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

Francesco Decarolis Gabriele Rovigatti

October 23, 2018

Abstract

This paper analyzes the impact of intermediaries' concentration on the allocation of revenues in online platforms. We study *sponsored search* - the sale ad space on search engines through online auctions - documenting how advertisers increasingly bid through an 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's keyword-auctions, we first apply machine learning algorithms to cluster keywords into thematic groups serving as relevant markets. Then, through an instrumental variable strategy, we quantify a negative and sizeable impact of intermediaries' concentration on platform's revenues.

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

JEL Classification: C72, D44, L81.

---

*Francesco Decarolis: Bocconi University, Department of Economics and IGIER. Via Reontgen 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@esterni.bancaditalia.it.* The authors thank the NET Institute for financial support. The authors would also like to thank their colleagues at their institutions as well and seminar participants at the EU DG Competition Economic Workshop, MaCCI Annual Conference, NET Institute Conference, NBER-Economics of Information Technology and Digitization Workshop and at the UK Competition and Markets Authority.

"*....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

The impact of technological change on the advertisement industry is well summarized in the opening quote. Advertising in the internet era is mostly about capturing the attention of consumers browsing the web and this requires both detailed data to tailor the ad to the right consumers (i.e., *targeting*) and fast algorithms to bid on the online auction platforms where ad space is sold. These needs have given rise to a major shift from advertisers' individual bidding to intermediated bidding by specialized and highly concentrated firms.

The resulting increase in buyer power is happening within an industry that is also clearly concentrated in the hands of a single firm: Google. Thanks to its technological innovations, Google has come to dominate the sale of online ad space becoming a leading example of those "superstar firms" at the center of the academic discussion on concentration and market power. In economics, as well as in the media, new evidence has been hotly debated regarding the increased concentration in important European and US industries.[1] The profound implications of industry concentration on firms' competition, workers' salaries and, ultimately, consumers' welfare explain this revived interest in a classical industrial organization topic.

This study contributes to two key issues of this debate by looking at the case of the online ad industry. The first is the quantification of concentration increases, with an emphasis on what is the proper use of industry data to identify markets. Starting from granular data on individual keywords, we propose a machine learning method based on state of the art natural language processing algorithms to cluster them into thematic groups representing markets. The second regards the effects of concentration. While recent studies look at impacts on competition, workers or consumers, we focus on a feature that has been so far overlooked: how downstream buyers respond to high concentration in an upstream industry. This is the old, but powerful idea of countervailing power that leads us to see the emergence of buyer power in the hands of intermediaries as a response to Google's concentration. Galbraith [1952] notoriously remarked that "the best and established answer to economic power is the building of countervailing power: the trade union remains an equalizing

[1] See, among others, Autor et al. [2017], De Loecker and Eeckhout [2017], 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].

force in the labor markets, and the chain store is the best answer to the market power of big food companies." A more recent, but no less egregious example in the case of the US healthcare is the insurers' introduction of HMOs and PPOs which is credited to have dramatically rebalanced power in favor of insurers after decades of hospitals' increased concentration [Gaynor and Town, 2012]. For online markets, these experiences have led national and international competition agencies to identify buyers' concentration as a possible solution to the imbalance of bargaining power in favour of platforms (EU Commission [2018], Mullan and Timan [2018]). This study is the first to quantify the effectiveness of this remedy.

But how and to what extent countervailing power can emerge? We look at the sale of ad space on search pages (*sponsored search*), which represents about half of all internet advertising revenues, or about $40 billion dollars in 2017.[2] This is a market that for nearly twenty years has been highly concentrated with Google earning a share between 75% and 80% of the total US search ad revenues in the period 2016-2018 [eMarketer, 2018]. Advertisers are the demand side. They seek attracting the attention of users querying search engines. To do so, they compete against other advertisers to buy one of a limited number of 'slots' available on the search engine result pages.[3] While in the early days of this market advertisers used to operate individually, most ad buying is currently taking place through intermediaries; in fact, they are involved in the slots sold in about 75% of Google's keywords in our data.

To understand why this shift can have profound implications on 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, DMAs) to which they give the mandate to bid on their behalf in the auctions; iii) most DMAs further delegate bidding to specialized entities which, for those DMAs belonging to an agency network (henceforth *network*),[4] are the network's agency trading desks. While thousands of DMAs operate in the market, essentially only seven networks are responsible for colletcing data and optimizing bidding algorithms for most advertisers. Their specialized activity can serve many useful purposes in terms of allowing for economies of scale and scope in the use of data and the implementation of automated bidding, but from the perspective of the search engine it can also trigger major revenue losses if the network lessens competition among its clients bidding on keywords pertaining to the same market.

Our data covers Google's search auctions involving 6,000 large US advertisers in the period 2014-2017. Seven networks bid for about 50% of all these advertisers, or for about 40% of the nearly 40 million keyword-advertiser combinations observed, and their concentration grows over time. The main objective of our empirical analysis is to quantify whether such increase in intermediaries' concentration affects Google' revenues and to what extent. Our strategy to find answers is based on three ingredients. First, a novel dataset built

[2]The Internet Advertising Board evaluates in $88 billion dollars the revenues in 2017 of the US internet ad industry, with the main tiers being sponsored search (46%), banner (31%) and video (14%) [IAB, 2018].

[3]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).

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

by combining multiple data sources: we have obtained from Redbooks - the most comprehensive database on marketing agencies - the list of the 6,000 largest US advertisers active in online marketing. For all these advertisers, Redbooks data give us the full list of DMAs affiliated with them, as well as the link of each individual agency to the network to which it belongs. We have combined these data on agencies with data on the Google's sponsored search auctions from SEMrush, a major data provider for DMAs. For all Redbooks advertisers, we know which keywords, if any, they bid on via Google. For each keyword and year, we know the position of the domain in the search outcome page, the volume of searches (i.e., the average number of search queries for the given keyword in the last 12 months); the visible URL of the ad; the content of the ad; and the keyword-specific average price advertisers pay for a user's click on the ad (Cost-Per-Click, or CPC).

The second element is the definition of relevant markets. We exploit the richness of our data to move from the advertisers' industry definition (23 industries, provided by Redbooks) to more granular clusters of keywords representing markets. By running an algorithm which features the use of a machine learning model with an unsupervised clustering technique, we identify about 23,000 markets. The approach involves two steps. We first vectorize the K keywords bid by Redbooks advertisers through a state-of-the-art natural language processing model (*GloVe*, by Pennington, Socher and Manning [2014*a*]) trained on more than 840 billion documents; we then use the distance between the resulting vectors to obtain a similarity measure that we use to cluster the keywords together. Albeit not in a strict antitrust sense, we can treat these groups as relevant markets because all cluster components share an underlying topic so that advertisers seeking to appear in searches involving this topic should compete on all keywords within the cluster.

The third element is an instrumental variable strategy. Instruments are needed because changes in intermediaries' concentration might be due to unobserved factors also affecting search engines' revenues. We address this problem by exploiting the variation in intermediaries' concentration driven by changes in the DMAs' affiliation to networks: in our sample period, 21 acquisitions and 2 divestures take place, affecting 6 out of the 7 networks. These episodes involve hundreds of markets for which they drastically alter the degree of concentration, but can also be considered as exogenous events to the extent that the local markets are too small to drive 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 growth of the search engine's revenues. Under our baseline IV model, a one standard deviation increase in concentration leads to a decrease of 0.42 standard deviations in the proxy of revenue changes, corresponding to an 8% drop. This effect appears to be robust to a number of sensitivity analyses, including different market definitions, the exclusion of outlier industries or major networks and a different set of controls used in model specifications. We complement the baseline analysis with several exercises aimed at identifying the main drivers of the changes in market revenues. On the one hand, the decline in revenues is in part due to changes in the set of chosen keywords - they tend to become more specialized, more precisely targeted and, as a result, less expensive on average;[5] on the other hand, we find a high degree of heterogeneity among industries, which reflects the variety of strategies followed by DMAs and networks 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 represent a relevant contribution 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. The first 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 second aspect, also well known, is that the profitability of even the most concentrated industry crucially depends on other forces, among which the degree of buyer power. The analysis illustrates how the market power of Google in the search auctions has been partially eroded by technological innovations and concentration among buyers. Thus, the old Galbraith [1952]'s idea of countervailing power still matters for one of the most modern and dynamic industries, confirming the relevance of recent policy recommendations [Mullan and Timan, 2018].

Second, the paper contributes to the understanding of a particularly complex, economically relevant and rapidly evolving market. The existing studies on online ad have mostly focused either on their effectiveness (see, among others, Blake, Nosko and Tadelis [2015], Golden and Horton [2018] and Johnson, Lewis and Reiley [2017]) or on the functioning of the selling mechanisms (see, among others, Edelman, Ostrovsky and Schwarz [2007], Varian [2007], Athey and Nekipelov [2014], Borgers et al. [2013], Balseiro, Besbes and Weintraub [2015] andCelis et al. [2015]). By focusing on the role of intermediaries, we complement a small number of recent studies that have looked at these players emphaisizing their positive role for either imporoving the use of information and thus limiting winners' curse risks (McAfee [2011]) or more effectively administering their clients budget avoiding the potential 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.

