in the estimation of equation (1).

Step 1: Market Definition

Potential definitions of markets range from granular – the single keyword – to aggregate – the 23 industries provided by Redbooks. The latter help to identify the agency/network sector of specialization, but contain keywords that are too heterogeneous to analyse competitive and strategic effects (as discussed more generally in Werden and Froeb [2018]). In order to find a useful middle-ground, we apply state-of-the-art natural language processing methods and unsupervised clustering techniques to form keyword groups interpretable as markets. The method entails two steps: first, we use an unsupervised learning algorithm to represent keywords as numerical vectors (*keyword vectorization*); second, we group the vectorized keywords into clusters according to a two-layer clustering, the first based on their semantic similarity (*thematic clustering*) and the second based on their proximity in terms of advertiser co-occurrences (*competitive clustering*).

A key element for the first step is the availability of a *corpus* (i.e., body of text) on which the algorithm learns the association between words. Given the goal of identifying relevant markets within the online advertisement industry, the ideal corpus should be informative about how consumers find products and services online. With such a corpus, the approach described below mimics what is sometimes done in antitrust cases: surveying consumers about the products they see as belonging to the same product space. Without aiming for the same accuracy required for competition cases, we nevertheless see this approach as a valuable contribution. We first detail how it works and then discuss some of its limitations.

Keyword vectorization For each keyword appearing in SEMrush data, we need a vector representation. The reason is straightforward: "red car," "blue car" and "automobile" are three keywords that we would like to see grouped together, but using keyword match approaches (e.g., using matches between single words), only "red car" and "blue car" would be pooled together. The vector representation systems developed in the natural language processing literature are meant to directly address the issues related to synonyms and antonyms in text clustering or semantic similarity exercises. We use an unsupervised learning algorithm (GloVe, developed by Pennington, Socher and Manning [2014]) to obtain vector representations for each term within the keywords. The GloVe model is a word embedding system which builds on the classical matrix of word co-occurrences in a corpus – i.e., a sparse matrix with one row per document in the corpus, and one column per word, populated with the number of occurrences (see details in the appendix). We use a GloVe dataset pre-trained on 840 billion documents, corresponding to approximately 2.2 million unique terms, from Common Crawl in English, featuring 300 dimensions. Such an extensive corpus originating from mimicking the web surfing behavior of typical internet users makes the resulting vectorization analogous to surveying people about the proximity between keywords.[13] Similarly, when applied to the sponsored search keywords in our data, the vectorization should reflect the proximity between products and services identified by the semantic similarity between keywords. Once every keyword is split into its constituent *terms*, we proceed by merging every term with the corresponding GloVe vector. Finally, we obtain the vector representation of each keyword by summing together the vectors relative to all its constituent terms.

Layer 1 – Thematic Clustering. We perform the thematic clustering step within each of the 23 industries in which the advertisers are categorised in the Redbooks data. We use the GloVe vector representation of all the keywords belonging to all the advertisers within an industry by summing the vectors of their vectorized terms. Then, we run a spherical k-means clustering algorithm (Dhillon and Modha [2001]) on the vectorized keywords' matrix with 1,000 centroids, industry by industry, to group them into thematic clusters. As a result, we identify the semantic "themes" linking the keywords (robustness checks regarding the implementation of the k-means algorithm are discussed in the appendix). There are two main shortcomings of the thematic clustering approach. First, different geographical markets can be identified only up to the extent that the geographical aspect is explicit from the terms composing the keywords (and in the training corpus). Visual inspection of the clusters reveals that this is only sometimes the case (like "car rental Boston" and "car rental New York" being sometimes pooled together). Second, the thematic clusters pool together both substitute and complementary products/services. This is not necessarily a shortcoming: to the extent that the advertisers of complementary products are in competition for the limited ad space, the analysis would not be distorted. However, the possibility of joint marketing efforts by advertisers of complementary products is a concern (see Cao and Ke [2018] for a recent study of this type of marketing).

Layer 2 – Competitive Clustering. To incorporate supply side information into the clusters, we exploit the competitive structure within each thematic cluster to further subdivide them into what we will refer to as "markets." The basic idea is to pool together keywords that are close in terms of the set of advertisers bidding for them. This is implemented by constructing, separately for each thematic cluster, a sparse matrix whose rows correspond to the keywords in the cluster and whose columns match the advertisers that bid on at least one of these keywords. The resulting row vectors are projections of the keywords in the space spanned by the advertisers (which we consider, to all extent and purposes, the competitive structure space). Through such vectors, we build for each pair of keywords a measure of similarity (the Euclidean distance between the corresponding row vectors).[14] Finally, we feed the similarity matrix describing the proximity of each pair of keywords into a hierarchical clustering algorithm to partition the keywords into "markets."[15]

[13]The dataset, and GloVe code, are available at https://nlp.stanford.edu/projects/glove/. There are a number of other datasets available (e.g., trained on Wikipedia, on Twitter, with 25, 50, 100 or 200 dimensions, etc.). The results obtained using these alternative datasets are available from the authors upon request.

[14]That is, keywords showing similar sets of bidders are more likely to belong to the same competitive space and, hence, more likely to be in the same (unobservable) product space.

[15]In order to optimally prune the cluster tree we employ the Kelley, Gardner and Sutcliffe [1996] penalty

Table 2: Market-level descriptives, thematic and competitive – Analysis Sample

	Thematic Clusters				Competitive Clusters (Markets)			
	Mean	SD	Median	Observations	Mean	SD	Median	Observations
Market Characteristics								
# Advertisers	6.7	10.5	3.0	8,324	4.0	4.8	3.0	25,947
# Keywords	116.1	180.3	55.0	8,324	37.2	104.9	4.0	25,947
# Networks	2.79	1.77	2	8,324	2.22	1.26	2	25,947
Competitive Clusters	5	5	3	8,324	-	-	-	-
Market Variables								
$log(R_{m,t})$	10.89	2.27	10.92	29,796	10.41	1.96	10.37	52,476
$HHI_{m,t}$	2,765	2,311	2,000	29,899	2,740	2,257	2,000	52,476
Long Tail	0.32	0.35	0.18	29,899	0.27	0.37	0.01	52,476
$\Delta R_{m,t}$	-0.05	1.78	0.00	21,256	0.40	1.53	0.28	43,973
# of Results (mil)	76.93	269.19	21.52	29,899	75.97	231.28	19.7	52,476
# Clusters			8,324				25,947	

Notes: top panels (*Market Characteristics*) report the features of the thematic – left panel – and competitive clusters – right panel. *# Advertisers* is the total number of advertisers bidding in keyword auctions, *# Keywords* is self-explanatory, and *Competitive clusters* is the number of clusters identified by the hierachical clustering algorithm in the second layer. In below panels (*Market Variables*) we report relevant outcome and explanatory variables relative to the estimation sample: $log(R_{m,t})$ stands for search engine's market revenues, $HHI_{m,t}$ is our measure of demand concentration, *Long Tail* is an indicator for keyword with four or more terms, $\Delta R_{m,t}$ is the yearly change in revenues and *# of Results* is the number of organic results – in millions.

Table 2 reports summary statistics for the subset of thematic clusters and markets that we will use for the analysis. Since many clusters are composed of keywords that contribute either very little or nothing to the search engine's revenues, and are never involved in any of the changes in MA ownership that we exploit for the IV strategy, we keep in the analysis sample only markets that either experience variation in the instrument at least once during the sample period or, for the remaining ones, that are in the top quartile of revenues. This leads us to dropping markets that represent between 1 percent and 2 percent of the total yearly revenues. In the appendix, we report robustness checks regarding this sample selection.

In the top-right panel of Table 2, the summary statistics indicate that an average market has 37 keywords and 4 competing advertisers, with the number of competing advertisers within single keywords (not reported) being on average 1.62. The statistics in the top-left panel further show that there are on average 5 markets within a thematic cluster. The bottom panel of this table reports summary statistics for the market-level variables that we describe below. Before moving to that, however, we stress that we cannot directly test the quality of the clusters obtained as that would require a reference sample where keywords and markets are correctly associated. Nevertheless, lacking this type of sample, we resorted to random inspection of the cluster quality. Overall we find very satisfactory results with our initial motivating concern of related but

function. Further details on the algorithm are in the appendix.

different keywords (like "car" and "automobile") systematically pooled together. Moreover, we designed and implemented a simple task aimed at testing the reliability of the clusters, and we ran it through *Amazon Mechanical Turk* (see the web appendix for a description of the test design with some examples and the results). With the exception of the residual industry that pools together many heterogenous advertisers (*miscellaneous*), for all other industries the share of correctly classified keywords is between 80 percent and 90 percent.[16]

Step 2: Measurement of the Main Variables

Having defined markets, we can now proceed to measure the main dependent and independent variables.

Outcome Variable. Suppose that the clustering procedure has identified M markets, $m = 1, ..., M$. Denote as K_m the set of k keywords in market m. We can use our keyword-level data to construct a measure of search engine's revenues (R) in market m in period t by aggregating revenues over keywords:

$$R_{mt} = \sum_{k \in K_m} CPC_{kmt} * Volume_{kmt} * CTR_{kmt} \tag{2}$$

where CPC_{kmt} is the average cost-per-click of keyword k (belonging to the set K_m in market m) at time t, $Volume_{kmt}$ is its overall number of searches and CTR_{kmt} is the cumulative click-through rate of all the sponsored ad slots shown for keyword k.[17] There is substantial heterogeneity in the levels of revenues across markets, mostly driven by heterogeneity in volume and CPC. To perform a meaningful analysis of the association of the revenue's level and the level of concentration, we thus work with $\log(R)$.

Concentration Measure. Suppose we have a market m defined by the set of keywords K_m. For each keyword $k \in K_m$, there are J_k sponsored ad slots, each occupied by an advertiser a. Each of these slots brings a certain number of clicks, which are ultimately the advertisers' object of interest. We therefore measure the "market size" (S_{mt}) as the sum of all the clicks of all the ad slots allocated in all the keywords in market m. That is: $S_{mt} = \sum_{k \in K_m} Volume_{kmt} * CTR_{kmt}$. The intermediaries' market share is measured accordingly by summing together all the clicks of all the market keywords associated with the slots occupied by each of the advertisers that the intermediary represents. That is, for intermediary i, representing the set of advertisers A_i, the market share in market m at time t is:

$$s^i_{mt} = \frac{1}{S_{mt}} \sum_{a \in A_i} \sum_{k \in K_m} \sum_{j \in J_k} CTR_{jkmt} * Volume_{kt} * 1\{a \text{ occupies } j \in J_k\}. \tag{3}$$

[16] A random set of clusters is available for download and inspection at `https://github.com/GabrieleRovigatti/adNets_clusters`.

[17] For each k, the overall CTR_{kmt} is the cumulative sum of the number of clicks across all j ad slots appearing on the search outcomes page of keyword k: $CTR_{kmt} = \sum_{j \in J_k} CTR_{jkmt}$. Since CTRs are not part of our dataset, we supplement this information from *Advanced Web Ranking*. As discussed in the appendix, although the CTR data is likely to involve measurement error, our baseline findings are qualitatively robust to two sets of robustness checks. First, we exclude entirely the CTR from the analysis by setting all CTRs to 1 (see Table F.1 and F.2) and, second, we randomly re-match CTRs to keywords (see Figure F.1).

Thus, our concentration measure for market m at time t is the squared sum of each intermediary's market share, or: $HHI_{mt} = \sum_{i=1}^{I}(s^i{}_{mt})^2$.[18] As discussed earlier, the intermediary is the network, or, if not present, the MA. Figure 4 shows the dynamics of HHI in our sample; more specifically, in this figure for each market we take the difference between HHI_{2017} and HHI_{2014}. The figure makes it evident that, although many markets experience an HHI decline, the majority experience concentration increases and about 6,000 markets have an HHI increase of more than 1,500 points.

Figure 4: Change in HHI – 2014 to 2017

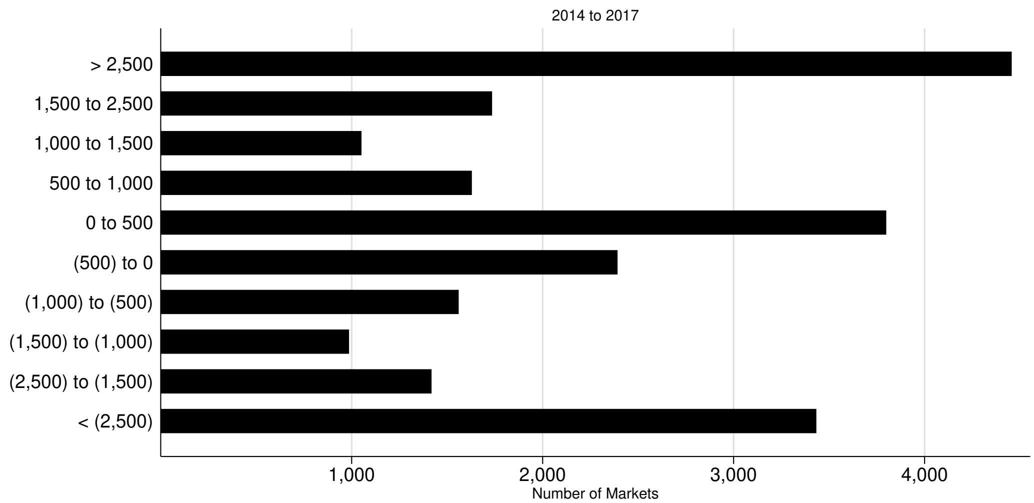

Notes: The bars report the number of markets – on the x-axis – grouped according to the differences between the HHI in 2017 and in 2014, clustered in ten classes. The HHI scale ranges 0 to 10,000.

Having defined the main variables, we can now return to the bottom panel of Table 2. There, we present basic summary statistics for the main variables entering our market level analysis. There we see, for instance, that the average market is highly concentrated with an HHI of 2,740. On average, the share of highly concentrated markets (i.e., those with an HHI of at least 2,500 points) is 40 percent and this share is increasing over time: from 37 percent in the first two sample years to 47 percent in the last year. Thus, while the overall market does not appear to be highly concentrated, the trend is in this direction.[19]

Step 3: Identification Strategy

There are two main reasons why the OLS estimation of equation (1) might lead to biased estimates of β. The

[18]Despite several theoretical and practical shortcomings of the HHI (see O'Brien [2017]), it is commonly used in both academia and competition policy to proxy for concentration (see Hastings and Gilbert [2005], Dafny, Duggan and Ramanarayanan [2012] and the US Horizontal Merger Guidelines). In our setting, the use of the HHI as a proxy for demand concentration has a theoretical foundation in the results of Decarolis, Goldmanis and Penta [2019] and, moreover, it will be empirically implemented through an IV strategy to control for measurement error problems.

[19]We remark that concentration in these advertising markets does not mechanically imply concentration in the intermediaries' market. We will return to this issue in the conclusions.

first is the measurement error problem associated with the HHI being only an imperfect proxy of demand concentration. The second is the risk of an omitted variable bias. For instance, a keyword k might have become suddenly fashionable for some exogenous reasons, such as changes in consumer taste; advertisers that were previously not interested in k now hire an intermediary to bid for it; moreover, they all hire the same intermediary as it is the one specialized in the market to which k belongs. This situation would likely induce observation of a positive association between intermediary concentration and the growth of search engine revenues, but it does not imply the existence of a causal relationship between the two phenomena. In practice, the available data allows us to reduce the risk of an omitted variable bias in two ways. First, we can include among the set of covariates market-time varying observables (like the number of organic links) that can likely control for phenomena such as the sudden change in appeal of a keyword, as mentioned above. Second, we can include fixed effects for the thematic clusters, thus exploiting the cross-sectional variation across markets within a cluster. This clearly reduces the extent to which relevant factors might be omitted since, for instance, omitted demand factors should be controlled through the thematic cluster fixed effects.

Nevertheless, since these fixed effects neither eliminate all risks of omitted variable bias nor deal with the measurement error bias, we use an IV strategy to estimate β. This strategy is inspired by the work of Dafny, Duggan and Ramanarayanan [2012] on the health insurance markets (also followed in Carril and Duggan [2018] to study concentration among suppliers of the US Department of Defence). It exploits changes in market structure originating from mergers and acquisitions (M&A) between intermediaries as a source of exogenous shock to local concentration. The idea is that M&A operations between intermediaries, especially the larger ones, are unlikely to be driven by the expectation of how the CPC would evolve in specific markets as a consequence of a merger. Indeed, M&A operations are a pervasive element of the ad network business. Individual agencies (MAs) are continuously purchased by the growing networks, often with hostile takeovers and exploiting moments of weaknesses of the agencies, such as when the founder is approaching retirement age or suddenly dies.[20]

Given that two merging intermediaries might have clients in a plethora of markets with possibly different starting levels of concentration, then the M&A operation generates useful local variation in the HHI. More specifically, for each market-time we compute the "simulated change in HHI" ($sim\Delta HHI_{mt}$) being the difference between the actual HHI and the counterfactual HHI (absent the merger) interacted with a post merger indicator. That is, we compute the change in concentration of market m at time t *induced* by the merger, *ceteris paribus*. Consider the merger between α and β in market m at time t^*. The contribution of the new entity to the concentration measure amounts to the squared sum of the shares of the merged firms,

[20] An important feature of this strategy is that, by isolating variation in the HHI that can be credibly attributed to changes in competition, it overcomes the problem stressed in the literature that the reduced-form nature of equation (1) makes it hard to identify the causal impacts of competition on market outcomes, see O'Brien and Waehrer [2017] and Berry, Gaynor and Scott Morton [2019].

which is by construction greater or equal than the contribution of the counterfactual with unmerged firms:

$$sim\Delta HHI_{mt} = \underbrace{(s^{\alpha}_{m,0} + s^{\beta}_{m,0})^2}_{\text{Share of merged firm}} - \underbrace{((s^{\alpha}_{m,0})^2 + (s^{\beta}_{m,0})^2)}_{\text{Sum of single firms' shares}} \times \mathbb{1}(t \geq t^*) = 2s^{\alpha}_{m,0}s^{\beta}_{m,0} \times \mathbb{1}(t \geq t^*), \tag{4}$$

where the subscript 0 denotes the year before the merger year t^*. We use, for each market-year, the variable $sim\Delta HHI_{mt}$ as instrument for HHI_{mt}. There are in total 21 mergers in our sample (details on each merger are in Table A.2).[21] Across networks, there is heterogeneity both in the number and the size of the MAs acquired. While Dentsu-Aegis appears to be the most "active" network with 8 acquisitions – including the one with most clients, Merkle, – WPP secured the largest acquisition in terms of presence in the markets (*SHIFT Communications* with clients active across 1,049 different markets). Some acquisitions take the form of hostile takeovers, with subsequent attempts to buy back independence and, as mentioned above, we observe two cases of divestitures. The effects of these M&A's on the HHI measure described above are substantial: across markets affected by mergers, the average HHI increase between the year of the merger and the proceeding year is 245.10 points.[22] For our baseline estimates, we will use an IV that exploits the variation from the whole set of M&A episodes. Clearly, the instrument's validity would be violated if the M&A operations were driven by expectations about revenue performance in the search auctions. Thus, we also look in isolation at the large merger episodes involving several clients active in many markets, as they are the least likely to be endogenously driven by revenues in local markets. Furthermore, the larger the merger the more likely the companies interested will do advertisement activities outside Google's search auctions, thus making less likely their endogenous determination within our empirical framework.

Using $sim\Delta HHI_{mt}$ as instrument for HHI_{mt} therefore entails the following first-stage regression:

$$HHI_{mt} = \beta^{FS} sim\Delta HHI_{mt} + \phi X_{mt} + \tau_t + \gamma_z + \epsilon_{mt}, \tag{5}$$

where the set of covariates entering the regression along with $sim\Delta HHI_{mt}$ is the same as those in equation (1). What sign to expect on the estimate of β^{FS} is an interesting question. To the extent that there is persistency in the market shares, we would expect a positive sign and this is indeed the finding of Dafny, Duggan and Ramanarayanan [2012] for the US health insurance market. Nevertheless, it would not be unreasonable to see a negative sign if a merger between intermediaries leads some clients to leave in order to avoid sharing a marketing resource with rivals (i.e., avoiding "sleeping with the enemy" Villas-Boas [1994]).

[21] When a market is affected by more than one merger, $sim\Delta HHI_{mt}$ is the sum of the values that it would assume were each merger considered separately.

[22] If we look separately at each merger, then the HHI increase across all the 21 mergers is 120.06 points, while it is 417.20 points across the four largest mergers. To put these numbers in perspective, consider that, according to the US Horizontal Merger Guidelines, when a merger results in an HHI increase of more than 200 points and a highly concentrated market, it will be "*presumed to be likely to enhance market power.*"

The remaining step of the empirical strategy requires estimating the following reduced-form regression:

$$\log(R)_{mt} = \beta^{RF} sim\Delta HHI_{mt} + \phi X_{mt} + \tau_t + \gamma_z + \epsilon_{mt}. \tag{6}$$

VII Results

We begin the presentation of our results from the first-stage and reduced-form estimates in Table 3. It reports the estimates for five different model specifications, gradually expanding the set of covariates. Model (1) includes demand concentration only, while model (2) adds thematic clusters fixed effects. Model (3) adds year fixed effects, while model (4) also adds a control for the number of organic results, which captures the "popularity" of the keywords in the market, thus reflecting the appeal to customers. This latter model is our baseline. Indeed, while model (5) includes further controls for the types of keyword composing the market (i.e., the average number of *long-tail* and *branded* keywords), we know from the earlier discussion that these might be endogenously determined by the strategies of intermediaries. Nevertheless, by way of comparison it is useful to report the estimates of model (5) as they offer a way to check whether these keyword choices affect revenues through increases in concentration.

Table 3: Reduced Form and First Stage Estimates

	(1)		(2)		(3)		(4)		(5)	
	RF	FS	RF	FS	RF	FS	RF	FS	RF	FS
sim$\Delta \hat{HHI}$	-6.761***	0.618***	-4.070***	0.957***	-3.842***	0.830***	-3.831***	0.829***	-3.723***	0.831***
	(1.110)	(0.170)	(1.133)	(0.0790)	(1.162)	(0.0914)	(1.165)	(0.0915)	(1.165)	(0.0913)
Weak Id. F-Test		13.21		146.99		82.37		82.18		82.94
Underid. F-test		4.56		13.67		11.02		11.01		11.02
Observations	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476
Cluster FE			✓		✓		✓		✓	
Year FE					✓		✓		✓	
Organic Results							✓		✓	
Keyword Characteristics									✓	

Notes: the dependent variable is the (log) revenues, R_{mt}. RF columns report the reduced-form estimates, FS columns the first-stage ones; the models – (1) to (5) – have an increasing number of controls and fixed effects. Model (1) includes industry fixed effects. In the baseline model, reported in column (4), we control for the average number of organic results, thematic cluster and year fixed effects. Model (5), in which we add keyword characteristics such as the share of long tail and branded keywords, is likely to suffer from an additional source of endogeneity. In all models the standard errors are clustered at the thematic clusters level.

Both the first stage and reduced form estimates in Table 3 are rather stable across model specifications. As expected, magnitudes are impacted the most by the addition of thematic cluster FE between model (1) and model (2). We consider the latter level of clustering quite useful to control for most of the potential omitted variable bias and, therefore, rely on this cross-sectional variation within clusters as a main source of causal identification. In terms of the results, the positive sign of the $sim\Delta HHI$ estimate in the first-stage regression indicates that the HHI increases in the markets where the simulated HHI grows the most. This implyies that the clients of a MA acquired by a network tend to remain within the acquired network. Hence, although the estimated coefficient of 0.829 falls short of 1, its large magnitude indicates that the "sleeping with the enemy" concern does not appear to drive a reshuffling of clients among acquired MAs.[23] The large value of the F-statistics also confirms the relevance of the proposed instrument. On the other hand, the reduced form estimates indicate a negative and statistically significant relationship between (log) revenues and the simulated change in HHI. Considered together, reduced form and first stage regressions inform us that the IV estimate indicates a negative impact of intermediary concentration on the platform's revenues.[24]

Table 4: Effect of Concentration on Search Engine Revenues - OLS and IV Estimates

	OLS					IV				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\hat{HHI}$	-2.217***	-2.120***	-2.129***	-2.122***	-2.130***	-10.93***	-4.252***	-4.630***	-4.620***	-4.479***
	(0.0718)	(0.0567)	(0.0573)	(0.0572)	(0.0569)	(2.902)	(1.068)	(1.200)	(1.204)	(1.201)
Organic Results (billion)				0.252***	0.263***				0.206***	0.225***
				(0.0437)	(0.0458)				(0.0463)	(0.0477)
Keywords Characteristics										
Branded Keyword					0.396***					0.458***
					(0.0537)					(0.0639)
Long-tail Keywords					-0.0908**					-0.0491
					(0.0367)					(0.0423)
R^2	0.07	0.62	0.62	0.62	0.62					
Observations	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476
Cluster FE		✓	✓	✓	✓		✓	✓	✓	✓
Year FE			✓	✓	✓			✓	✓	✓

Notes: the dependent variable is the (log) revenues, R_{mt}. Columns (1) to (5): OLS estimates of equation (1), with an increasing number of fixed effects and controls. Columns (6) to (10): IV estimates – where we instrumented HHI_{mt} with the merger-induced change in concentration as defined in equation (4). In all models the standard errors are clustered at the thematic clusters level.

This is indeed what we observe in Table 4. This table reports OLS (columns 1 to 5) and IV (columns 6 to 10) estimates. Both sets of coefficients are negative and statistically significant. IV coefficients are larger, being about twice the corresponding OLS ones. This is compatible with both measurement error in the demand concentration proxy and with residual omitted variable bias. As expected from the estimates in Table 3, there is a large drop in the magnitude of the coefficient of the IV estimates when controlling for thematic cluster fixed effects. With these fixed effects, the estimates are remarkably stable across all models, in terms of both magnitude and significance. Controlling for either organic results or keyword characteristics

[23]In the appendix, Figure J.1 reports the results of the instrument's monotonicity test proposed by [Angrist and Imbens, 1995]. Verifying that monotonicity holds – as Figure J.1 indicates – is important because the sign of the first stage regression is theoretically unclear and, also, because splitting the market by keyword may create a negative relationship between HHI and simulated HHI over some of the latter's range.

[24]In the appendix, Figure G.1 allows us to visualize the changes in log revenues before and after an acquisition-driven change in concentration. Although, due to the limited time length of our data, this falls short of a proper event study analysis, the drop in the average revenues post-merger displayed in this figure is consistent with the econometric estimates presented in Table 3.

has quantitatively no impact on the findings.

To ease the economic interpretation of the estimates, it is useful to consider the average HHI increase induced by mergers of 245.10 points. In this case, the baseline estimate (column 9) implies a decrease in revenues of 11.32 percent (that is $4.62 \times 100 \times 0.02451$). This is an economically relevant result indicating that intermediary concentration shifts a significant share of the value created by search advertising away from the platform and toward buyers. If we look separately at the merger episodes, then the decline in revenues amounts to 5.55 percent for the average merger and to 19.27 percent for the four largest mergers (recall that the corresponding increases in HHI are 120.06 points and 417.20 points respectively). Overall, these figures describe a sector in which the networks are eroding the search engine's revenues.

Table 5: Individual Mergers

Panel a): Individual Mergers – Reduced Forms and First Stages								
	Sapient		Merkle		Shift		Forsman & Bodenfors	
	RF	FS	RF	FS	RF	FS	RF	FS
sim$\Delta \hat{HHI}$	-4.911*	1.026***	-5.981***	1.388***	4.536	0.707***	-16.30**	6.357***
	(2.882)	(0.387)	(1.181)	(0.0386)	(2.998)	(0.230)	(6.388)	(0.159)
Observations	4,776	4,776	3,047	3,047	3,013	3,013	981	981

Panel b): Individual Mergers – OLS and IV Estimates								
	Sapient		Merkle		Shift		Forsman & Bodenfors	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
$\hat{HHI}$	-5.302***	-4.786*	-4.516***	-4.308***	-3.823***	6.415	-5.236***	-2.563**
	(0.208)	(2.547)	(0.293)	(0.871)	(0.175)	(4.963)	(0.672)	(0.999)
Observations	4,776	4,776	3,047	3,047	3,013	3,013	981	981
Industry FE	✓		✓		✓		✓	
Year FE	✓		✓		✓		✓	
Organic Results	✓		✓		✓		✓	

Notes: the dependent variable is the (log) revenues, R_{mt}. For each reported M&A action (*Sapient, Merkle, Shift and Forsman & Bodenfors*), the estimation sample amounts to all markets involved – i.e., all markets in which at least one of agency's clients was bidding before the merger. In panel a) odd columns report the reduced form and even columns the first stage estimates, respectively. In panel b), odd columns report the OLS and even columns the IV estimates. All models feature controls for the average number of organic results, industry and year fixed effects, and the standard errors are clustered at the thematic clusters level.