Indeed, the third stream of the literature to which we contribute is the long and well established literature on collusion in auctions (pioneered by studies including Graham and Marshall [1987] and , and Hendricks, Porter and Tan [2008]itethendricks1989collusion). Before our empirical contribution, several theoretical studies had highlighted the potential vulnerability of online ad auctions to collusive bidding through intermediaries (Bachrach [2010], Mansour, Muthukrishnan and Nisan [2012] and Decarolis, Goldmanis and Penta [2017]). Given the ample scope for potential efficiencies driven by intermediaries, it is interesting that on balance the evidence points to a prevalence of the negative effects of intermediary's concentration on the search engine's revenues. Nevertheless, this is not surprising in light of practitioners' perceptions and the occurrence of at least three recent shifts in the market organization that can be interpreted in terms of a

[5] Goldfarb [2014] argues that targeting is the fundamental difference between online and offline marketing.

response to the increased strength of intermediaries. The first is the switch of Facebook (and of Google in a subset of markets) from the Generalized Second Price (GSP) auction to the VCG mechanism. As shown in Decarolis, Goldmanis and Penta [2017], in the presence of collusive bidding by an intermediary the VCG outperforms the GSP for both revenues produced for the search engine and allocative efficiency. Second, starting from May 2017 Google has begun increasing the reserve price applied to its search auctions. This is compatible with the standard recommendation from economics on how to curb collusion. Nevertheless, as we discuss in greater detail below, it is troubling in an environment where multiple slots are typically up for sale as the advertisers ending up worse off from an increase in the reserve price might not be those colluding. Third, Google, as well as similar companies, are actively trying to offer services in direct competition with those offered by the ad networks. "Disintermediation," as this practice is called, is not proving fully successful as advertisers perceive the benefits of using a more neutral platform, like a network, to bid on Google, instead of using a Google's own service. But this further highlights the ongoing arms race in this industry.

Finally, while we use the term collusion, the behaviors that we describe below do not qualify as such from a legal perspective. Antitrust laws typically require explicit coordination between the parties. What we describe is, instead, achievable simply through the behavior of automated pricing algorithms. This thus places our study also at the forefront of a very active debate in the antitrust literature on the potential risk that algorithmic pricing represents for collusive practices (see OECD [2017] and Calvano et al. [2018]).

The remainder of the paper is structured as follows: section II shows through a simple example the basic incentives at play, section III presents the industry, section IV describes the dataset, section V summarizes the empirical strategy describing both the clustering definition and the empirical strategy; section VI reports the main results and section VII concludes discussing policy implications and avenues for future research.

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 some keyword k. 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.

Now suppose that advertisers do not bid directly on the search auction. They submit their bid to an intermediary who runs internally a second price auction among 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. In panel B, we plot the same game, but with 2 intermediaries: both q and j are linked to agency α. This intermediary can now alter the search auction outcomes by retaining or emending the bids it places on behalf of its two clients. It might report just the highest bid among the two, $b_{q,\alpha} = 4$, or both bids, but setting $b_{j,\alpha} \in [0, 1]$, as indicated in panel B by the interval to which the bid $b_{j,\alpha}$ leads. In all cases, q wins the slot. It pays 1 instead of 3 thanks to the role of intermediary α.

Figure 1: An Example of Bidding through Intermediaries

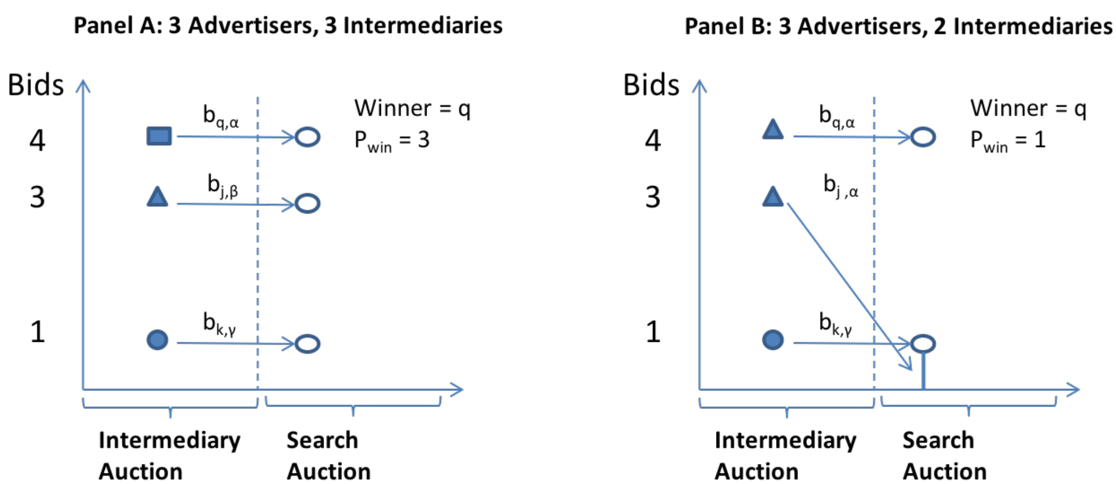

Notes: There are three advertisers (q, j and k) submitting arbitrary bids ($b_q = 4$, $b_j = 3$ and $b_k = 1$) to a second price auction held by the intermediary to which they are affiliated. In panel A, each advertiser has a different intermediary (α, β and γ). In panel B, q and j share intermediary 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 transmitted 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).

Now suppose that each advertisers is assigned to an intermediary by randomly drawing (with replacement) the identity of the intermediary out of a set K. Then, the search engine's revenues depend on random occurrence of the highest bidding advertiser being pooled together under the same intermediary with the next highest bidders. In Panel A with $K = \{\alpha, \beta, \gamma\}$, the search engine's expected revenues are $E(R) = 2.22$, but in Panel B with $K = \{\alpha, \gamma\}$ they fall to $E(R) = 1.75$. Indeed, for the general case in which N advertisers with their arbitrary bids are randomly assigned to K intermediaries, the search engine's expected revenues decline as the number of intermediaries falls, being minimized when there is a single intermediary.[6]

[6] The general expression is $E(R) = \sum_{n=2}^{N} b_n (\frac{K-1}{K^{n-1}})$. This assumes intermediaries shading bids other than the highest, as in panel B. Hence, if the highest bidder is with intermediary 1, then $R = b_2$ whenever advertiser 2 is not with intermediary 1, which happens with probability $(K-1)/K$; $R = b_3$ if advertiser 2 is with intermediary 1, but advertiser 3 is not, which happens with probability $(K-1)/K^2$; etc..

Mansour, Muthukrishnan and Nisan [2012] were the first to identify this theoretical problem in the context of ad exchanges where one slot is sold at a time. Decarolis, Goldmanis and Penta [2017] further extend the analysis to the situation of multiple, heterogenous slots that is typical in sponsored search auctions. They show that when multiple slots are sold via the Generalized Second Price (GSP) auction - as done by Google - both search engine's revenues and allocative efficiency are damaged by intermediaries' concentration. Their most surprising result is that both these effects are even more pronounced under the GSP relative to a benchmark system - the VCG auction - that is known to perform poorly when bidders play coordinated strategies. The reason being that the VCG is strategy proof: bidders outside the coalition have no incentive to shade their true value. In the GSP, instead, all bids are interlinked in equilibrium so that bidders outside the coalition will typically respond to others' collusive bids by lowering their own bids.

In the typical situation where advertisers bid over a multitude of keyword auctions, the scope for coordinating bids is even stronger. The markets can be split in simple ways (by keyword, geography, time of the day, consumers' demographics, etc.). The feasibility and payoffs of these strategies and the intermediary's incentive to implement them depend on technological, contractual and strategic considerations.[7] But at the same time intermediaries can bolster search engines' revenues through their ability to bid on more and better targeted keywords. Next we illustrate a few key facts of the industry helpful to understand this nuanced relationship between intermediaries and search engines that entails both conflict and partnership.

III Industry Background

The premise to what follows is acknowledging the impossibility of fully describing an heterogenous and rapidly evolving industry. The lack of public regulation also means that both search engines and networks are essentially 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, matured from interactions with market participants and direct involvement as both bidders and intermediaries' clients.

Internet advertising is split into sponsored search and display advertising. Our study focuses on the former.[8] In essence, an advertiser should open an account on the platform auctioning off ad space on the

[7]On top of that, an intermediary must both set up an internal system to select among its clients and contractually specify how the outcomes of this system will determine actions in the search auctions. Details can be crucial: if in panel B of Figure 1 part of q's surplus from having $b_{j,\alpha} \leq 1$ in the search auction is rebated back to advertiser j in proportion to j's bid in the intermediary's auction, then j might have an incentive to overstate his bid in the intermediary's auction (as in Asker [2010]'s cartel case). This would reduce the extent to which intermediation can reduce prices in the search auctions.

[8]Display advertising entails the sale of ad space on web pages, videos and apps. The owners of these spaces (*publishers*) connect to advertisers through a *display network*. Google's *AdSense* is an example of such network. In it, advertisers select the contextual environments in which they are interested and bid to show their ad there. Auctions are used to select the winning ad (typically only one) and determine the

search engine results page (for instance, *Google Ads*, formerly *AdWords*) and enter 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 eight) among the interested advertisers. The slot order reflects the ranking emerging from the bids (reweighed by ad quality in the case of Google and of some other platforms like Bing), and the payment occurs only if the user clicks on the ad.

In recent years, the market has evolved becoming more complex. "Ad exchanges" (ADX) have emerged as marketplaces connecting the demand of ad space with the supply by many, differentiated publishers (Google, Bing, Facebook, Amazon, Twitter, etc.).[9] Therefore, even though the search auctions remain firmly dominated by Google, advertisers often reach these auctions through ad exchanges. But unlike in the early days of the search auctions when advertisers directly bid, the ad exchanges – like in the typical financial exchanges – can be accessed only by qualified bidders and specialized intermediaries.

These marketplace changes have thus gone hand in hand with changes on the demand side. As mentioned in the introduction, advertisers typically contract out marketing campaigns to a digital marketing agency (DMA). Thousands of DMAs operate in the US market, but most of them belong to one of the seven agency networks, the DMA's holding companies. Within networks, bidding happens through centralized and specialized entities. There are a few kinds of such intermediaries, but the most relevant for our study are 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. Automated systems for bidding is what crucially characterizes networks' bidding: they allow faster bidding, to the point that a large portion of internet ad are currently traded in real time.[10] They also allow networks to more effectively use large amount of consumers' profiling data, both internal and purchased from third parties (*data exchanges*), possibly - again - in real time. The need for speed and data explain why the backbone of online ad bidding takes the form of the simplified two-auction structure of the example in the previous section [McAfee, 2011]. On the one hand, an (automated) auction within the network is run to select the advertisers, with their relative bids, to participate to the ad exchange auction. On the other hand, pooling all advertisers' data into a single network allows immediate access to relevant data and this is profitable both to save on the costs charged by the *data exchanges* and to improve speed.[11]

A feature of the intermediary sector relevant for our analysis is the networks' expansion through acquisition of DMAs. Table 1 reports the full list of the M&A operations that we observe in our data. For each case, it reports the name of the acquiring network, the number of advertisers linked to the agency at

payments. See Choi et al. [2018] for a recent, detailed review of the literature on display advertising.

[9] *Google Marketing Platform*, formerly *DoubleClick*, and Microsoft AD-ECN are examples of ad exchanges.

[10] Indeed, each network has its own ATD: for IPG, WPP, Publicis Groupe, Omnicom Group, Dentsu-Aegis, Havas and MDC, the corresponding ATDs are Cadreon, Xaxis, Vivaki (closed in 2014 and replaced by *Precision* in 2017), Accuen, Amnet, Affiperf and Varick Media.

[11] To grasp the value of higher speed, consider for instance that as page load time goes from 1 second to 3 seconds, the probability of bounce increases 32%, see `https://www.thinkwithgoogle.com/data-gallery`.

the time of the acquisition, the number of industries in which they operate and number of markets (a key measure that we describe in the next section). Across networks, there is heterogeneity both in the number and the size of the DMAs 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 700 different markets). Some acquisitions take the form of hostile takeovers, with subsequent attempts to buy back independence and, indeed, we observe two cases of divestitures. As explained in greater detail below, we will exploit changes in the market structure produced by these M&A operations to develop our IV strategy.

Table 1: 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	19
Essence Digital Limited	WPP	2015	1	1	74
Quirk	WPP	2015	5	2	272
SHIFT Communications	WPP	2017	13	8	700
Deeplocal Inc.	WPP	2017	5	1	74
Maruri GREY	WPP	2017	1	1	133
Zubi Advertising Services, Inc.	WPP	2017	3	2	185
Campfire	Publicis	2015	3	1	21
La Comunidad	Publicis	2015	9	5	181
Sapient Corporation	Publicis	2015	17	6	514
Blue 449	Publicis	2016	4	2	76
Forsman & Bodenfors	MDC	2017	5	1	155
Formula PR	Havas	2015	6	4	189
FoxP2	Dentsu-Aegis	2015	1	2	31
Rockett Interactive	Dentsu-Aegis	2015	1	1	12
Covario, Inc.	Dentsu-Aegis	2015	3	1	54
Achtung	Dentsu-Aegis	2016	2	1	100
Gravity Media	Dentsu-Aegis	2016	5	3	249
Grip Ltd.	Dentsu-Aegis	2016	3	2	70
Merkle	Dentsu-Aegis	2017	18	7	567
Gyro	Dentsu-Aegis	2017	12	6	270

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.

Before turning to the IV strategy, however, it is useful to describe what our data can tell about how intermediaries operate. First we want to contrast intermediaries with individual advertisers. Figure 2 reports the evolution of four variable as advertisers transition from bidding individually to bidding through a DMA. The four variables are: i) the number of keywords, ii) the cost-per-click, iii) a dummy for *long tail keywords*,[12] and iv) the number of markets, as defined in the next section.[13] To the left of the red, vertical line we report the yearly average (and standard errors) of the variables across all the advertisers that at time $t = 0$ we observe transitioning to a DMA. The next section describes more in detail the data, some simple indications are immediately apparent: on average, when an advertiser joins a DMA the number of keywords on which it bids increases, their average cost-per-click however declines. The nature of the keywords also changes with more of them becoming more specialized (long tail keywords), thus typically guaranteeing less competition (lower cost) and higher clicks. Finally and relatedly, the greater specialization is also reflected in the reduction in the markets to which the keywords belong. Although we have not yet explained how keywords are clustered into markets, it is remarkable that to a strong growth in the number of keywords it corresponds a marked decline in the number of markets spanned by these keywords.

Figure 2: Strategies of Advertisers and DMAs

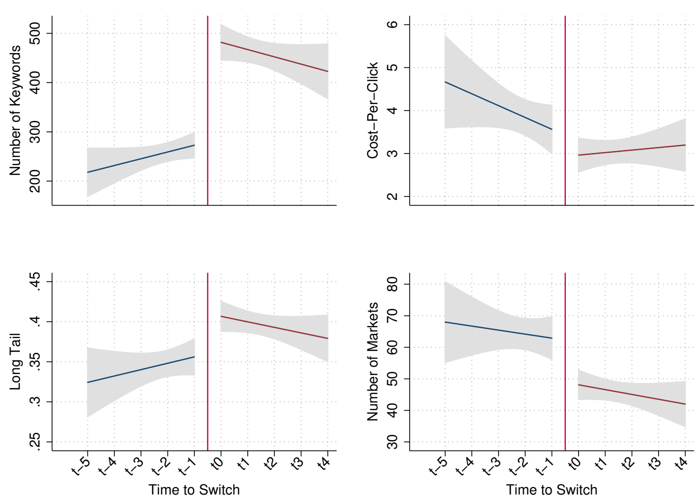

Notes: blue (maroon) lines are linear fits of average values before (after) joining a DMA at *t0* (red vertical line). The reported variables are *Number of Keywords*, *Cost-per-Click*, *Long-tail Keywords* and *Number of Markets* - panels a) to d). *Cost-per-Click* value is reported in USD, the shaded area corresponds to the standard deviation of the mean. Evidence for an additional set of variables is reported in the web appendix.

The above descriptive evidence points to some intrinsic difference in how intermediaries bid relative to advertisers. The greater specialization is compatible with more effective and better targeted bidding,

[12]These 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 funnel. For instance, while an advertiser might bid for "charity donations," an ATD might bid on thousands of more specific variants, one of which could be "charity donations furniture pickup."

[13]Additional results involving volume, position and other keyword characteristics like *branded* (see Golden and Horton [2018]) are reported in the web appendix.

possibly due to a better data usage. These are indeed some of the efficiencies that intermediation might bolster. Furthermore, the literature on sponsored search has identified the presence of strong externalities effects across ads that agencies could help internalize: 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]. Moreover, in the closely related context of the ad exchanges, the literature has identified further problems pertaining to limited information driving to winners' curse [McAfee, 2011] and budget constrains leading to inefficiencies [Balseiro and Candogan, 2017]. For all the above, the presence of an intermediary can be a source of efficiencies.

Figure 3: Keyword Splitted and Shared

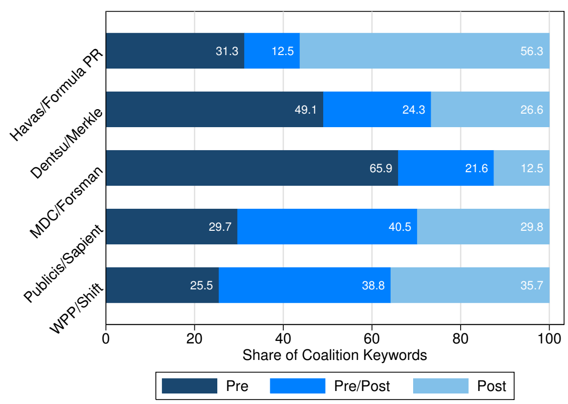

Notes: share of *coalition* keywords - i.e., keywords bid by both the advertisers in the DMA acquired and the ones in the acquiring network - before and after the merger. Shares are calculated on the keywords bid both before and after. *Pre* are keywords in coalition only the year before the merger; similarly, *Post* refers to keywords in coalition only after the merger, and *Pre/post* are keywords in coalition both before and after.