Robustness. We assess the reliability of the above estimates to several modifications. First, to ensure the reasonableness of the IV approach, we repeat the analysis looking exclusively at the largest mergers. We perform this analysis separately for each one of the four largest mergers, involving four different networks. The results reported in Table 5 are broadly consistent with the baseline estimates presented earlier. The top panel reports reduced form and first stage estimates, while the bottom panel reports OLS and IV estimates. In all cases the model specification is that of the baseline estimates (model (9) in the previous table). For the mergers involving Sapient, Merkle and Forsman & Bodenfors, both the significance and the magnitude of the estimates track closely what is reported in Table 4 (although the IV estimates are smaller for the Forsman

& Bodenfors merger). For the Shift merger, however, the reduced form is not statistically significant. Thus, while the OLS estimates are in line with those of the other mergers, this is the only IV estimate that is not significant. Possibly this is because WPP never fully integrated Shift into its systems as this company entered the WPP network indirectly through an acquisition by a large Canadian affiliate of WPP, National Public Relations, that maintained Shift as its MA for its US clients. Despite some heterogeneity across the cases, the overall takeaway is that, even narrowing down the analysis to the subset of the data where the IV strategy is the most reasonable, the results are qualitatively close to those of the baseline estimates.

Table 6: Robustness Checks

	Market Definition		Two-layers Clustering		
	Industry Level	Thematic Clusters	GloVe mean	500K	N/30K
	(1)	(2)	(3)	(4)	(5)
$\hat{HHI}$	9063.3	-10.75***	-3.870***	-2.600***	-3.404***
	(1427185.6)	(1.572)	(0.600)	(0.716)	(0.860)
Observations	68	22,353	68,368	54,621	52,867
Industry FE	✓				
Cluster FE		✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓

Notes: the dependent variable is the (log) revenues, R_{mt}. The definition of m changes across models. In column (1), we do not perform any clustering exercise, and m is the industry level (there are up to 23 industries per year). In column (2), m is the thematic clusters level. In columns (3) to (5) m is a competitive cluster, but the clustering algorithm used is not the same as in the baseline estimates. In column (3), we average over GloVe-vectorized terms – instead of summing up the vectors – before performing the step-1 clustering exercise, column (4) features 500 clusters *per industry* in step-1, while in column (5) we repeat the exercise with a size-dependent number of clusters – i.e., with 1 cluster every 30 unique keywords in the sample. All models feature controls for the average number of organic results, industry – (1) – or thematic clusters and year fixed effects. Standard errors are clustered at the industry or thematic clusters level.

We consider next five sets of robustness checks presented in Table 6. All estimates reported in this table are the IV estimates of the baseline model specification. In the first two columns, we explore the effects of using alternative definitions of "markets." In column (1), markets are defined as the advertisers' industries. Earlier we discussed why this is likely to be problematic as industries are an excessively broad category and, indeed, the estimates in column 1 indicate a very unreasonable IV estimate. In the following column, we thus return to a definition of market based on the 2-layer keyword clustering procedure, but use as markets the thematic clusters. The qualitative insight of a negative and significant β is maintained, but the magnitude is substantially larger, which is reassuring with regards to the fact that our baseline are a conservative estimate of the true effect. The following three columns explore the robustness of the estimates to the details of the proposed 2-layer approach. In column (3), instead of using the term-by-term sums of GloVe vectors, the thematic clusters are built by averaging GloVe vectors within keywords. Intuitively, averaging the vectors attenuates the effects of "topical" terms, whose weight is instead amplified by the sum;

moreover, the latter method tends to isolate long tail keywords – keywords with more terms face a higher likelihood of being positioned "far away" in the vector space. As a result, the averaged GloVe keywords are less sparse, and possibly harder to cluster. Despite this, the estimates are very close to the baseline ones. The next two columns, (4) and (5), explore related modifications of the clustering approach involving the number of centroids of the k-means algorithm, using either 500 centroids or the number of keywords in the industry divided by 30. Again, the baseline estimates appear robust to these modifications.[25]

Channels. The findings above indicate that the effects of increased buyer power dominate the efficiency gains from which the search engine might benefit. To better understand our findings, here we analyze the channels through which competition impacts revenues. In Table 7 we explore the relationship between market concentration and changes in the average CPC (columns 2), search volume (column 3) and number of keywords (columns 4). The estimates are noisy and not statistically significant for the latter two, but they are negative and strongly significant for the CPC. These point estimates are hence suggestive of the way in which market concentration affects market revenues: the number of keywords and volume show a null effect and the cost-per-click, in accordance with the theoretical predictions about the incentive to coordinate prices, decreases substantially.

Table 7: Revenues Components – IV Estimates

	$log(R)$	$log(cpc)$	$log(vol)$	$log(\#keys)$
	(1)	(2)	(3)	(4)
$H\hat{H}I$	-4.408***	-1.286***	-0.546	-0.956
	(1.164)	(0.447)	(1.015)	(0.726)
Observations	52,476	52,476	52,476	52,476
Cluster FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: IV estimates using five different outcomes: $log(R)$ (1), $log(cpc)$ (2), $log(vol)$ (3), and $log(\#keys)$ (4). All models feature controls for the average number of organic results, thematic clusters and year fixed effects. Standard errors are clustered at the thematic clusters level.

In interpreting this evidence, it is interesting to recall the descriptive evidence presented earlier. While the estimates in Table 7 exploit variation driven by network M&A activity, the graphical evidence in Figure 3 is based on what happens when individual advertisers join agencies. As that figure shows, when an individual advertiser joins an agency, a multiplicity of changes occur and some, such as the expansion in the number of keywords, are clearly beneficial for the search engine. But when an intermediary is acquired by a network, the changes in the types of bidding behavior are more subtle and pertain exclusively to what is allowed by greater concentration within an intermediary: bid coordination (on the negative side of the search engine's

[25] Additional robustness checks are presented in the appendix.

revenues) or also budget coordination and internalization of the externalities across slots (on the positive side of the search engine's revenues). Thus, the lack of effects on the number and volume of keywords in Table 7 is indicative of demand concentration by itself which has mostly negative effects on the search engine revenues, whereas the activity of MAs more generally has both positive and negative effects on revenues.

The capacity of concentrated networks to lower the CPC helps to explain why advertisers use them instead of replacing them with their own bidding algorithms, despite the ease of developing such algorithms and the networks' hefty fees (of the order of 17 percent of ad spending).[26] But what are the means through which networks can lower the CPC? We discussed this question with industry experts. Some experts highlighted a mechanical effect linked to the quality scores: demand concentration allows the larger intermediaries to pool together relevant data from rival advertisers and this allows a better optimization of the quality scores of their clients, which mechanically implies lowering the CPC. The other answers that we got can be grouped into two broad strategies: easing competition among the network's clients and bolstering competition between ad selling platforms.[27]

The first type of strategy involves employing bidding algorithms that exploit the targeting features of Google Ads to reduce (or even eliminate) competition among clients of the same network. Recall the example earlier about splitting the market by targeting on two dimensions (geography and timing). On Google Ads, the set of targeting dimensions is extensive and includes: demographics (6 groups for age, 6 for income, 2 for parental status and 2 for gender), device (computer, tablet or mobile phone) and audiences (i.e., groups of people with specific interests, intents, and demographics, as estimated by Google.). This makes algorithms capable of implementing nearly infinite forms of market splits. But this also makes it nearly impossible to detect this type of coordination, absent the granular data available to Google only (which has never been shared for research). Market segmentation might also be implemented by splitting keywords. In this case, our data allow the capture of some interesting patterns. For instance, significant shares of marketing budgets are spent for own brands and those of rivals [Blake, Nosko and Tadelis, 2015]. As mentioned earlier, advertisers' explicit coordination to stop bidding on each other's brands is unlawful, but it is not if autonomously implemented by their intermediary. Interestingly, our data indicate that after a merger, among the merged entity's clients, there is a decline in the probability of an advertiser obtaining an ad slot on the branded keywords of same-network advertisers, and an increase in the probability of obtaining a slot in one's own branded keywords (see Figure H.2 in the appendix).

Algorithms might also be written with the different goal of reducing the size of bids by competing advertisers for the same keyword. A glimpse of what might be happening in practice can be grasped by

[26] The 17 percent figure is obtained as the sum of the fees for the agency of record (5 percent) and of the trading desk (12 percent) reported in Figure 6 in Adshead et al. [2019]. ISBA [2020] also finds similarly large fees, as well as denouncing the presence of large hidden fees. Both studies, however, are based on display advertising where the collaboration of both advertisers and publishers allows measuring intermediary fees.

[27] Selected quotes from the interviews are reported in an ad hoc web appendix.

looking at the case of iProspect – a leading independent MA, later acquired by the Dentsu-Aegis network. This company is credited with having developed one of the earliest automated bidding systems for search auctions. It is thus intriguing that the scientist who developed this algorithm is also the leading author of a computer-science paper, Kitts, Laxminarayan and Leblanc [2005], on cooperative strategies for search auctions that proposes "a coordination algorithm that optimally distributes profit on the auction between participating players" and shows its implementation in real data. Finally, an interesting question related to algorithms is whether the algorithms that we are discussing entail AI. All the experts that we consulted think that, while this is unlikely to be the case for most of the algorithms that are currently used, there is a rapidly growing trend toward AI (see footnote 9). But AI bidding algorithms are not seen as a panacea for search for three main reasons: the likely large costs required to make these algorithm learn effective bidding strategies in the face of the amount of data required, the likely low rate of convergence and the systematic presence of shocks in keyword value.

The second type of strategy involves increasing the competition level between ad selling platforms. The most straightforward way to do this is by splitting the advertisers' marketing budget across more digital ad platforms. This form of market segmentation differs from those described earlier because its efficacy hinges on the availability of alternatives to search ads on Google. These alternatives mainly involve other search platforms (primarily Amazon and Bing), ad platforms in display advertising (where there are a handful of competitors to Google) and social media advertising (mostly Facebook). There is also a second way though which large networks can exploit the presence of competing platforms to reduce the cost of search ads, which is bargaining. Within Google's rigid auction system, there seems to be no room for bargaining, but this is a naive view, according to the experts we spoke to. There are simple tweaks to the auction algorithm that may implement side deals with networks. Two examples are i) excluding the losing bid(s) coming from of a network's own client(s) when calculating the CPC applied to the winning bid(s) of the same network's client(s); and ii) bolstering the quality scores. The first method is essentially equivalent to (selectively) lowering the reserve price, while the second exploits the black-box nature of the quality scores which only Google fully controls. Although bargaining with Google might not seem feasible, it should be recalled that the networks themselves have a huge economic size – WPP, Omnicon and IPG all earn yearly revenues in excess of $15 billion each – and they are essential partners in order for Google to reach millions of advertisers.[28] There is, however, no guarantee that deals negotiated by the networks will benefit advertisers, as we discuss below in the conclusion.

[28]This might also explain why over the years each of these networks has been able to develop special arrangements with Google to buy ad space on its platforms (see quotes in the appendix). Bargaining might be feasible also for the largest advertisers, as suggested by the fact that the experiment in Blake, Nosko and Tadelis [2015] was run when the eBay marketing department was clashing with Bing-Yahoo! to negotiate lower ad prices and eBay tried proving its bargaining strength by ending ads for the keyword 'eBay'.

VIII Conclusions

The findings we present indicate that concentration among the intermediaries bidding on behalf of advertisers in sponsored search auctions negatively and significantly impacts search engine revenues. Despite the potential benefits for the search engine from the increased efficiency that intermediaries bring, especially through enhanced speed and better data, the negative revenue result is indicative of the capability of intermediaries to reduce average prices. This is a novel insight into what is currently one of the largest and fastest growing advertising markets and underscores the relevant role of intermediaries. This has received limited attention in the literature so far. The three key elements of our analysis are first, a novel dataset linking together keywords, advertisers and intermediaries; second, a new approach to defining markets by aggregating keywords through a 2-layer machine learning algorithm incorporating both demand and supply information; and third, the application of an IV strategy based on intermediary mergers.

Several questions are left open for future research. We conclude by briefly exploring two questions whose answers are particularly important in interpreting the broader impacts of our findings. The first question is about the internal or external frictions that could slow down, or even revert the processes discussed here. Internal frictions would involve the advertisers' choice to forego the benefits of joint bidding in order to avoid sharing intermediaries (and data) with rivals. But this type of friction does not appear to be salient according to our analysis. Instead, external frictions can derive from the actions of either antitrust authorities or the platform. The former are limited to the very specific cases mentioned in the introduction, whilst the latter could involve a large spectrum of actions initiated by the ad selling platform. Four industry trends might reveal what the selling platforms are doing to reduce their loss of market power: increasing the auction reserve price, reducing the number of ad slots offered, promoting disintermediation services and lastly – as done most notably by Facebook – changing the auction format. Among these four changes, market efficiency is more likely endangered by the first two. In May 2017, Google introduced higher reserve prices differentiated by keyword. In a market dominated by concentrated intermediaries, however, substantial reserve price increases might be required to increase the average CPC. But this would likely hurt the "wrong" advertisers (i.e., those not sharing a common intermediary). Small advertisers placing low bids near the reserve price might find themselves either paying substantially higher prices or being outright excluded from the set of ad that is displayed, thus undermining market efficiency. Over the last few years, Google also started reducing the available ad slots (by eliminating the side-bar and adding a bottom-bar with fewer ads).[29] But clearly this approach to increasing competition, by making ad space more scarce, might have the same perverse effect of hurting the "wrong" advertisers mentioned above in relation to the reserve price.[30]

[29] In the buyer power literature, this is remindful of how Snyder [1996] suggests dealing with collusion.