Nevertheless, when competing advertisers bid through the same network the collusive bidding issues described earlier can emerge: bids might be strategically retained or, at least, reduced. But also other collusive strategies like bid rotation or market split could be implemented. To assess the extent to which networks tend to bid on the same keywords on behalf of multiple advertisers instead of splitting keywords among advertisers, we compare what happens before and after a DMA joins a network. Let us define *coalition* keywords those keywords bid by different advertisers under the same network. Figure 3 reports the changes in the share of coalition keywords before and after mergers episodes between networks and DMAs in a subset of our data, distinguishing the keywords in coalition only before the event (*Pre*, deep blue), those only after the event (*Post*, light blue) and those which have been in coalition both before and after the merger (*Pre/Post*, marine blue). Both splitting keywords among clients and haVING coalition keywords are relevant strategies, with different networks making using them differentially. For instance, MDC resorted to a keyword split following its acquisition of Forsman & Bodenfors (abandoning most coalition keywords), but WPP maintained and even increased its coalition keywords following its acquisition of SHIFT Communications.

IV Data

The minimal data requirements to test the effects of market 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 analysis is based on a new dataset that contains all this information, and more. Figure 4 shows the hierarchical structure of the data we use in the analysis: the highest level (the networks, for non-independent agencies) groups the individual DMAs (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: for example, DMA 2 participates to the auction for *key5* on behalf of both Advertiser 3 and Advertiser 4, but we also consider *key2* in a coalition, given that Advertiser 1 and Advertiser 4 both bid through Network 1.

Figure 4: Redbooks-SEMrush Data Structure

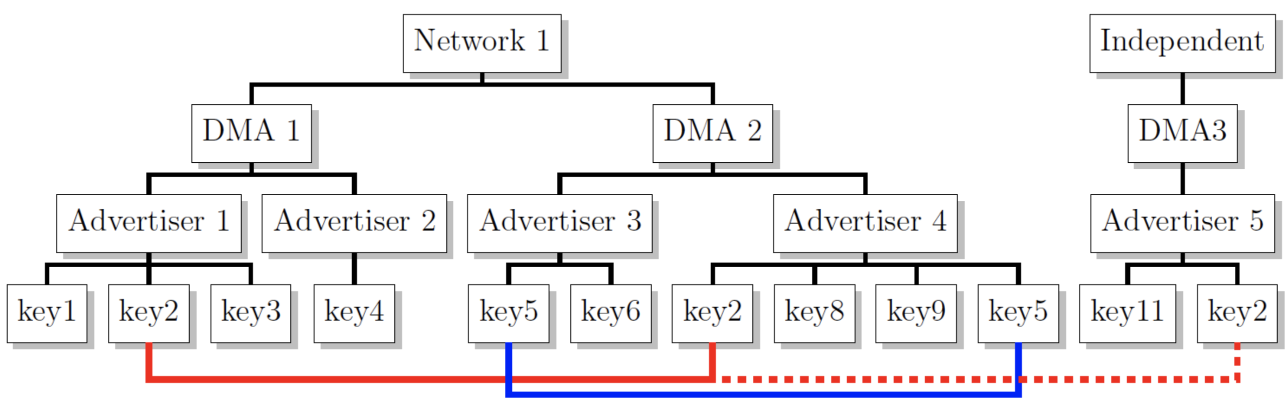

Notes: Hierarchical structure of the data. From bottom to top: keywords (SEMrush), advertisers (Redbooks/SEMrush), agencies and networks (Redbooks). Solid lines represent examples of coalitions: within DMA (blue) and network (red).

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 Day [2014] for other studies using these data). For each of these advertisers, the Redbooks data gives us the full list of DMAs. The data is yearly for the period 2012-2017 and covers around 6,000 advertisers (i.e., web domains) per year active in all sectors of the economy. For the years 2014 and 2017 only, we also have access to a linkage variable that relates each individual DMA to its network, if any. Overall, there are seven networks and about a thousand independent agencies; as for the advertisers with no agency affiliation, we consider them as bidding autonomously.

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 each keyword searched on Google, it collects the identity and the linked website of advertisers appearing in the sponsored ad slots. Moreover, it gathers information on the keyword-specific average CPC (i.e., the price advertisers pay for a user's click on an ad triggered by the given keyword), the position of the ad in the search outcome page, the volume of searches associated with the keyword (i.e., the average number of search queries for the given keyword in the last 12 months); the visible URL of the ad; the content of the ad; the traffic (that is, the share of the advertiser's traffic associated with the specific keyword) and also generate a *competition* index, ranging from 0 to 1, which reflects how tough is the competition within the auction to appear in the top slot.[14] Thanks to the visible url and the advertiser name, we are able to link Redbooks and SEMrush data.

Table 2 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 advertisers; in the right columns, those for keywords with at least one independent advertiser. The two groups are thus not mutually exclusive. For both of them, we see a similar CPC; although the mean and median CPC is lower for the network bidders. In terms of volume, for both groups the substantially lower value of the mean relative to the median indicates a tendency to bid on keywords that are infrequently searched. The lower value of *traffic* (1%) observed for the network advertisers relative to the 6% of the non network advertisers is compatible with the former placing ad over more keywords. *Competition* is rather similar for the two groups despite the number of bidders being slightly higher for keywords on which the networks advertisers bid. *Coalition* measures the number of keywords where more than one of the ad 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.[15]

Panel B shows the relative size of each one 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) indicates an high degree of concentration: it reaches 72% 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 networks' presence across keywords (the values within each column do not add up to one since the same keyword can display ad from multiple networks, as well as non-network bidders). The sheer prominence of networks in the data, together with bidding centralization at the networks' ATD level induce us to consider networks as the key players in our analysis. The next step is then to define the right level of keyword aggregation at which to study networks' concentration. We first do so and then return to the bottom panel of Table 2.

[14] Since the Redbooks data are a snapshot of the agency affiliation during the first half of January each year, we also downloaded data relative to January 15^{th} from SEMrush.

[15] Although the vast majority of cases involves coalitions of size 2, there are a few examples of larger coalitions. For instance, for the keyword "online banking" there are four advertisers (Bank of America, Travelers, Geico and State Farm) all affiliated with a single DMA (The Martin Agency, a major marketing agency that is also part of the Interpublic Group of Companies). Furthermore, not reported in the table is another interesting statistic that shows that there is typically a single multi-advertiser DMA per keyword.

Table 2: 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
Competition	0.58	0.69	0.39	15,383,769	0.59	0.73	0.39	21,525,056
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

Panel C. Statistics by Market

	Not-merged Markets					Merged Markets			
	Mean	SD	Median	Obs		Mean	SD	Median	Obs
Revenues	298	2,198	13	64,039	ΔR	-0.09	2.05	-0.04	39,179
HHI	3,568	2,731	2,607	64,039	ΔV	0.02	1.03	0.05	39,179
Long-tail	0.41	0.38	0.31	64,039	ΔK	-0.31	1.52	-0.29	39,179

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 DMA bid (columns 5 to 8). The variables included are: *Cost-per-click* is reported in USD, *Competition* is provided by SEMrush and is generated through an undisclosed algorithm, *Coalition* is an indicator function for the presence of multiple bidders affiliated with the same network participating to the keyword auction, while *Coalition size* is populated for keywords with coalitions only. Organic results are 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 the non-network DMAs; 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 do not add up to one since the same keyword can display ad from multiple network and non-network bidders). Panel C: statistics of the sample used in the IV analysis. The unit of observations are markets-years (i.e., clusters of keywords, see section V.1). There are about 23,000 markets per year. On the left, we report statistics for the search engine's revenues (in $1,000), the market's HHI and the share of long-tail keyword. On the right, we report the first differences between year t and $t-1$ for the revenues, the volumes and the number of keywords.

V Empirical Strategy

V.1 Market Definition Via Thematic Clustering

Potential definitions of markets range from granular - the single keyword - to aggregate - the 23 industries provided by Redbooks. The latter help identifying the agency/network sector of specialization, but contain too heterogeneous keywords to serve the purpose of the analysis of competitive and strategic effects (as discussed more generally in Werden and Froeb [2018]). In order to find a useful middle-ground solution, we apply-state-of-the-art natural language processing methods to form keyword clusters 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 their semantic similarity (thematic 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 on how consumers find products and services online. With such corpus, the approach described below can mimic what is sometimes done in antitrust cases by surveying consumers about what products they see as living in the same product space. Without aiming for the same accuracy required for competition cases, we nevertheless see this approach as a valuable contribution.[16] We first detail how it works and then discuss some of its limitations.