[30] Regarding the reasons why we see the other two trends as less problematic, they are are as follows. First, a change in the auction format might serve to fix only the most detrimental types of effects produced by bid coordination on revenues and efficiency. This is, for instance, what Decarolis, Goldmanis and Penta [2019] argue about the switch from GSP to VCG auctions, like that implemented by Facebook in 2011. Second,

The second question is the extent to which the drop in Google's revenues may be passed on to consumers and, hence, positively contribute to consumer welfare. Since most advertisers operate in markets more competitive than internet search, a transfer of revenues from Google to the advertisers should induce a drop in their costs and, consequently, in consumer prices. If that were the case, increasing buyer power would represent a particularly desirable policy to address the concerns associated with platform concentration. In particular, it might reduce the platform market power without affecting market shares. This is important for search as the market size mirrors the extent of the within-group network effects [Belleflamme and Peitz, 2018]: the quality of search outcomes depends on the size of the user base. Hence, there is an evident risk with the alternative policies currently debated which involve either helping consumers switch between search engines or improving the quality of smaller search engines through mandated access to Google's data.[31]

The positive effects on welfare, however, require that advertisers benefit from intermediary concentration in the form of lower ad prices. The extent of this benefit depends on the degree of competition among intermediaries. To the best of our knowledge, there is no conclusive evidence on this issue. Silk and King [2013], in a landmark study on concentration in the US advertising and marketing services agency industry, find the industry to be reasonably competitive. But, as mentioned earlier, intermediary commissions are fairly high [Adshead et al., 2019]. Although the industry experts that we consulted (quoted in the web appendix) differed in their views about the extent of the competition, there are multiple reasons to consider the market to be reasonably competitive. In our data, when we look at the ad markets (i.e., the competitive clusters), there are typically only 2 networks per market, but the markets where intermediaries compete are likely to be broader than that. For instance, if we take the relevant market definition to be the advertisers' industry classification, then our data indicates that on average 6 out of the 7 networks are simultaneously present (moreover, for 13 out of the 23 advertisers' industries each network is present representing at least one advertiser). Furthermore, it is important to stress that the networks face competition from a competitive fringe of independent agencies and, more recently, also from consulting firms. In fact, all of the major consulting firms – especially Accenture, Deloitte and McKinsey – have stolen customers from the MAs by offering specialized support for digital advertising integrated with their other consulting services.

---

disintermediation – the practice by the selling platform of offering services in direct competition with those of the intermediaries – entails a choice by advertisers and, hence, we should expect the platform to offer valuable options to induce the advertisers to abandon their MA. For instance, Google is promoting its bidding services by emphasizing the good outcomes it achieved for advertisers, like Hewlett-Packard that switched from its MA to Google's DoubleClick Search. Clearly, trusting Google to bid on its own auctions as well as on rival ad selling platforms might seem problematic to some advertisers, and this will likely pose a limit to disintermediation. A related issue involves *smart bidding*, the suite of AI-bidding options that Google offers to advertisers and agencies to bid in its search auctions. The depth of Google's data might allow *smart bidding* to outperform the intermediary algorithms, leading to the intermediary exit form bidding services.

[31] For an overview of the policy proposals currently being debated to deal with market power by online platforms see the Stigler Report, 2019, Furman Review, 2019, Competition Policy for the Digital Era, 2019, CMA Interim Report on Online Platforms and Digital Advertising, 2019. Enhancing buyer power as a policy tool to deal with online platforms has been discussed by the EU Commission [2017] and Mullan and Timan [2018]. Our study adds to the limited set of cases upon which these recommendations were based.

There are, however, at least five elements that are possibly at play that limit the extent of competition and that deserve further analysis. First, intermediary competition is limited by the well known difficulties in measuring the returns to advertising, which Lewis and Rao [2015] indicate to be severe for online advertising. Second, a closely related feature regards the lack of transparency in intermediary reports to advertisers about how their budget is spent [ISBA, 2020]. While the former issue relates to an intrinsic difficulty in advertising, the second relates to the behavior of intermediaries who typically report very aggregated measures of how they allocated their clients' money. This contributes to explaining why advertisers may fail to optimize their bidding campaigns; something that powerfully emerges from Blake, Nosko and Tadelis [2015]. Third, the exact same features mentioned in our study for why bid coordination by a common intermediary can be valuable, all imply that advertisers might become locked-in. This is because obtaining the same benefits would require a joint deviation by competing advertisers, from their current intermediary toward a different one. Fourth, a collusive conduct between some of seven agency networks might be aided by some features, like their common ownership. Similar to the discussion in Azar, Schmalz and Tecu [2018], we find systematic common ownership among the 5 publicly listed networks, with BlackRock and Vanguard among the main shareholders in each of them (see Figure K.1 in the appendix). Fifth, some industry observers suggest even more complex forms of collaboration with Google having found ways to cooperate with the intermediaries in order to ensure its long run dominance, even at a short term cost. In recent years, there have been multiple revelations about agency kickbacks. An investigation by the US Association of National Advertisers, [ANA, 2016], states that "numerous non-transparent business practices, including cash rebates to media agencies, were found to be pervasive in the US." Nevertheless, some of the networks, like WPP, have responded by saying that they did not take part in the Google's US media rebate program, while others, like Omnicon, admitted to being part of it but argued that the rebate was passed down to clients. Overall, monitoring these five areas of concern would be an essential component of a policy intervention that seeks to make good use of advertising intermediaries as a remedy to the dominance of the largest online platforms.

The final concern worth mentioning regards dynamic implications. The previous discussion on the positive effects of buyer power ignores the potential downsides on consumer welfare that might result from two dynamic channels. The first regards innovation: increased buyer power by the merged networks may lead to reduced incentives to innovate by suppliers (Google). The second entails waterbed effects: increased buyer power by the merged networks may increase costs for other competing intermediaries, for instance due to a relative worsening of their clients' quality scores. This would lead to a worsening in choice (or service) for advertisers and, through their higher costs, harm consumers. Regarding these dynamic considerations, however, more than thirty years after the breakup of the Bell System in 1982 – the dominant firm of the US telecommunications sector at that time – how an economist should look at the long run effects of the loss of power by dominant firms like Google, or the Bell System is still an open and key question.

References

Adshead, Stephen, Grant Forsyth, Sam Wood, and Laura Wilkinson. 2019. "Online advertising in the UK." *PLUM Consulting report for the UK Department for Digital, Culture, Media and Sport.*

ANA. 2016. "Media Transparency Initiative: K2 Report." *K2 report for the Association of National Advertisers*, June.

Angrist, Joshua D, and Guido W Imbens. 1995. "Two-stage least squares estimation of average causal effects in models with variable treatment intensity." *Journal of the American statistical Association*, 90(430): 431–442.

Athey, Susan, and Denis Nekipelov. 2014. "A Structural Model of Sponsored Search Advertising Auctions." *mimeo.*

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. "Concentrating on the Fall of the Labor Share." *American Economic Review*, 107(5): 180–85.

Azar, Josa, Martin C Schmalz, and Isabel Tecu. 2018. "Anticompetitive Effects of Common Ownership." *The Journal of Finance*, 73(4): 1513–1565.

Bachrach, Yoram. 2010. "Honor among thieves: collusion in multi-unit auctions." *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagents Systems (AAMAS).*

Balseiro, Santiago R., and Ozan Candogan. 2017. "Optimal Contracts for Intermediaries in Online Advertising." *Operations Research*, 65(4): 878–896.

Balseiro, Santiago R, Omar Besbes, and Gabriel Y Weintraub. 2015. "Repeated auctions with budgets in ad exchanges: Approximations and design." *Management Science*, 61(4): 864–884.

Belleflamme, Paul, and Martin Peitz. 2018. "Platform and Network Effects." In *Handbook of Game Theory and Industrial Organization.* , ed. Corchon and Marini. Cheltenham: Edward Elgar.

Berry, Steven, Martin Gaynor, and Fiona Scott Morton. 2019. "Do Increasing Markups Matter? Lessons from Empirical Industrial Organization." *Journal of Economic Perspectives*, 33(3): 44–68.

Bird, Steven, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit.* " O'Reilly Media, Inc.".

Blake, Thomas, Chris Nosko, and Steven Tadelis. 2015. "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment." *Econometrica*, 83(1): 155–174.

Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. 2013. "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence." *American Economic Journal: Microeconomics*, 5(4): 163–187.

Calvano, E., G. Calzolari, V. Denicoló, and S. Pastorello. 2018. "Algorithmic Pricing and Collusion: What Implications for Competition Policy?" *mimeo.*

Cao, Xinyu, and Tony Ke. 2018. "Cooperative Search Advertising." *Marketing Science*, forthcoming.

Carril, Rodrigo, and Mark Duggan. 2018. "The Impact of Industry Consolidation on Government Procurement: Evidence from Department of Defense Contracting." National Bureau of Economic Research Working Paper 25160.

CEA. 2016. "Benefits of Competition and Indicators of Market Power." Council of Economic Advisers 20160414.

Celis, L. Elisa, Gregory Lewis, Markus Mobius, and Hamid Nazerzadeh. 2015. "Buy-It-Now or Take-a-Chance: Price Discrimination Through Randomized Auctions." *Management Science*, 60(12): 2927–2948.

Chipty, Tasneem, and Christopher M. Snyder. 1999. "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry." *The Review of Economics and Statistics*, 81(2): 326–340.

Choi, Hana, Carl F. Mela, Santiago Balseiro, and Adam Leary. 2019. "Online Display Advertising Markets: A Literature Review and Future Directions." Columbia Business School Research Paper, 18-1.

Craig, Stuart, Matthew Grennan, and Ashley Swanson. 2018. "Mergers and Marginal Costs: New Evidence on Hospital Buyer Power." National Bureau of Economic Research Working Paper 24926.

Dafny, Leemore, Kate Ho, and Robin S. Lee. 2019. "The price effects of cross-market mergers: theory and evidence from the hospital industry." *The RAND Journal of Economics*, 50(2): 286–325.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry." *American Economic Review*, 102(2): 1161–85.

Dai, Weija. 2014. "Matching with Conflicts: An Application to the Advertising Industry." mimeo.

Decarolis, Francesco, Maris Goldmanis, and Antonio Penta. 2019. "Marketing Agencies and Collusive Bidding in Online Ad Auctions." *Management Science*, forthcoming.

De Loecker, Jan, and Jan Eeckhout. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics*, 135(2): 561–644.

Dhillon, Inderjit S, and Dharmendra S Modha. 2001. "Concept decompositions for large sparse text data using clustering." *Machine learning*, 42(1-2): 143–175.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo. 2018. "The economic consequences of hospital admissions." *American Economic Review*, 108(2): 308–52.

Economist. 2016*a*. "The rise of the superstars." 26 March.

Economist, The. 2016*b*. "Too much of a good thing." 17 September.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review*, 97(1): 242–259.

Ellison, Sara Fisher, and Christopher M. Snyder. 2010. "Countervailing Power in Wholesale Pharmaceuticals." *The Journal of Industrial Economics*, 58(1): 32–53.

eMarketer. 2018. "US Media and Entertainment Industries StatPack 2018: Digital Ad Spending Forecast and Trends."

EU Commission. 2017. "Platforms to business relations in online platform ecosystems." *JRC Digital Economy Working Paper 2007-07.*

Galbraith, John Kenneth. 1952. *American capitalism: the concept of countervailing power.* Houghton Mifflin, Massachusetts.

Gaynor, Martin, and Robert J. Town. 2012. "Competition in Health Care Markets." *Handbook of Health Economics* Vol. 2, Chapter 9, 499–637. Elsevier.

Golden, Joseph, and John Horton. 2018. "The Effects of Search Advertising on Competitors: An Experiment Before a Merger." mimeo.

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

Gomes, Renato, Nicole Immorlica, and Evangelos Markakis. 2009. "Externalities in keyword auctions: An empirical and theoretical assessment." 172–183, Springer.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry." *American Economic Review*, 105(1): 172–203.

Gutierrez, German, and Thomas Philippon. 2017. "Declining Competition and Investment in the U.S." National Bureau of Economic Research Working Paper 23583.

Harrington, Joseph. 2018. "How Do Hub-and-Spoke Cartels Operate? Lessons from Nine Case Studies." *Competition Policy International*, September.

Hastings, Justine S., and Richard J. Gilbert. 2005. "Market Power, Vertical Integration and the Wholesale Price of Gasoline." *The Journal of Industrial Economics*, 53(4): 469–492.

Hoberg, Gerard, and Gordon Phillips. 2016. "Text-based network industries and endogenous product differentiation." *Journal of Political Economy*, 124(5): 1423–1465.

IAB. 2018. "Internet Advertising Revenue Report, Full Year 2017."

ISBA. 2020. "Programmatic Supply Chain Transparency Study." *PwC Report for the Incorporated Society of British Advertisers.*

Jeziorski, Przemyslaw, and Ilya Segal. 2015. "What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising." *American Economic Journal: Microeconomics*, 7(3): 24–53.

Johnson, Garrett, Randall A. Lewis, and David H. Reiley. 2017. "When Less is More: Data and Power in Advertising Experiments." *Marketing Science*, 36(1): 43–53.

Kelley, Lawrence A, Stephen P Gardner, and Michael J Sutcliffe. 1996. "An automated approach for clustering an ensemble of NMR-derived protein structures into conformationally related subfamilies." *Protein Engineering, Design and Selection*, 9(11): 1063–1065.

Kitts, Brendan, Parameshvyas Laxminarayan, and Benjamin Leblanc. 2005. "Cooperative Strategies for Search Auctions." First International Conference on Internet Technologies and Applications.

Lewis, R. A., and J. M. Rao. 2015. "The Unfavorable Economics of Measuring the Returns to Advertising." *Quarterly Journal of Economics.*

Mansour, Yishay, S. Muthukrishnan, and Noam Nisan. 2012. "Doubleclick Ad Exchange Auction." *Computing Research Repository*, abs/1204.0535.

McAfee, R. Preston. 2011. "The Design of Advertising Exchanges." *Review of Industrial Organization*, 39(3): 169–185.