Step 1 - 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 synonims and antonyms in text clustering or semantic similarity exercies. We use an unsupervised learning algorithm (GloVe, developed by Pennington, Socher and Manning [2014*a*]) to obtain vector representations for each word 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. In particular, the novelty of

[16] The definition of the relevant market is one of the main challenges in both the applied Industrial Organization literature and competition policy. It is the intersection of a *relevant product market* and a *relevant geographic market*. The European Commission states that the former “comprises all those products and/or services which are regarded as interchangeable or substitutable by the consumer by reason of the products’ characteristics, their prices and their intended use”, whereas the latter “comprises the area in which the firms concerned are involved in the supply of products or services and in which the conditions of competition are sufficiently homogeneous”.

the GloVe model with respect to previous approaches is that it combines the benefits of matrix factorization approaches - i.e., reduce the dimensionality of co-occurrence matrices - with the good performances of skip-gram models (like Google's *word2vec*) in word analogy tasks.[17]

We use a GloVe dataset pre-trained on 840B documents, corresponding to $\approx$ 2.2M unique terms, from Common Crawl in English, featuring 300 dimensions. Using as corpus such an extensive body of text which originates from mimicking the web crawling behavior of typical internet users is what makes the resulting vectorization analogous to surveying people about proximity between keywords.[18] Similarly, when applied to the sponsored search keywords in our data, the vectorization shall reflect the proximity between products and services identified by the semantic similarity between keywords. Once split every keyword in its constituent *terms*, we proceed in 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 underlying terms (see further details in the appendix).

Step 2 - Thematic Clustering We perform the thematic clustering step within each one of the 23 industries in which the advertisers are partitioned 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 compute the cosine distance matrix between all keywords in the industry; finally, we run a spherical k-means clustering algorithm (Dhillon and Modha [2001]) on the cosine distance matrix with 1,000 centroids, industry by industry, to group keywords into thematic clusters - that is, we identify the semantic "themes" linking the keywords. We treat the resulting clusters as markets.

The above approach suffers from two potential drawbacks that are intrinsic to the method. The first is that both substitute and complement products/services are likely to be pooled together. Nevertheless, to the extent that the advertisers of complement products are in competition among them for the limited ad space, our analysis would not be distorted. But the possibility of joint marketing efforts by advertisers of complementary products would clearly be a concern.[19] Similarly, our keyword-based method is able to identify different geographical markets only to the extent that the geographical aspect is particularly salient in 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).[20]

[17]In the Appendix we provide additional details on GloVe and its main features. Interested readers could refer to sthe original paper by Pennington, Socher and Manning [2014*a*] for the full development of the model. For applications of GloVe in diverse scientific fields, see e.g. Caliskan, Bryson and Narayanan [2017].

[18]The dataset, and GloVe code, are open source and freely available at https://nlp.stanford.edu/projects/glove/. There is a number of other datasets available (e.g., trained on Wikipedia, on Twitter, with 25, 50, 100 or 200 dimensions, etc.): results obtained re-running all analyses using different datasets are available from the authors upon request.

[19]Cao and Ke [2018] is a recent theoretical contribution on how in sponsored search a manufacturer should optimally coordinate with its retailers by sharing a fixed percentage of each retailer's advertising cost.

[20]We experimented with our algorithms with different number of clusters and changing the way in which

In general, testing the quality of the clusters obtained requires a reference sample where keywords and markets are correctly associated. Lacking this type of sample, we resorted to random inspection of the clusters' quality. We find overall very satisfactory results with our initial motivating concern of related but 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 run it through *Amazon Mechanical Turk* (see appendix E 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 ranges between 80 and 90%.

V.2 Instrumental Variable Estimation

Having defined markets, we can now quantify intermediaries' concentration within these markets and analyze its implications. We begin by describing the main variables used and then present the empirical strategy.

A. Outcome Variables - 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:

$$R_{mt} = \sum_{k \in K_m} CPC_{kmt} * Volume_{kmt} * CTR_{kmt} \tag{1}$$

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 yearly searches and CTR_{kmt} is the cumulative Click-through-Rate of all the sponsored ad slots shown for keyword k.[21]

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 should successfully standardize markets. Rather than attempting this route, we focus on the growth rate of revenues for the same keywords in a specific market over time ($\Delta ln(R_{mt})$) and analyze how this relates to the intermediaries' concentration in the market. By focusing on growth, we also avoid concerns related to time-invariant unobservable differences in the characteristics of the keywords that may be correlated with revenue levels. We will also control for the influence of time-varying features at the level of keywords by using information on the organic results (i.e., the non-sponsored links) of the query.

we built the keyword distance metrics. See appendix.

[21]For each k, the overall CTR_{kmt} is the cumulative sum of all the ad slots j appearing on the search outcomes page of keyword k: $CTR_{kmt} = \sum_{j \in k} CTR_{jkmt}$. We use the industry-averaged click-through rate for each slot/month provided by *Advanced Web Ranking*, the leading provider of this type of data.

Other outcomes that will be interesting to explore in addition to the revenue measure R are the average CPC and volume, as well as the total number of keywords, separating branded and long tail. These additional outcomes allow us to explore the channels through which concentration affects the search engine revenue and they can offer direct evidence of the different types of strategies previously discussed.

B. Concentration Measure We measure intermediaries' concentration computing the Herfindahl-Hirschman Index (HHI) for each market/year. Suppose that we have a market m defined by the set of keywords K_m. For each keyword $k \in K_m$, there are sponsored ad slots j, 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' concentration 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:

$$s^i_{mt} = \frac{1}{S_{mt}} \sum_{a \in A_i} \sum_{k \in K_m} \sum_{j \in J_k} CTR_{jkmt} * Volume_{kt} * 1\{a \text{ occupies } j \in J_k\}. \tag{2}$$

Thus, our concentration measure for market m at time t is the squared sum of each intermediary's market share, or: $HHI_{mt} = \sum_{i=1}^{I} ({s^i}_{mt})^2$.[22] As discussed earlier, the intermediary is the network, or, if not present, the DMA; alternatively, it is the individual advertiser itself. Figure 5 shows the dynamics of HHI in our sample; more specifically, for each market we take the difference between HHI_{2017} and HHI_{2014}. The plot shows, from top to bottom, the number of markets with changes in HHI of 1,500+ base points, from 1,000 to 1,500, down to negative values (i.e., experienced decreases in concentration). This figure makes evident how a majority of markets experienced a tremendous concentration increase: about 14,000 markets experienced an HHI increase equal or greater than 1,500 points - even though the plot also highlights a substantial heterogeneity in HHI dynamics with some markets experiencing HHI declines. As a reference, consider that the US Horizontal Mergers Guidelines indicate as potentially problematic those mergers that, within already moderately or highly concentrated markets (i.e., those with an HHI of at least 1,500), lead to an HHI increase of at least 100 points.

Having defined both the unit of observation and the the main variables, we can now return to the bottom panel of Table 2. In panel C, we present basic summary statistics for the main variables entering our IV analysis. There we see, for instance, that the average market is highly concentrated with an HHI of 3,568

[22]Despite several theoretical and practical shortcomings of this measure, it is commonly used in both academia and competition policy to proxy for concentration. See, among others Hastings and Gilbert [2005], Dafny, Duggan and Ramanarayanan [2012] and the US Horizontal Merger Guidelines.

points. We also see that across markets long tail keywords represent 41% of the cluster's keywords. For those markets observed in consecutive years, the variables on the right hand side of panel C reveal that revenues tend to decline over time, as does the number of keywords, while instead volumes grow. As stated above, to perform an analysis across market, but without attempting to homogenize the vastly heterogenous markets observed, the sample in first differences is the one we use in the IV analysis below.

Figure 5: Change in HHI - 2014 to 2017

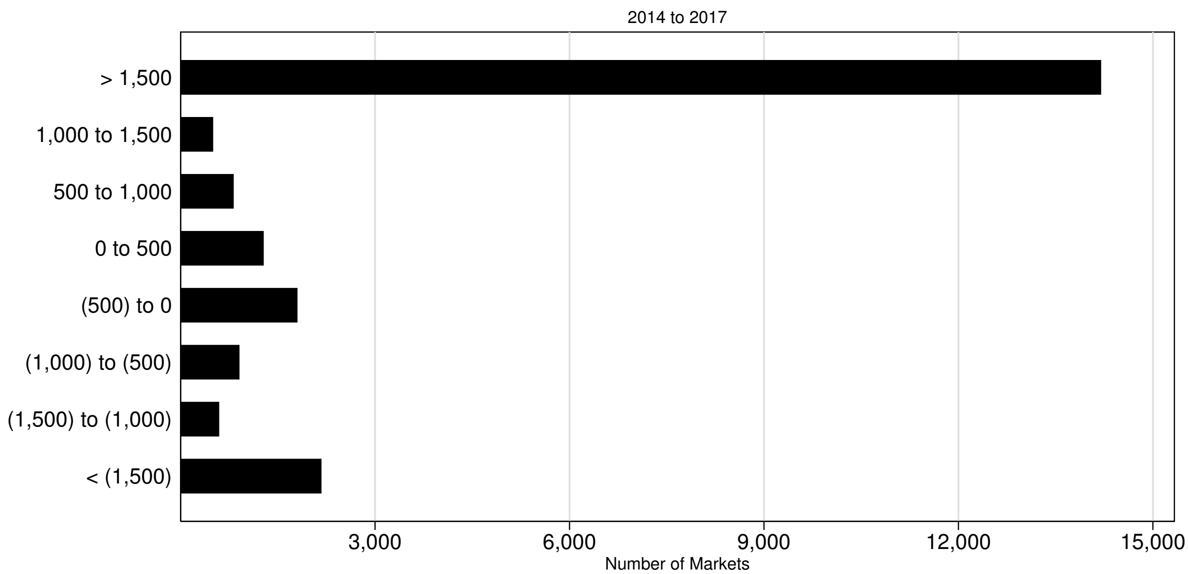

Notes: HHI is scaled 0 to 10,000. We report the number of markets on the x-axis, grouped according to the level of ΔHHI.