Mullan, Hugh, and Natalie Timan. 2018. "Strengthening buyer power as a solution to platform market power." *CPI Antitrust Chronicle.*

O'Brien, Daniel P. 2017. "Price-Concentration Analysis: Ending the Myth, and Moving Forward." *Working Paper.*

O'Brien, Daniel P., and Keith Waehrer. 2017. "The Competitive Effects of Common Ownership: We Know Less than We Think." *mimeo*.

OECD. 2017. "Algorithms and Collusion: Competition Policy in the Digital Age." http://www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. "Scikit-learn: Machine learning in Python." *Journal of machine learning research*, 12(Oct): 2825–2830.

Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. "GloVe: Global Vectors for Word Representation." 1532–1543.

Prat, Andrea, and Tommaso M Valletti. 2019. "Attention oligopoly." *Working Paper*.

Schmitt, Matt. 2017. "Do hospital mergers reduce costs?" *Journal of Health Economics*, 52(C): 74–94.

Silk, Alvin J., and Charles III King. 2013. "How Concentrated Is the U.S. Advertising and Marketing Services Industry? Myth vs. Reality." *Journal of Current Issues & Research in Advertising*, 34.

Snyder, Christopher M. 1996. "A Dynamic Theory of Countervailing Power." *The RAND Journal of Economics*, 27(4): 747–769.

Stiglitz, Joseph. 2016. "The new era of monopoly is here." The Guardian May, 13.

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

Villas-Boas, J. Miguel. 1994. "Sleeping with the Enemy: Should Competitors Share the Same Advertising Agency?" *Marketing Science*, 13(2): 190–202.

Villas-Boas, Sofia Berto. 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *The Review of Economic Studies*, 74(2): 625–652.

Weche, John P., and Achim Wambach. 2018. "The Fall and Rise of Market Power in Europe." ZEW Paper Series, Discussion Paper 18-003.

Werden, Gregory J., and Luke Froeb. 2018. "Don't Panic: A Guide to Claims of Increasing Concentration." *Antitrust Magazine*, forthcoming.

For Publication on the Authors' Web Pages

From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising

Web Appendix

A) Data Details

The data used in the paper come from several sources. First, from Redbooks we obtained data on advertisers, their MAs and their network affiliations. Access to the data is available at:

https://www.redbooks.com/

In order to benchmark the information on M&A contained in the Redbooks data we relied on the *Zephyr* dataset on M&A, IPO, Private Equity and Venture Capital by Bureau Van Dijk. Data accessible at:

https://www.bvdinfo.com/it-it/our-products/economic-and-m-a/m-a-data/zephyr

We complement the data on advertisers with information on bids, keywords and advertisers provided by SEMrush, the most important and renowned provider of SEM data and related services:

https://www.semrush.com/

We obtained the data on Click-Through Rate at the industry/month level – position by position – by AdvancedWebRanking, a provider of SEM services, whose data are available at:

https://www.advancedwebranking.com/

Finally, in order to proceed with the thematic clustering we used a pre-trained set of GloVe word vectors – more specifically, we used the Common Crawl, 840B tokens, 2.2 million words, 300d vectors – publicly available and open source at:

https://nlp.stanford.edu/projects/glove/

In Table A.1, we summarize all the main variables used in this study, reporting their source, frequency and their short description. While the Redbook data have been previously used in economics and marketing studies – see Dai [2014] for a recent example – to the best of our knowledge the SEMrush data are new to the literature. We therefore discuss in more details their nature and limitations.

SEMrush is a leading provider of sponsored search data and this is how we selected it for this study.[32] Importantly, this implies that the data that we use tend to be the same of that used by many players in

[32]SEMrush was launched in 2008. It gained and maintained a leadership position, frequently winning

Table A.1: Raw variables' description and sources

Variable Name	Source	Frequency	Definition
Semrush			
keyword	www.semrush.com	year/advertiser	The keyword bringing users to the website via search results – that is, the keyword advertisers bid on
position	-	year/keyword/advertiser	The position of the domain in paid search for the given keyword at the specified period
searchvolume	-	year/keyword	Number of search queries for the given keyword in the last 12 months
CPC	-	year/keyword/position	Average price advertisers pay for a users click on an ad triggered by the given keyword
traffic	-	year/keyword/advertiser	The share of traffic driven to the website with the given keyword for the specified period
competition	-	year/keyword	Competitive density of advertisers using the given term for their ads
results	-	year/keyword	The number of URLs displayed in organic search results for the given keyword
Redbooks			
enterprise_nbr	Redbooks files	year/advertiser	Advertiser's ID code
company_name	-	year/advertiser	Advertiser's business name
website	-	year/advertiser	Advertiser's website
agency_ID	-	year/advertiser	Digital Marketing Agency (MA) ID code - possibly with multiple matches per advertiser
agency_name	-	year/advertiser	Digital Marketing Agency (MA) business name
digital	-	year/agency	Indicator function for digital agency
parent_ent	-	year/agency	Agency owner ID code - mainly belonging to 7 networks
industry	-	year/advertiser	Core business industry of the advertiser
Advanced Web Ranking			
CTR	www.advancedwebranking.com	month/industry/position	Click-through rate: average number of clicks per impression
GloVe			
key_vec	nlp.stanford.edu/projects/glove/	keyword tokens	Set of GloVe vectors pre-trained on Common Crawl, 840B tokens, 2.2 million words, reported in 300 dimensions

Notes: summary of the raw variables that we use in the paper. We report the variable name, the data source, the raw frequency – as used for the analysis – and a brief description.

this market to set their strategies. Data like those we obtained from SEMrush represent a way to have an overview of the entire market – like those that the internal data from search engines would give – but without the limitations that might be posed by using the internal records of search engines in terms of advertisers' identities and prices.[33] A limitation of the data is, however, the non-fully transparent way that the yearly averages are calculated: proprietary algorithms are used to aggregate data from multiple providers and assemble the SEMrush data. As typical in this industry, Google's Keyword Planner is a key source for accessing CPC data which would be otherwise not observable.[34]. Although Google's Keyword Planner itself

awards as a top SEO and SEM tool in the last years, including best SEO suite 2017, "US & UK search awards" and "European Search Awards."

[33]To the best of our knowledge, no published study using internal search engines' data contains this type of information.

[34]This is typically done programmatically using the services of the like of TargetingIdeaService API,

does not report the exact algorithms used to calculate the average CPC, its data are accurate and all the rich dynamics that might characterize bidding on a keyword throughout an year should contribute to the formation of the average. Averaging, while leading to some information loss, is needed to form an overall view of such an highly dynamic and fragmented market. In our study, this is made even more necessary by the yearly nature of the Redbooks data. Although these are important limitations, the data that we use are likely representative of those available to many advertisers and intermediaries and are of comparable quality and extent to what might be available from other publicly accessible sources.[35]

Regarding the CTR data that we use, there are a few limitations worth discussing. In particular, since we lack keyword-level click through rates, we impute this from a market average using data from Advanced Web Rankings. However, the research question in this study involves structures that are more aggregate than individual keywords, thus an aggregation is unavoidable.[36] Moreover, keyword-specific CTRs are in most cases useless as they all are just zeros for the obvious reason that most keywords are infrequently searched and even less frequently generating any click. Hence, using CTRs typically requires substantial aggregation across large sets of keywords and/or over long period of times. In appendix F below, we return to the issue of the reliability of our CTR measure by evaluating the robustness of our estimates to measurement errors in the CTR.

Finally, regarding the mergers, we first identified the changes in the MAs' affiliation to networks in the Redbooks data. Then, we inspected each of these episodes using the Zephyr data (Bureau Van Dijk) to confirm that there was an M&A operation where the MA control passed to the network. In Table A.2 we report the full set of mergers observed in the data. For each case, it reports the name of the acquiring network, the number of advertisers linked to the agency at the time of the acquisition, the number of industries in which they operate and number of markets. Across networks, there is heterogeneity both in the number and the size of the MAs acquired. While Dentsu-Aegis appears to be the most "active" network with 8 acquisitions – including the one with most clients, Merkle, – WPP secured the largest acquisition in terms of presence in the markets (*SHIFT Communications* with clients active across 1,049 different markets).

see `https://developers.google.com/adwords/api/docs/guides/targeting-idea-service`. SEMrush's CPC is an average of the past 12 months and is updated on a monthly basis. See: `https://www.semrush.com/kb/162-monthly-numbers`

[35]Indeed, to further ensure that we were not missing on some important (and possibly more disaggregated) data, we compared our SEMrush data to what available from SpyFu, one of the main SEMrush' s competitors. We found that the variables available are essentially the same, but that the CPC data is reported in a more informative way on SEMrush than on SpyFu: SEMrush reports the CPC across all the keyword's advertisers, while SpyFu reports that associated with being (on average) in the second position.

[36]If different intermediaries were representing clients that, despite operating in the same industry, were facing systematically different CTRs conditional on position-year, then aggregation would be problematic. However, this appears as an unlikely situation because all advertisers, apart from operating in the same industry, are also all large firms active in online advertising in the US market.

Table A.2: M&A Operations across All Networks, 2014-2017

Agency	Acquiring Network	Acquisition year	Number of Advertisers	Number of Industries	Number of Markets
The Brooklyn Brothers	IPG	2016	6	2	23
Essence Digital Limited	WPP	2015	1	1	145
Quirk	WPP	2015	5	2	272
SHIFT Communications	WPP	2017	13	8	1,049
Deeplocal Inc.	WPP	2017	5	1	117
Maruri GREY	WPP	2017	1	1	150
Zubi Advertising Services, Inc.	WPP	2017	3	2	345
Campfire	Publicis	2015	3	1	27
La Comunidad	Publicis	2015	9	5	271
Sapient Corporation	Publicis	2015	17	6	1,038
Blue 449	Publicis	2016	4	2	93
Forsman & Bodenfors	MDC	2017	5	1	315
Formula PR	Havas	2015	6	4	309
FoxP2	Dentsu-Aegis	2015	1	2	42
Rockett Interactive	Dentsu-Aegis	2015	1	1	22
Covario, Inc.	Dentsu-Aegis	2015	3	1	78
Achtung	Dentsu-Aegis	2016	2	1	226
Gravity Media	Dentsu-Aegis	2016	5	3	433
Grip Ltd.	Dentsu-Aegis	2016	3	2	92
Merkle	Dentsu-Aegis	2017	18	7	973
Gyro	Dentsu-Aegis	2017	12	6	363

Notes: the table reports the set of acquisitions 2014-2017 by the networks. To identify these events, we used Redbooks data and confirmed them through Zephyr data (Bureau Van Dijk). The table only reports acquisition involving at least 51%+ of the acquired agency. Acquisition prices are typically not disclosed. Exceptions are the cases of *Sapient Corporation*, acquired for $3.7 billion by Publicis Groupe, and *Merkle* acquired for $1.5 billion by Dentsu-Aegis in 2016. Furthermore, not listed in the table are two divestitures cases: TM Advertising and Moroch returned independent by buying themselves back from the networks.

B) Redbooks Industries and Imputation

Figure B.1 reports the primary industry on the y axis, and the number of advertisers on the x axis (dark grey) as reported in the Redbooks data: 23 different macro-sectors are represented, with the three largest ones being *Media*, *Industrial* and *Financial services*. The number of advertisers ranges from a handful (*Tobacco* and *Telecom*) to several hundreds. For a third of its advertisers, however, Redbooks does not report the information on industry affiliation; hence, we exploited SEMrush data to impute it. In particular, we matched all keywords by advertisers without a reported industry with the keywords by all advertisers for which this information is available: advertisers with a missing industry are then assigned to the industry with which they share most keywords. Light grey bars in the figure indicate the number of imputed advertisers, per industry.

Figure B.1: Number of Advertisers per Industry: Redbooks data

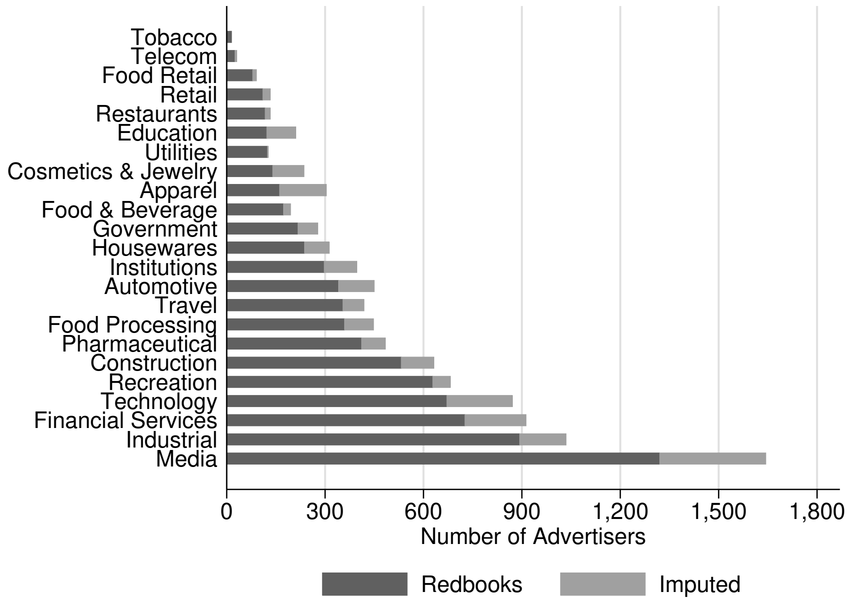

Notes: Horizontal bars report the number of advertisers reported in Redbooks (dark grey) and imputed through SEMrush keyword data (light grey), for each industry.

C) Vector Representation and Clustering

We proceed in generating vector representations of the keywords by splitting the keywords in our sample, term by term, and by merging them with the GloVe pre-trained set of words. More specifically, we split each keyword $k \in [1, .., K]$ into its consitutent terms $t_k \in [1, .., T_k]$, where T_k is the number of terms in the k^{th} keyword. After stemming, we then matched each term with the corresponding GloVe term $t_g \in [1, .., G]$ – in our application $G \approx 2.2$ million, and each t_g is a vector in $J = 300$ dimensions. Each vector locates the term/keyword into the GloVe vector space, which is a sub-structure of the classic word-word co-occurrence matrix ([Pennington, Socher and Manning, 2014]). For each keyword, we generate a single vector in J dimensions by summing up all the T_k vectors. If any term was not matched with the GloVe pre-trained sample (it covers $\approx 80\%$ of the terms in our sample), we input a vector of zeros, which does not impact the total sum.