C. Empirical strategy In an ideal setting, the market-level OLS regression of revenue growth on the level of concentration would reveal the causal effect of concentration. In practice, it is not possible to assign a causal interpretation to this conditional correlation. For instance, a keyword k might have become suddenly fashionable for some exogenous reasons; 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 observing a positive association between intermediaries' concentration and the growth of search engine's revenues, but it does not imply the existence of a causal relationship between the two phenomena. We, nevertheless, report below this OLS regression which will represent the starting point in the analysis:

$$\Delta log(R)_{mt} = \beta HHI_{mt} + \phi X_{mt} + \tau_t + \gamma_z + \zeta_{at} + \epsilon_{mt}. \tag{3}$$

The dependent variable is the logartihmic change in revenues in market m (of industry z) between t and $t-1$. This regression controls for year (τ_t), industy (γ_z) and merger (ζ_{at}) fixed effects,[23] as well as for characteristics of the market-time included in X_{mt} (we use the number of organic links, plus a series of keyword- and market-related controls). To deal with the endogeneity, we use an IV strategy inspired by that of Dafny, Duggan and Ramanarayanan [2012] on the health insurance markets. It exploits the changes in the market structure originating from mergers and acquisitions (M&A) between intermediaries as a source of exogenous shock to the 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 the merger. Given that two merging intermediaries might have clients in a plethora of markets with possibly quite 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" ($sim\Delta HHI_{mt}$) as the difference between the actual HHI and the counterfactual HHI, absent the merger. 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 t. The contribution of the new entity to the concentration measure amounts to the squared sum of the shares of the merged firms, $(s^{\alpha}_{mt}+s^{\beta}_{mt})^2$, which is by construction greater or equal than the contribution of the counterfactual with unmerged firms, $(s^{\alpha}_{mt})^2+(s^{\beta}_{mt})^2$. Therefore, the merger-induced change in market m's HHI reads:

$$sim\Delta HHI_{mt} = \underbrace{(s^{\alpha}_{m,t}+s^{\beta}_{m,t})^2}_{\text{Share of merged firm } \alpha+\beta} - \underbrace{((s^{\alpha}_{m,t})^2+(s^{\beta}_{m,t})^2)}_{\text{Sum of single firm } \alpha \text{ and } \beta \text{ shares}} = 2s^{\alpha}_{m,t}s^{\beta}_{m,t} \tag{4}$$

We use, for each market-year, the variable $sim\Delta HHI_{mt}$ as instrument for HHI_{mt}. In figure 6 we report the distribution of the proposed instrument in our sample of markets interested by the mergers. We observe a large number of markets experiencing substantial increases in concentration, but also about 200 markets with a negative $sim\Delta HHI$, these are the markets interested by DMA returning independent after a buyback. Using $sim\Delta HHI_{mt}$ as instrument for HHI_{mt} therefore entails the following first-stage regression:

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

where the set of covariates entering the regression along with $sim\Delta HHI_{mt}$ is the same of those in equation (3). 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]).

[23] ζ_{at} captures the effect of M&As on revenues via an indicator function for the presence of the intermediary a after the date of its acquisition (t^*_a). More specifically, $\zeta_{at} = 1(a \in m, t > t^*_a)$.

The next step of the empirical strategy entails estimating the following reduced-form regression:

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

The proposed instrument captures large variations in the degree of concentration and hence strongly influences $\Delta log(R)_{mt}$. Furthermore, while not directly testable, the exclusion restriction requiring that all the effects of $sim\Delta HHI_{mt}$ on $\Delta log(R)_{mt}$ pass through HHI_{mt}, is plausibly satisfied given the nature of the chosen instrument which is itself constructed as a function of the lagged HHI.

Figure 6: Distribution of Simulated Changes in HHI

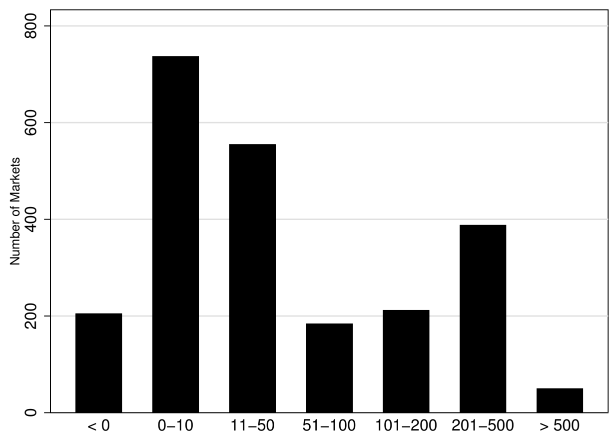

Notes: the sample includes markets with non-zero $sim\Delta HHI$ only. Negative values refer to the divestitures.

VI Results

A. Baseline We begin the presentation of our results from the first-stage and reduced-form estimates. In Table 3, odd columns report the reduced-form estimates, while first-stage are on even columns: models across columns feature an increasing set of controls (*Organic Results*, which capture the "popularity" of the keywords in the market, hence reflecting how appealing it is to customers, from model (2) on), fixed effects (industry, merger, and time from model (3)) and keyword-level characteristics (model (4)). We estimate a positive first-stage, and a strongly negative reduced-form regression: in all cases the effects are significant and more so for the baseline model - (3). Model (4) adds potentially endogenous regressors which are typically considered in the optimization of advertising campaigns by DMAs: the average number of *long-tail* and *branded* keywords.[24] As discussed when presenting Figure 2, when an advertiser switches from

[24]In model (4), by controlling for both long tail and branded keywords, we test whether these keyword choices affect revenues through increases in concentration. More specifically, the model tests whether the

running directly tis campaigns to delegating them to a DMA, several aspects of the bidding strategy can change (including the utilization of long tail and branded keywords). Since we want to isolate the effects of concentration and avoid picking up, for instance, that individual advertisers commit mistakes and fail to bid strategically, the estimation sample used throughout this section includes only advertisers bidding via DMAs, either network or independent. That is, we exclude advertisers bidding autonomously. The robustness checks presented later show that, however, this sample choice makes little difference on the estimates.

Table 3: Reduced Form and First Stage Estimates

Panel a) GloVe - Sum

	(1)		(2)		(3)		(4)	
	RF	FS	RF	FS	RF	FS	RF	FS
$\text{sim}\Delta \hat{HHI}$	-4.517***	1.193***	-4.495***	1.183***	-4.409***	1.375***	-4.379***	1.368***
	(1.695)	(0.127)	(1.694)	(0.126)	(1.702)	(0.107)	(1.692)	(0.111)
Weak Id. F-Test	88.12	88.12	88.36	88.36	165.17	165.17	152.48	152.48
Underid. F-test	7.51	7.51	7.51	7.51	8.26	8.26	8.09	8.09
Observations	39,179		39,179		39,179		39,179	
Organic Results			✓		✓		✓	
Industry FE					✓		✓	
Merger FE					✓		✓	
Year FE					✓		✓	
Keyword Characteristics							✓	

Notes: Reduced-Form (odd columns) and First-Stage (even columns) estimates. *Keyword characteristics* are *branded* and *long-tail* keywords. Robust standard errors in parentheses.

Both the first stage and reduced form estimates in Table 3 are rather stable across model specifications. In particular, the comparison of the models (3) and (4) reveals that the estimates are little affected by controlling fro branded and long-tail keywords. The positive sign of the $sim\Delta HHI$ estimate in the first-stage regression is in line with the idea that advertisers in this market maintain the same intermediary after, through an intermediary's M&A operation, it belongs part of a network representing competing advertisers. This is also in line with what shown in Decarolis, Goldmanis and Penta [2018] through a discrete choice framework, that is that an advertisers is more likely to select an intermediary the more of its rivals use the same intermediary. The large magnitude indicates that mergers have strong effects on concentration at the local market level and indeed the large value of the F-statistics confirms the relevance of the proposed instrument.As regards, the reduced form estimates, the estimates indicate a negative and statistically significant relationship between the change in revenues and the simulated change in HHI. Considered together, these regressions already inform us that the IV estimates indicate a negative impact of intermediaries' concentration on the search engine's revenues' growth.