The resulting vector representation ($\overrightarrow{\boldsymbol{d_k}}$) of the K keywords read:

$$\overrightarrow{\boldsymbol{d_1}} = (d_{1,1}, d_{2,1}, ..., d_{J,1}),$$
$$\vdots$$
$$\overrightarrow{\boldsymbol{d_k}} = (d_{1,k}, d_{2,k}, ..., d_{J,k}),$$
$$\vdots$$
$$\overrightarrow{\boldsymbol{d_K}} = (d_{1,K}, d_{2,K}, ..., d_{J,K}),$$

Step 1: for each industry defined by Redbooks we run a spherical k-means algorithm ($k = 1,000$ in the baseline model) on the matrix of vectorized keywords in order to group them according to their Euclidean distance.[37] Hence, through the first layer of the algorithm we are able to capture the similarities between keywords (i.e., their "distance" in GloVe terms) and make the underlying semantic themes emerge from the data structure itself. The well-known drawback of the k-means algorithm, though, is that the number of clusters is pre-specified and might not reflect the "real" number of topics; in order to address the issue, we run several checks on clustering quality – and we show the robustness of the results to different choices of K.

Step 2: we add a second clustering layer exploiting the structure of the competition within the thematic clusters. More specifically, for each cluster c, we build a $K_c \times N_c$ sparse matrix, whose rows correspond to the keywords in the cluster, and whose columns match the advertisers which, at least once in the data, have participated in one of those keyword auctions – panel B in table C.1. The resulting row vectors, akin term vectors in text analyses, are projections of the keywords in the space span by the advertisers – i.e., the competitive structure space. The underlying assumption is that keywords showing similar patterns of bidders are more likely to belong to the same competitive space, and that the latter has substantial overlaps with the – unobserved – product space. In order to exploit the keyword similarity, we build a matrix of pairwise Euclidean distances among the keywords, in terms of co-occurrency – panel C in table C.1. Each non-diagonal cell a_{ij} represents the distance between keywords i and j, computed with the L2 norm $d()$, that is

$$a_{i,j} = d(\overrightarrow{\boldsymbol{k^i}}, \overrightarrow{\boldsymbol{k^j}}) = \sqrt{\sum_{v=1}^{N_c} (k_v^i - k_v^j)^2}$$

where N_c is the number of advertisers in cluster c.[38]

[37] In the code, we use the standard python libraries `nltk` [Bird, Klein and Loper, 2009] and `sklearn` [Pedregosa et al., 2011], which feature function for NLP and unsupervised clustering.

[38] In the code, we use the `R` base functions `dist` and `hclust` (package `stats`).

Table C.1: Layer 2 clustering: data preparation

Keyword	Advertiser
key 1	Adv 1
key 1	Adv 2
key 1	Adv 3
key 2	Adv 2
key 3	Adv 2
key 3	Adv 3

A. Actual Data

⇒

	Adv 1	Adv 2	Adv 3
key 1	1	1	1
key 2	0	1	0
key 3	0	1	1

B. Advertisers' co-occurrency

⇒

	key 1	key 2	key 3
key 1	0	$\sqrt{2}$	1
key 2	$\sqrt{2}$	0	1
key 3	1	1	0

C. Keyword distance metric

Notes: data preparation for layer 2 clustering. For each thematic cluster, from the keyword auction data listing keywords and advertisers (panel A) we build a matrix of advertisers' co-occurrency (panel B). Through that, we can compute the pairwise Euclidean distance between keyword vectors in the advertisers' space and build the distance matrix (panel C).

Finally, we select the best-fitting definition of competitive clustering through a hierarchical clustering algorithm run on the distance matrix. In order to optimally prune the cluster tree we employ the Kelley, Gardner and Sutcliffe [1996] penalty function.

Figure C.1: Hierarchical clustering

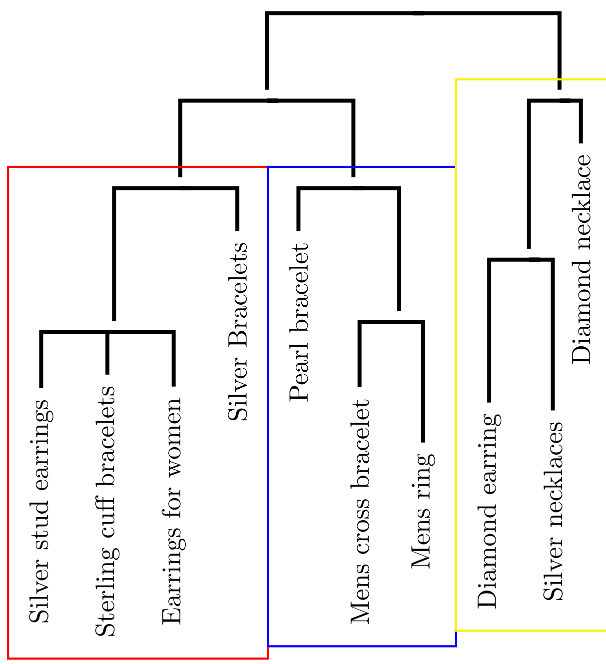

Notes: Graphical representation of the structure of competitive clusters. The three clusters – red, blue, and yellow boxes – are identified by optimally pruning the thematic clusters through the Kelley, Gardner and Sutcliffe [1996] penalty parameter applied to the keyword distance matrices.

D) Cluster Validity

In order to test the reliability of the clustering exercise, we implemented a task to validate them. With no training samples needed, we relied on human intervention only at the very end of the funnel – i.e., we checked the "quality" of the clusters ex-post by designing a series of simple tasks that we submitted to human testers. More specifically, for each cluster $c \in [1, .., 1,000]$, within industry i, we randomly picked a reference keyword $refK_{ci}$ and two test keywords $testK_{ci}$ and $testK_{-ci}$, from c and from one of the other clusters in i, respectively. Figure D.1 is a graphical representation of the task we submitted to the human testers: the user is asked whether, given that she searched for $refK_{ci}$, she would be more likely to search for $testK_{ci}$ or $testK_{-ci}$, or none of them. The task yields three potential outcomes: i) the user chooses $testK_{ci}$ (*success*), ii) the user chooses $testK_{-ci}$ (*failure*), and iii) the user cannot choose either option (*no answer*).

Figure D.1: Cluster Quality Checking Task

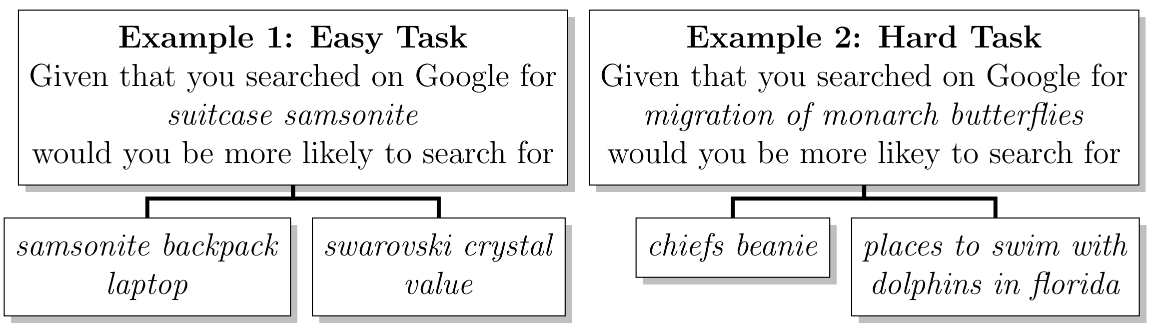

Notes: Amazon Mechanical Turk task representation. First, the user is given a reference keyword belonging to cluster c (*suitcase samsonite* in Example 1) which she is supposed to have searched on Google. Then, she is asked to identify out of two additional keywords which of the two she considers more likely to be searched given the initial search. One of the two keywords proposed belongs to cluster c (*samsonite backpack laptop* in Example 1), while the other belongs to the same industry but to a different cluster (*swarovski crystal value* in Example 1). Example 2 is analogous, but representative of a more difficult case for the tester.

The question is designed to check whether the keyword links emerging from the thematic clustering are effectively mimicking the users' behavior when surfing the web. In the figure, example 1 is an "easy task" – from the *Apparel* industry – and had a very high hit rate in the test: the presence of the brand name within both $refK_c$ and $testK_c$ helps delimiting the market (and enhances the similarity, too). Example 2 is relative to the *Travel & Leisure* industry, and experienced a high rate of non-respondence by the testers: the underlying theme linking *migration of monarch butterflies* and *places to swim with dolphins in Florida* is the Florida Keys, which are both one of the destinations of monarch butterfly migrations and a renowned place to swim with dolphins. While this theme is known to real users, it has not been identified by most of our human testers, nonetheless GloVe correctly highlighted their similarity. We submitted the tests through *Amazon Mechanical Turk*, a marketplace for work that requires human intelligence. In table D.1 we report the share of successes, failures and no answers in a sample of industries. Our initial design of the test did not allow the user to skip answers (i.e., *No answer = 0* by design for the first five industries in the table); however, when subsequently we introduced the option we recorded an average of one third of non-responses. The success rate is consistently high and evenly distributed among industries. Moreover, it does not appear to be influenced by the rate of non-response.

Table D.1: *Amazon Mechanical Turk* Test

Industry	Answer		No Answer
	Success	Failure	
Technology	.84	.16	0
Travel & Leisure	.88	.12	0
Media	.82	.18	0
Food Processing	.89	.11	0
Micellaneous	.60	.40	0
Utilities	.90	.10	.30
Apparel	.89	.11	.32
Retail	.81	.19	.33
Industrial	.82	.18	.38

Notes: clustering test results on a subset of industries. For *Technology*, *Travel & Leisure*, *Media* and *Food Processing* we did not allow the user to leave the question blank.

E) Sample Selection

Following on from the discussion in section 6, we report in Figure E.1 how β^{IV} changes with the dimension of the analysis sample. Indeed, among the competitive clusters, many are composed of keywords that both contribute very little, or not at all, to the search engine's revenues and are never involved in any of the M&As that we exploit for the IV strategy. Therefore, we keep in the baseline analysis sample only markets that either experience variation in the instrument at least once during the sample period or, for the remaining ones, those that are in the top quartile of revenues. This leads us to drop markets that represent between 1% and 2% of the total yearly revenues. In Figure E.1, the baseline sample – with the corresponding IV estimate – is marked by the vertical dashed line. As this figure illustrates, after we drop less than 50% of the lowest revenue markets, the IV estimates become fairly constant and similar to the baseline ones. Dropping 50% (or less) of the lowest revenue markets corresponds to dropping less than 1% of the total yearly revenues. Thus, for the purpose of our analysis, we consider these small markets as not a valuable source of variation in the data, but rather as a source of noise that makes it impossible to detect the causal association between demand concentration and revenues. This is especially the case because these zero (or nearly so) revenue markets are often very small, possibly made up of one or very few keywords and, crucially, with a single advertiser bidding on them.

Figure E.1: Effects of Sample Selection on the IV Estimates

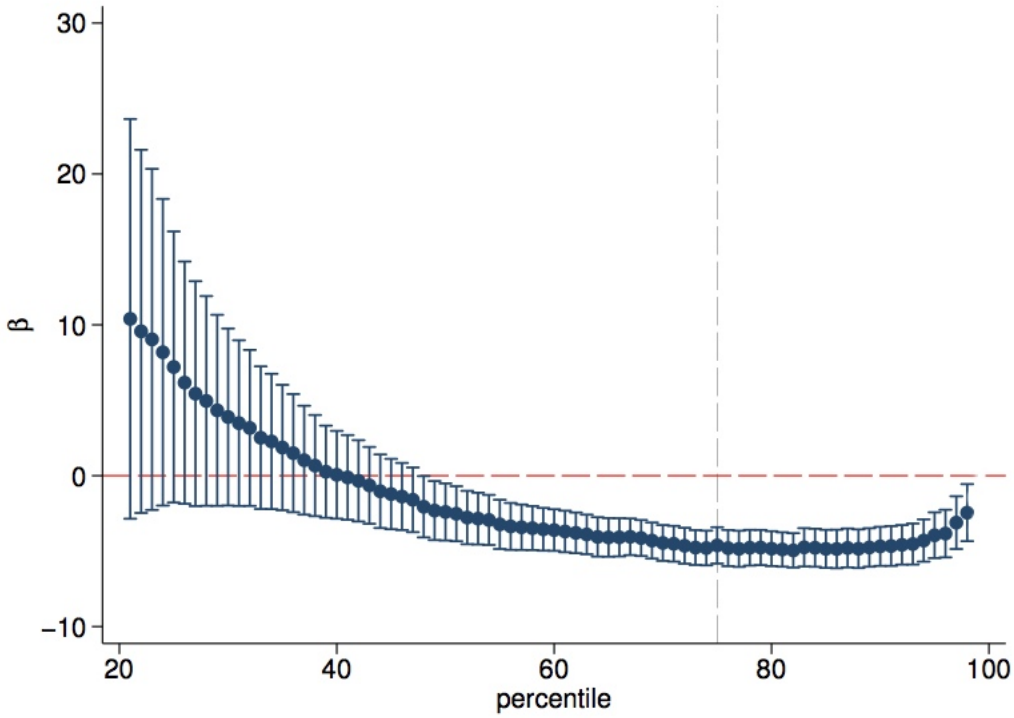

Notes: points estimates – blue dots – and their confidence intervals – blue caps – on samples of different sizes. The dotted gray line at the 75^{th} percentile marks the sample used in the baseline analysis.

F) Robustness Check: CTR Measurement Error

The CTR measure that we use presents a measurement error problem, as discussed both in the text and in section A of this appendix. In this section, we explore the robustness of our baseline estimates to this problem. In particular, we consider two sets of robustness checks: first, we exclude the CTR from the analysis by setting all CTRs to 1 and, second, we randomly re-match CTRs to keywords. The first exercise consists of estimating the same regression models presented in Table 4 using modified versions of the main variables: in the case of Table F.1, the CTRs are set 1 only for the dependent variable, while in the case of Table F.2, they are set to 1 for all variables whose calculation involves the CTR. To distinguish these modified variables from those used earlier, we indicate the former with an upper bar: $\bar{R}_{mt}$ is thus R_{mt} recalculated without CTRs.[39] The reason why it is interesting to present the two sets of estimates in Table F.1 and F.2 is that estimating the effect of HHI_{mt} on $\bar{R}_{mt}$ can also serve as a check of the robustness of our analysis to an alternative measure of the revenues: an upper bound on the revenues attainable when

[39] $\bar{R}_{mt} = \sum_{k \in K_m} CPC_{kmt} * Volume_{kmt}$, $\bar{s}^i_{mt} = \frac{1}{\bar{S}_{mt}} \sum_{a \in A_i} \sum_{k \in K_m} Volume_{kt}$, and $H\bar{H}I_{mt} = \sum_{i=1}^{I} (\bar{s}^i{}_{mt})^2$.

all ads generate the same number of clicks per time. In any case, the estimates in both Table F.1 and F.2 are quite close to each other and also close to the estimates in Table 4 in the text, although systematically smaller. For instance, relative to our benchmark estimated effect of an 11.32 percent drop in revenues, the corresponding estimate in Table F.1 indicates a drop of 8.54 percent and that in Table F.2 a drop of 8.77 percent. Thus, the qualitative implications of our analysis are robust to this type of alternative use of CTRs.

In the second set of robustness checks, we consider a different approach aimed at assessing how the variation of our CTR measure across markets might impact our findings. We proceed by setting up a bootstrap procedure that, at the beginning of each repetition, randomly assigns to each keyword a vector of industry-year CTRs (i.e., the CTRs of positions 1 to 11 for the specific industry-year, from AWR data) that is drawn (with replacement) from the whole set of industry-year CTR vectors in the data. Then we calculate the baseline estimate (corresponding to the model of column 9 of Table 4). Figure F.1 reports the IV estimates obtained on 500 samples with the block-bootstrapped CTR data. The figure reports each repetition on the x-axis. On the y-axis, it reports the estimates: the point estimate (red solid square), and the 95% confidence interval (blue spikes). The dashed white line marks the baseline $\hat{\beta}^{IV}$ (from column 9 in Table 4), whereas the dashed grey line, $\hat{\beta}^{boot}$, reports the average bootstrapped $\hat{\beta}^{IV}$. Although there is variability in the estimates across the 500 repetitions, all point estimates are close to the baseline estimate. The average estimate across the samples, $\hat{\beta}^{boot}$, is in fact very close to $\hat{\beta}^{IV}$. Furthermore, $\hat{\beta}^{IV}$ always falls within the 95% confidence interval of the bootstrapped parameter, while zero (or positive values) are never contained. Therefore, the results in Figure F.1 confirm the robustness of the main estimates in the text.

Table F.1: Effect of Concentration on Search Engine Revenues - $log(\bar{R})$

	OLS					IV				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\hat{HHI}$	-2.314***	-2.278***	-2.228***	-2.220***	-2.215***	-12.07***	-5.012***	-3.495***	-3.483***	-3.456***
	(0.0646)	(0.0529)	(0.0525)	(0.0524)	(0.0525)	(3.682)	(0.846)	(1.075)	(1.079)	(1.079)
Organic Results (billion)				0.261***	0.252***				0.238***	0.232***
				(0.0582)	(0.0564)				(0.0595)	(0.0577)
Keywords Characteristics										
Branded Keyword					-0.0406					-0.00740
					(0.0534)					(0.0608)
Long-tail Keywords					-0.121***					-0.0993**
					(0.0357)					(0.0403)
R^2	0.07	0.62	0.63	0.63	0.63	-1.25	0.57	0.62	0.62	0.62
Observations	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476
Cluster FE		✓	✓	✓	✓		✓	✓	✓	✓
Year FE			✓	✓	✓			✓	✓	✓

Notes: the dependent variable is the (log) revenues, $\bar{R}_{mt}$. Columns (1) to (5): OLS estimates, with an increasing number of fixed effects and controls. Columns (6) to (10): IV estimates – where we instrumented HHI_{mt} with the merger-induced change in concentration. In all models the standard errors are clustered at the thematic clusters level.

Table F.2: Effect of Concentration on Search Engine Revenues - $log(\bar{R})$ on $\bar{HHI}$

	OLS					IV				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\bar{HHI}$	-2.347***	-2.252***	-2.188***	-2.181***	-2.176***	-11.45***	-4.931***	-3.592***	-3.580***	-3.544***
	(0.0648)	(0.0536)	(0.0533)	(0.0533)	(0.0533)	(3.303)	(0.825)	(1.081)	(1.084)	(1.081)
Organic Results (billion)				0.262***	0.253***				0.236***	0.232***
				(0.0577)	(0.0561)				(0.0586)	(0.0570)
Keywords Characteristics										
Branded Keyword					-0.0203					0.0296
					(0.0533)					(0.0662)
Long-tail Keywords					-0.119***					-0.0921**
					(0.0357)					(0.0412)
R^2	0.08	0.62	0.63	0.63	0.63	-1.06	0.57	0.61	0.61	0.61
Observations	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476	52,476
Cluster FE		✓	✓	✓	✓		✓	✓	✓	✓
Year FE			✓	✓	✓			✓	✓	✓

Notes: the dependent variable is the (log) revenues, $\bar{R}_{mt}$. Columns (1) to (5): OLS estimates, with an increasing number of fixed effects and controls. Columns (6) to (10): IV estimates – where we instrumented $\bar{HHI}_{mt}$ with the merger-induced change in concentration. In all models the standard errors are clustered at the thematic clusters level.

Figure F.1: CTR Bootstrap repetitions

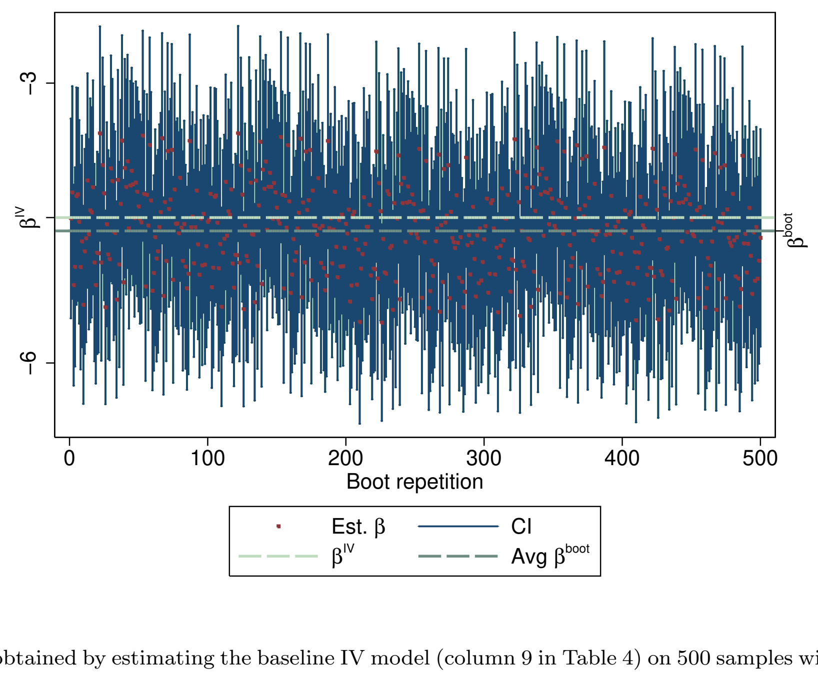

Notes: IV estimates obtained by estimating the baseline IV model (column 9 in Table 4) on 500 samples with block-bootstrapped CTR data. For each industry-year, we draw (with replacement) the distribution of CTR - positions 1 to 11 - from AWR data, then randomly merge them to the SEMrush data before aggregating at the market level, and run the estimation. For each repetition, reported on the x axis, we plot the point estimate (red solid square), and the 95% confidence interval (blue spikes). The dashed white line marks the baseline $\hat{\beta}^{IV}$, whereas the dashed grey line reports the average bootstrapped $\hat{\beta}^{IV}$.

G) Falsification

In a setting like the one analyzed, it seems useful to visualize changes in the outcome variable before and after an acquisition-driven change in concentration. In figure G.1, we report a graph built as in Dobkin et al. [2018].[40] Specifically, in order to show the impact of the mergers on total revenues, we first build indicator variables for time relative to the event at the market level (i.e., time from the first M&A which involved any MAs in the competitive cluster), then we estimate a nonparametric event study of the form:

$$log(R_{mt}) = \alpha + X_{mt}\gamma + \sum_{r=-3}^{-2} \mu_r + \sum_{r=0}^{2} \mu_r + \varepsilon_{mt}$$

where X_{mt} are market-level controls and μ_r are the coefficients on the relative time indicators (i.e., the key coefficients plotted in the figure, alongside their pre-merger linear trend – the dotted line). The vertical, dashed grey line indicates the first year after the merger. The upward sloping, dashed back line is the linear fit in the pre-merger period (as the figure suggests, the fit approximates these data quite well). The full dots are the period averages, while the hollow dots indicate the standard errors. There is rather clear graphical evidence: the drop in the average revenues post-merger indicates a negative association between the post-merger period and the log revenues, which is consistent with the estimates in the paper.

Figure G.1: Impact of Mergers on $log(R)$

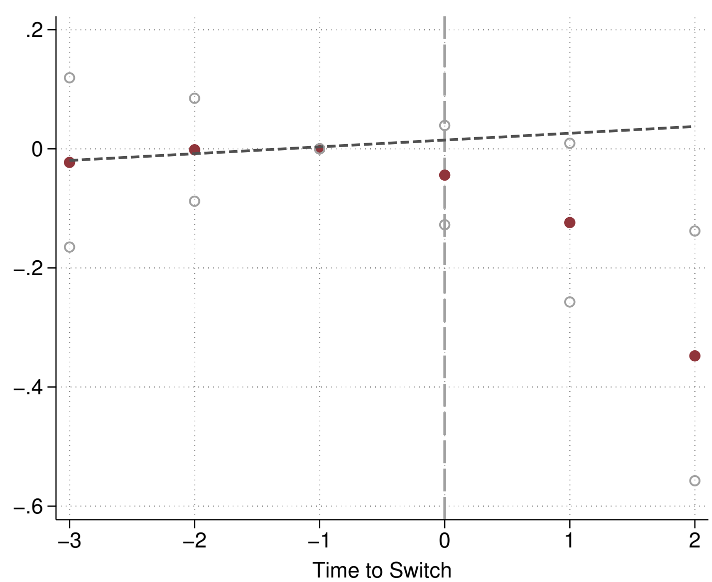

Notes: full dots are the averages of the time indicators' point estimates in a nonparametric event study estimation of $log(R)$ on M&A events, while the hollow dots indicate their confidence intervals. The upward sloping, dashed back line is the linear fit of the pre-merger years, projected on the post-merger period.

[40] `https://pubs.aeaweb.org/doi/pdfplus/10.1257/aer.20161038`

H) Mechanisms: Segmentation by Keywords and by Branded Keywords

In this section, we present additional material regarding the issue of the mechanisms through which concentration of intermediaries lowers the CPC. In particular, we explore two aspects related to market segmentation via the division of keywords. The first result that we present is the representation in Figure H.1 and complements the keyword-level descriptive evidence presented in section V. As explained there, for six large mergers involving different networks, we look at whether, after the acquisition of a MA by a network, there is any change in the overlap in the sets of keywords of the clients of either the network or the acquired MA. If the overlap declines, it might indicate that the intermediary splits the market by keywords, while if it stays identical (or grows) it might indicate that most of what the intermediary does takes place within-auctions. As discussed in the text, the evidence in Figure H.1 is suggestive that both strategies are adopted, although to different extents across the seven networks.

Figure H.1: Venn Diagram: all mergers

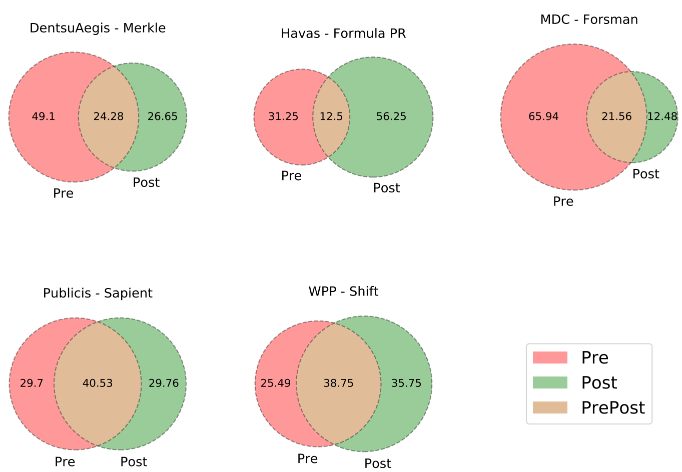

Notes: share of coalition keywords – i.e., keywords bid by both the advertisers in the acquired MA and those in the acquiring network – before and after the merger. Shares are computed on the overall number of coalition keywords. “Pre” is the share of keywords in coalition in the year before the merger only; similarly, “Post” refers to the share of keywords in coalition only in the year after the merger, and “PrePost” are keywords in coalition both before and after.

The second result is more specific and concerns branded keywords. Advertisers spend significant portions of their marketing budgets on branded keywords: these are related to both their own brand and to their rivals’. Among the feasible coordination strategies, keyword splitting represents the easiest way to fully segment the market. Direct coordination by advertisers to stop bidding on each others brands, though, is unlawful. But the same bidding pattern would be legal if autonomously implemented by a network representing rival advertisers. Hence, from the advertisers' viewpoint, a hub-and-spoke model of coordination through network intermediaries might be two-fold optimal: on the one hand, it lowers keyword-level costs by decreasing auction competition; on the other hand, it guarantees lower marketing costs by preventing brand competition.

We test the implications of the above conjecture in our data: we find that after an M&A event, the merged agency's clients show a marked decline in the probability of an advertiser appearing on the branded keywords of same-network competitors, and an increase in the probability for the own branded keywords. In order to do that, we formally define the brand bidding, and we build a *branded* indicator variable for all keywords that contain one or more words related to a brand (e.g., the keyword "Volkswagen beetle" would be branded, given that it contains the brand "Volkswagen").