This is indeed what we observe in Table 4. This table reports OLS (columns 1 to 4) and IV (columns 5

effects of concentration bite only through the implementation of different strategies or better keyword optimizations (i.e., if the concentration leads to market split or changes in keyword composition).

to 8) estimates. Both sets of coefficients are negative and statistically significant. IV coefficients are larger, being about four times the corresponding OLS ones. The choice of using $\Delta log(R_{mt})$ as a dependent variable has the benefits described above, but makes the economic interpretation of the estimates not immediate. The baseline estimate (column 7) implies that one standard deviation change in HHI causes a decrease in Δ revenues of $\approx 8\%$ or, alternatively, of 0.42 standard deviations. Even though the measure of HHI shows quite a high degree of variability - the standard deviation in the estimation sample amounts to 2,600 base points - the above is an economically relevant result, and in connection with the major increase in HHI plotted in figure 5, it depicts a sector in which the networks are significantly eroding the search engines' revenues thanks to the use of centralized technology.

Similarly to the earlier table, estimates are rather robust across specifications. With respect to the most parsimonious model specification (columns 1 and 5), controlling for the average number of organic results (columns 2 and 6) has negligible effects on the estimated parameters, but including fixed effects (columns 3 and 7) both modifies the magnitude of $\hat{\beta}_{IV}$, decreasing its absolute value, and bolsters the significance of the organic results parameter. Adding fixed effects helps capturing the heterogeneity across years and among industries, which we analyse in details below (additional specifications exploiting fixed effects are discussed below). Finally, adding keyword characteristics - columns 4 and 8 - does not produce sizeable effects on either the magnitude or the significance of the HHI parameter. Moreover, we do not estimate significant effects of these strategic channels on changes in revenues.

Table 4: OLS and IV Baseline Estimates

Panel a) GloVe - Sum

	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{HHI}$	-0.945***	-0.937***	-0.765***	-0.754***	-3.785***	-3.801***	-3.206***	-3.202***
	(0.0435)	(0.0438)	(0.0463)	(0.0464)	(1.331)	(1.342)	(1.223)	(1.229)
Organic Results (billion)		0.200***	-0.0936	-0.119*		-0.319	-0.350**	-0.357**
		(0.0717)	(0.0627)	(0.0638)		(0.259)	(0.152)	(0.146)
Keywords Characteristics								
Branded Keyword				0.0122				-0.0304
				(0.0327)				(0.0433)
Long-tail Keywords				-0.116***				-0.0126
				(0.0269)				(0.0596)
Observations	39,179	39,179	39,179	39,179	39,179	39,179	39,179	39,179
Industry FE			✓	✓			✓	✓
Merger FE			✓	✓			✓	✓
Year FE			✓	✓			✓	✓

Notes: Columns 1 to 4 report OLS estimation of equation (3), columns (5 to 8) the IV version, with increasing number of controls and Fixed Effects. Robust standard errors in parentheses.

B. Robustness Checks We assess the reliability of the above estimates to several modifications. We report in this section five sets of robustness checks presented in Table 5, while additional ones are presented in the appendix. First, we explore to what extent using as markets groups of thematically connected keywords modifies the estimates compared to the advertisers' industries as markets. In columns 1 and 2 of Table 5, we report the industry-level IV estimates of the baseline model (model 3 of Table 3). The point estimates are similar to the baseline ones, but are more noisy and statistically indistinguishable from zero. This is also in line with the high degree of heterogeneity that we find among industries (see next paragraph). In columns 3 and 4, we further explore the relevance of our clustering procedure. In particular, to obtain the numerical representations of the keywords, the procedure described above employed term-by-term sums of GloVe vectors. We thus report in columns 3 and 4 a robustness check involving markets constructed by averaging GloVe vectors. 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 - trivially, keywords with more terms face higher likelyhood 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 the substantial robustness of the estimates, this is indeed reflected in noisier parameters, with a smaller $\hat{\beta}_{IV}$.

Table 5: Robustness Checks

	Industry Level		Clustering		Complete		No *Publicis*		No Outliers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\hat{HHI}$	-50.05	-3.613	-2.905**	-2.415**	-3.574**	-3.150**	-4.226***	-2.679**	-4.093***	-3.295***
	(48.88)	(7.341)	(1.170)	(0.974)	(1.564)	(1.399)	(1.296)	(1.224)	(1.274)	(1.232)
Organic Results (billion)		-10.63		-0.149		-0.290**		-0.306**		-0.364*
		(26.10)		(0.0956)		(0.145)		(0.156)		(0.206)
Observations	69	69	39,069	39,069	43,190	43,190	39,179	39,179	35,050	35,050
Industry FE		✓		✓		✓		✓		✓
Merger Dummies		✓		✓		✓		✓		✓
Year FE		✓		✓		✓		✓		✓

Notes: in columns 1-2 the estimates of model (3) on a sample without clustering - i.e., at the industry level. In columns 3-4 the market definition is obtained by averaging GloVe vectors at the keyword level before running the k-means on cosine distances; columns 5-6 report the estimates with the full sample (i.e., without excluding independent advertisers). In 7-8 we treated all *Publicis* agencies as independent DMAs, and in 9-10 we excluded *Media* and *Pharmaceutical & Health* from the analysis.

The third robustness check in Table 5, *tab:rob*, reports the IV estimates for the full sample, inclusive of those advertisers bidding autonomously. This group contains a relatively small, but also very heterogenous set of advertisers with both some large, sophisticated advertisers and some very small, likely unsophisticated ones. Columns 5 and 6, however, show that point estimates are, to all extents, equivalent to the baseline case. The fourth robustness check exploits an event involving one of the networks, Publicis Groupe. In 2015, the then CEO announced that Publicis was attempting to differentiate from the other six networks by closing its centralized ATD (called Vivaki AOD and founded about 7 years earlier). The bidding experts of this ATD were then relocated to multiple DMAs within the network. A detailed analysis of this strategy is beyond the scope of this paper, but we record that less than two years later, in 2017, Publicis had a new CEO who immediately announced the launch of a nATD, *Precision*.[25] We exploit this episode to assess the robustness of our estimates as follows: columns 7 and 8 replicate the baseline model, but in computing the concentration measures we ignore both the affiliations and the mergers involving Publicis. The results hold and - consistently with the hypothesis that Publicis mergers until 2018 were less effective in implementing bid coordination - the non-reported first-stage estimates are both bigger in absolute terms and more significant, with significant F-tests.

The fifth and last robustness check involves excluding two industries, *Media* and *Pharmaceutical & Beauty*, which the analysis in the next section indicates as different from all the others. Depending on the market structure, on the specific areas of expertise or on the nature of the customers, advertisers (and intermediaries) in different industries might follow diverse marketing strategies. Below we discuss the heterogenous effects of concentration across industries. Before that, however, we point out through the estimates in columns 9 and 10, that even excluding the two most extreme industries, the overall effects of concentration are fairly in line with the baseline estimates.

C. Refinements and Heterogeneity Analysis 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 briefly analyze both the channels through competition impacts revenues and the heterogeneity across industries and specific mergers.

Table 6: Analytical Refinements: IV Estimates on Different Outcomes

	$\Delta log(\hat{R})$		$\Delta log(\#keywords)$		$\Delta log(volume)$		$\Delta log(cpc)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{HHI}$	-3.785***	-3.206***	0.0754	0.0696	1.250	1.103	-0.786	-0.411
	(1.331)	(1.223)	(0.868)	(0.756)	(1.061)	(1.049)	(0.565)	(0.460)
Organic Results (billion)		-0.350**		-0.187*		0.293**		0.0318
		(0.152)		(0.0975)		(0.126)		(0.0561)
Observations	39,179	39,179	39,179	39,179	39,179	39,179	39,179	39,179
Industry FE		✓		✓		✓		✓
Merger Dummies		✓		✓		✓		✓
Year FE		✓		✓		✓		✓

Notes: IV estimates of (3) on different outcomes: $\Delta log(R_{mt})$ (1-2), $\Delta log(number\ of\ keywords_{mt})$ (3-4), $\Delta log(Volume_{mt})$ (5-6) and $\Delta log(cpc_{mt})$ (7-8). Odd columns refer to the estimation without controls and fixed effects, even columns report the full model. Robust standard errors in parentheses.

[25] Even though formally Precision offers support to programmatic teams in network agencies and also directly buys media for some clients, its core business seems to be to centralize the programmatic strategy of Publicis agencies, exactly what Vivaki was doing until 2015. See https://adexchanger.com/agencies/two-years-vivaki-aod-publicis-rolls-new-programmatic-hub/

As regards the channels, we look at CPC, search volume and the number of keywords. In table 6 we explore the association between market concentration and changes in the number of keywords (columns 3 and 4), in search volume (5-6) and in average CPC (7-8). Although these are the channels compose the revenues, the estimates are noisy and not statistically significant. Nonetheless, point estimates provide suggestive indications on the way in which market concentration affects market revenues: the number of keywords show a null effect - in terms of magnitude - the volume increases (compatibly with specialized networks identifying the most popular keywords) and the cost-per-click, in accordance with the theoretical predictions, decreases substantially. In part, the lower significance is mechanically due to the higher noise introduced by the fact that here we weight equally all keywords so that, for instance, the CPC estimates weight equally frequently and infrequently searched keywords. In part, however, the lower significance also reflects the fact that individual networks follow different strategies (for instance coalition keywords vs keyword split) so that capturing their aggregate effects on the individual channels is harder. To explore this latter aspect, appendix G reports the estimated merger fixed effects - the ζ_{at} in (3). We find that, while there is substantial heterogeneity for both the number of keywords and their volume, the effect on *cpc* is typically negative.