Figure H.2: Changes in Own and Other Brand Keywords

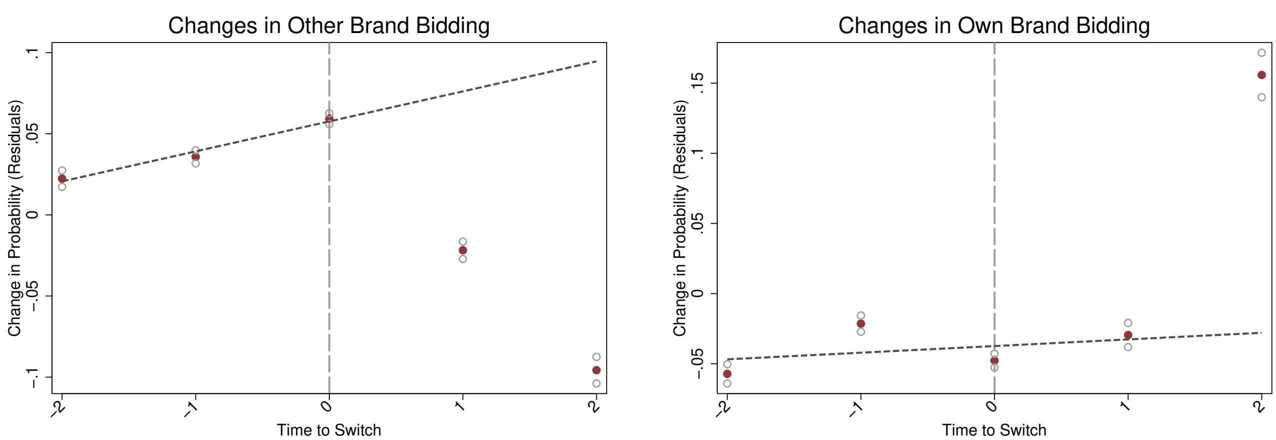

Notes: full dots are the demeaned values of $\Delta p(other\_branded)$ (left panel) and $\Delta p(own\_branded)$ (right panel) plotted against *timetoswitch* – i.e., distance in years from the merge (t^*, represented by the dashed vertical line). The hollow dots indicate standard deviations. The upward sloping, dashed back line is the linear fit of the pre-merger years, projected on the post-merger period.

We also define a few additional variables at different aggregation levels:

1. Keyword-auction level. Within the *branded* keywords we further distinguish two subdomains, depending on advertisers' identity: *own_branded* is an indicator for advertisers bidding on keywords related to their own brand ("Volkswagen beetle" when the advertiser is Volkswagen); *other_branded* indicates whether an advertisers bids on a keyword whose related brand is not its own ("Volkswagen beetle" when the bidder is Ford Motor Company);

2. Market-year level. For each agency-year pair (j, t) we define the variable *timetoswitch* as the distance – in years – to the relative M&A event (t^*). When aggregating at the market level, we use the *first* recorded event as the reference point in the definition of *timetoswitch*. We also define the indicators for the presence of branded, own branded and other branded at the market/year level (*dbranded*, *dother_branded* and *down_branded*, respectively);

3. Time-to-switch level. We aggregate the probability of being branded (total, own and other) at the time-toswitch level. We also compute the yearly change in probability as $\Delta P(branded)_t = \frac{(shareBranded_t - shareBranded_{t-1})}{shareBranded_{t-1}}$. Finally, in order to ensure the comparability of all measures, we de-meaned them.

In figure H.2, we apply the method by Dobkin et al. [2018] to describe the change in probability for both the other branding (left panel) and own branded (right panel) against *timetoswitch*; the dashed vertical line represents t^*. We also add the linear fit, estimated in the period *before* the treatment and projected in the post-treatment period. The quasi-linear increasing pattern of other brand bidding is clearly impacted negatively by the M&A event (left panel), with advertisers significantly bidding less on rivals' brands after the merger; on the other hand, the own brand bidding appears not to be impacted, and instead records a very neat upward jump at $t^* + 2$ (right panel).

I) Industry Heterogeneity

In Figure I.1, we explore differences among industries by showing the distribution of $\hat{\beta}_{IV}$ estimated at the industry level, for *log*(*cpc*) (left panel) and *log*(*#keywords*) (right panel). Although negative on average, the former features positive values for one sector, *Agriculture*. The estimated effect of concentration on changes in the number of keywords shows a higher degree of noise, with most industries characterized by an imprecisely estimated zero effect. A positive impact, however, is clear for three important industries, *Automotive*, *Recreation* and *Retail*. The resulting picture suggests that networks, and MAs, follow different strategies depending on the market structure and competitive pressures within industries. The overall effect on revenues hence emerges from multiple, different paths.

Figure I.1: Industry-level IV estimates distribution

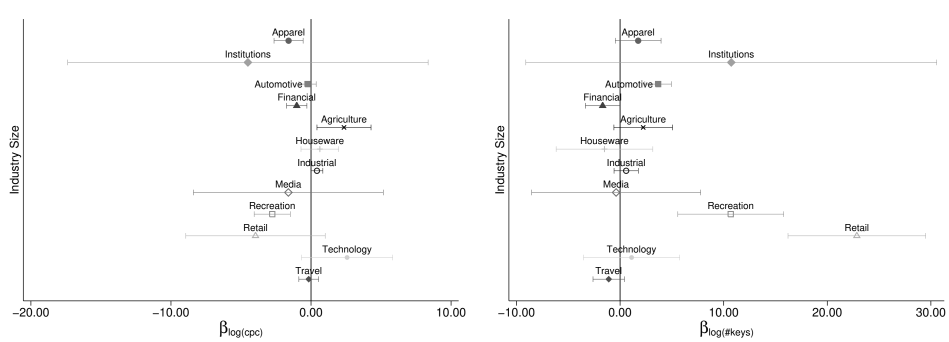

Notes: Industry-level IV estimates of $\hat{\beta}_{IV}$ with different outcomes: *log*(*cpc*) (left panel) and *log*(*number of keywords*) (right panel). Industries are ranked according to their size in terms of total search volume, and each point estimate is reported alongside its standard error. Industries with estimates 30 times bigger than the average point estimate have been excluded in order to ensure plot readability.

J) Monotonicity Test of the Instrument

In this section, we report the results of the instrument's monotonicity test proposed by [Angrist and Imbens, 1995]. Verifying that monotonicity holds is important because the sign of the first stage regression is theoretically unclear and, also, because splitting the market may create a negative relationship between HHI and simulated HHI over some of the latter's range. In fact, by instrumenting the HHI ($S^{\tilde{Z}}_{mt}$, $\tilde{Z} = [0, 1]$) with the merger-induced change in HHI (Z_{mt}), we are implicitly assuming that the merger effect is monotone – that is, either $S^1_{mt} \leq S^0_{mt}$ or $S^0_{mt} \geq S^1_{mt}$, $\forall m, t$. The assumption is not verifiable, but has testable implications on the CDFs of HHI for merged ($\tilde{Z}_{mt} = 1$) and unmerged markets ($\tilde{Z}_{mt} = 0$) – that is, they should never cross. In fact, if $S^1_{mt} \geq S^0_{mt}$ with probability 1, then $Pr(S^1_{mt} \geq j) \geq Pr(S^0_{mt} \geq j), \forall j \in \text{supp } S$. Figure J.1 plots the CDFs of markets subject (dashed red line) and not subject to any merger (solid green line). Since the two CDFs never cross, the instrument passes the test.

Figure J.1: Instrument Monotonicity Test

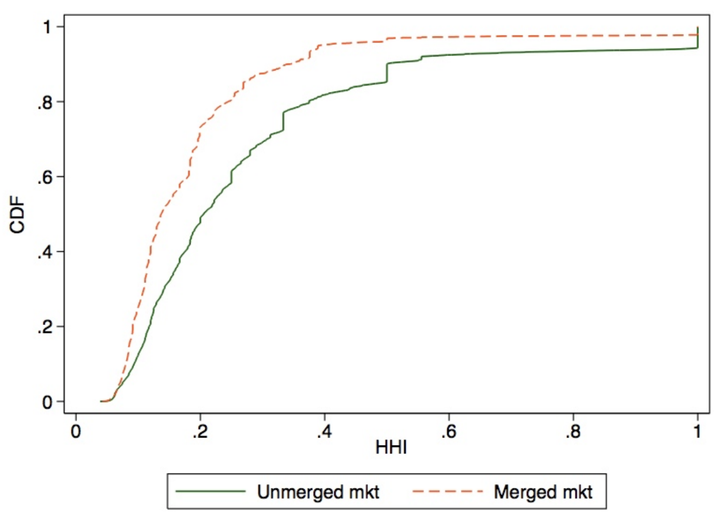

Notes: Instrument Monotonicity Test (Angrist and Imbens [1995]). By instrumenting the HHI ($S^{\tilde{Z}}_{mt}$, $\tilde{Z} = [0, 1]$) with the merger-induced change in HHI (Z_{mt}), we are implicitly assuming that the merger effect is monotone – that is, either $S^1_{mt} \leq S^0_{mt}$ or $S^0_{mt} \geq S^1_{mt}$, $\forall m, t$. The assumption is not verifiable, but has testable implications on the CDFs of HHI for merged ($\tilde{Z}_{mt} = 1$) and unmerged markets ($\tilde{Z}_{mt} = 0$) – that is, they should never cross. In fact, if $S^1_{mt} \geq S^0_{mt}$ with probability 1, then $Pr(S^1_{mt} \geq j) \geq Pr(S^0_{mt} \geq j), \forall j \in \text{supp } S$. The plot reports the CDFs of markets subject (dashed red line) and not subject to any merger (solid green line): indeed, they never cross.

K) Network Common Ownership and Competition

In this section, we provide the quantitative evidence to support the discussion in the conclusions about common ownership. Following Azar, Schmalz and Tecu [2018], we look for the presence of owners that are in common between the 5 publicly listed networks. In Figure K.1, for each of these 5 networks, we report the average ownership share in the 2010-2019 period for owners that, for at least 2 of the 5 networks, are among the 10 largest shareholders. Black Rock and Vanguard are among the top 10 shareholders for each of the 5 networks. The other investors are among the top 10 shareholders for 2 or 3 networks. It is important, however, not to overstate the meaning of this evidence on common ownership. As stressed in the main text, there are conflicting views on the extent of competition in the US advertising and marketing services agency industry, but the academic consensus is that the industry is reasonably competitive. In the text, we referenced Silk and King [2013], which is a landmark study on concentration in this industry. It reports a set of concentration measures for the various sector of advertising and marketing services industry (Tables 2, 4 and 5) along with addition measures that apply to the holding companies/networks, whose dominance has long be overstated (Table 6). In an earlier study (Silk and Berndt, 1994), evidence is presented that the industry's diversity and low level of concentration were consistent with MacDonald and Slivinski (1987) theory of the equilibrium structure of a competitive industry with multiproduct firms. King, Silk and Ketttelhohn (2003) investigated knowledge spillovers and externalities in the disagglomeration and growth of the advertising agency business. We found that a simple model of high demand, low wages, and externalities associated with clusters of related industries explained the dispersion of agency employment across states. Arzaghi, Berndt, Davis and Silk (2012) summarize a considerable body of stylized facts consistent with the market for advertising campaigns is contestable in the sense of Baumol et al. (1988).

Figure K.1: Common Ownership

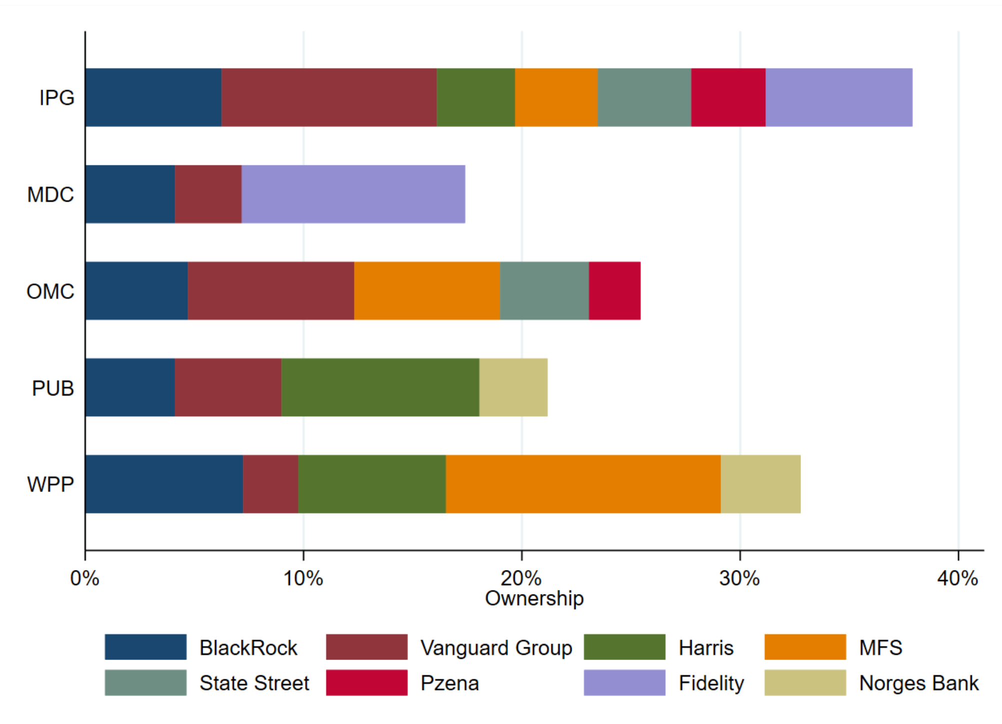

Notes: for the 5 networks that are publicly traded, the figure reports the average ownership share in the 2010-2019 period for owners that, for at least 2 of the 5 networks, are among the 10 largest shareholders. Black Rock and Vanguard are among the top 10 shareholders for each of the 5 networks. The other investors are among the top 10 shareholders for 2 or 3 networks. The data source is the Eikon dataset, `https://www.refinitiv.com/en/products/eikon-trading-software.`