Figure 7: Industry-level IV estimates distribution

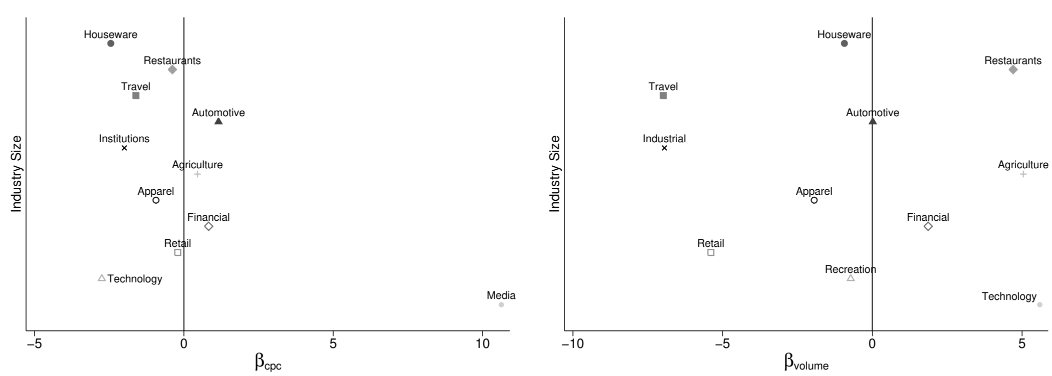

Notes: Industry-by-industry IV estimates. $\hat{\beta}_{IV}$ of the regressions for $\Delta log(cpc)$ (left panel) and $\Delta log(volume)$ (right panel).

Finally, in Figure 7 we explore differences among industries. It reports the distribution of $\hat{\beta}_{IV}$ estimated at the industry level, for $\Delta log(cpc)$ (left panel) and $\Delta log(volume)$ (right panel). Although negative on average, the former features positive values for some relatively important industries (*Automotive*, *Financial*, *Agriculture*) with two strongly positive outliers (*Media* and *Pharmaceutical & Health*, the latter excluded for readability). The estimated effect of concentration on changes in volume shows a higher degree of noise, with half industries characterized by a positive impact, and half by a negative one - with two negative outliers, *Media* and *Pharmaceutical & Health*, both excluded for readability. The resulting picture confirms that networks, and DMAs, follow different strategies depending on the market structure and competitive pressures within industries. The overall effect on revenues is hence emerging from multiple, different paths.

VII Conclusions

The findings in this study indicate that concentration among the intermediaries bidding on behalf of advertisers in sponsored search negatively and significantly impacts the growth of the search engine's revenues. Despite the potential benefits for the search engine from the increased efficiency that intermediaries bring to the market along many dimensions, especially through enhanced speed and better data, the negative revenue result is indicative of the intermediaries capability to reduce the average prices. This is a novel insight on what is currently one of the largest and fastest growing advertising markets.

In a period of increasing attention from the academic and policy worlds about industry concentration, our study reveals that technological innovation and countervailing power pose a limit to the market power attainable by dominant firms, like Google is in the sponsored search. Enhancing buyer power has been proposed as a preferable solution to increasing regulation of online platform markets (EU Commission [2018], Mullan and Timan [2018]). To the limited set of cases upon which these recommendations were based,[26] our study adds evidence from a major industry, within which we propose a rigorous way to define markets.

Furthermore, while multiple recent papers are exploiting the potential of field experiments to learn about online ad effectiveness,[27] we are the first to provide systematic evidence on intermediaries. We showed that campaigns run through intermediaries differ from those run by individual advertisers and such differences will be even deeper when the same intermediary bids on behalf of multiple, competing advertisers. We thus offer evidence linking the concerns of the theoretical literature on collusion to what is currently happening in one of the main online auction platforms, underscoring the potential benefits of a market design approach.

Numerous questions are left open for future research. In particular, it is unclear what internal or external frictions could slow down, or even revert the intermediaries' concentration process 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. External frictions can, instead, derive from actions of either antitrust authorities or the platform. The latter could implement responses like the ones mentioned in the introduction: changing the auction format, promoting disintermediation services and increasing the reserve price. If these three ongoing trends are happening in response to the lower revenue growth caused by intermediaries' concentration, then our analysis suggests that the latter is likely the most problematic of the three. Increasing the average CPC in a market dominated by concentrated intermediaries will require substantial reserve price increases and this will likely hurt, at least in part, the "wrong" advertisers (i.e., those not sharing a common intermediary).

[26] Mullan and Timan [2018] list three cases: the British Retail Consortium, a trade association representing retailers, successfully rejecting an ad-valorem pricing proposed by Mastercard, the OntheMarket platform created by a collective of estate agents to balance the strength of the two UK real estate platforms and TheAuctionRoom a bidding platform created by former clients of the online auction platform ATG Media.

[27] See Golden and Horton [2018] and Johnson, Lewis and Reiley [2017] for a review of the recent literature.

Small advertisers placing low bids near the reserve price and not bidding through intermediaries might find themselves either paying a substantially higher cost-per-click or being outright excluded from the set of ad displayed. This in turn can have detrimental long run effects for the efficiency of online ad platforms.[28]

The second source of external frictions involves actions by antitrust authorities to block or limit joint bidding by intermediaries. In fact, if on the one hand buyer power can constraint the dominance of some online platforms, on the other hand most antitrust statues forbid all forms of price coordination. Buyers' groups are typically treated differently from horizontal cartels, but past cases of sanctioned buyers' cartels exist. Crucial to guide the choices of an antitrust authority would be the clear identification of the final consequences of intermediaries' concentration for consumers. Future research should thus quantify to what extent a transfer of revenues away from Google and toward the more competitive industries where advertisers operate induces a pass-through of some of the benefits to consumers. But it is also unclear to what extent a decrease in Google's revenues might worsen the quality of services that consumers attain on its search engine, or through its other services. In this respect, more than thirty years after the breakup of the Bell System in 1982 - the dominant firm of the US telecommunications sector at that time - it is still an open and key question how should an economist look at the role of companies like Google, or the Bell System.

References

Asker, John. 2010. "A study of the internal organization of a bidding cartel." *American Economic Review*, 100(3): 724–62.

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.

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

[28]The reason why we see the other two trends as less problematic are as follows. First, the switch from GSP to VCG auctions is not troubling because it would fix only the most detrimental types of effects produced by bid coordination on revenues and efficiency (see Decarolis, Goldmanis and Penta [2017]). Second, disintermediation entails a choice by advertisers and, hence, we should expect the platform to offer valuable options to induce the advertisers to abandon their DMA. For instance, Google is promoting its bidding services in direct competition with the networks by emphasizing the good outcomes it achieved for advertisers, like Hewlett-Packard that switched from its DMA to Google's DoubleClick Search for all its real time bidding.

Balseiro, Santiago, and Ozan Candogan. 2017. “Optimal Contracts for Intermediaries in Online Advertising.” In *Operations Research.* Vol. 65, , ed. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, 878–896.

Balseiro, Santiago R., Omar Besbes, and Gabriel Weintraub. 2015. “Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design.” *Management Science*, 61(4): 864–884.

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.

Caliskan, Aylin, Joanna J Bryson, and Arvind Narayanan. 2017. “Semantics derived automatically from language corpora contain human-like biases.” *Science*, 356(6334): 183–186.

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.

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.

Choi, Hana, Carl F. Mela, Santiago Balseiro, and Adam Leary. 2018. “Online Display Advertising Markets: A Literature Review and Future Directions.” National Bureau of Economic Research Working Paper 18-1.

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.

Day, D. 2014. “Matching with Conflicts: An Application to the Advertising Industry.” mimeo.

Decarolis, Francesco, Maris Goldmanis, and Antonio Penta. 2017. “Marketing Agencies and Collusive Bidding in Online Ad Auctions.” National Bureau of Economic Research Working Paper 23962.

Decarolis, Francesco, Maris Goldmanis, and Antonio Penta. 2018. "Advertisers' Network Formation: Evidence from the Search Auctions."

De Loecker, Jan, and Jan Eeckhout. 2017. "The Rise of Market Power and the Macroeconomic Implications." National Bureau of Economic Research Working Paper 23687.

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

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.

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

EU Commission. 2018. "Fairness in platform-to-business relations." *Technical Report*.

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.

Graham, Daniel A., and Robert C. Marshall. 1987. "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions." *Journal of Political Economy*, 95(6): 1217–1239.

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

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.

Hendricks, Ken, Robert Porter, and Guofu Tan. 2008. "Bidding rings and the winner's curse." *The RAND Journal of Economics*, 39(4): 1018–1041.

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

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.

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.*

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

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

Pennington, Jeffrey, Richard Socher, and Christopher Manning. 2014*b*. "Glove: Global vectors for word representation." 1532–1543.

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.

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.