*American Economic Review 2021, 111(10): 1–30*
*https://doi.org/10.1257/aer.20190811*

# From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising[†]

*By* Francesco Decarolis and Gabriele Rovigatti*

*This paper analyzes the impact of intermediary concentration on the allocation of revenue in online platforms. We study sponsored search documenting how advertisers increasingly bid through a handful of specialized intermediaries. This enhances automated bidding and data pooling, but lessens competition whenever the intermediary represents competing advertisers. Using data on nearly 40 million Google keyword auctions, we first apply machine learning algorithms to cluster keywords into thematic groups serving as relevant markets. Using an instrumental variable strategy, we estimate a decline in the platform's revenue of approximately 11 percent due to the average rise in concentration associated with intermediary merger and acquisition activity.* (*JEL* C45, D44, G34, L13, L81, M37)

> *…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[1]

Online advertising sales are the main fuel of all of the major digital platforms. In the internet era, advertising means capturing the attention of consumers who are browsing the web and this requires both detailed data to effectively *target* the ad to the right customers and algorithms to bid in the online auctions where ad space is

* Decarolis: Department of Economics, Bocconi University, and IGIER (email: francesco.decarolis@unibocconi.it); Rovigatti: Research Department, Bank of Italy (email: gabriele.rovigatti@bancaditalia.it). Liran Einav was the coeditor for this article. We would like to thank anonymous referees, colleagues, and seminar participants in the EU, UK, and US for their many comments: competition authorities (DG Comp, CMA, and FTC), departmental seminars (BU Questrom, Duke Fuqua, MIT Sloan, Toulouse School of Economics, Paris School of Economics, Tilburg University, McGill University, and University of Mannheim), and conferences (NBER, MaCCI, NET Institute, IIOC, BECCLE) where earlier drafts of this study were presented. Decarolis gratefully acknowledges financial support from the Italian Ministry of Education grant PRIN (2017TMFPSH). IRB approval for the Amazon Mechanical Turk survey from Bocconi University ECR (SA000267).

[†] Go to https://doi.org/10.1257/aer.20190811 to visit the article page for additional materials and author disclosure statements.

[1] WPP founder and former CEO, WPP's 2012 Annual Report.

sold. These needs have led to a major, but understudied, shift in the industry: rather than bidding individually, advertisers increasingly delegate their bidding to highly specialized intermediaries. This concentration of demand within a few large intermediaries raises the question of countervailing buyer power. Can the emergence of intermediaries counterbalance the highly concentrated supply of online ads?

This study presents the first empirical assessment of how demand concentration at the level of intermediaries affects the creation and allocation of revenue from digital ads. Our setting is that of sponsored search ads, a market worth $40 billion in 2017 in the United States, which represents about half of all internet advertising revenue (IAB 2018). The supply side of this market is highly concentrated: 75–80 percent of search advertising revenue were earned by Google in 2016–2018 (eMarketer 2018). Advertisers, the demand side, compete against each other to buy one of a limited number of "slots" available on the search engine's result page for a given search term or *keyword*. In the early days of this market, advertisers used to operate individually but, over time, more and more ad buying is conducted through intermediaries. In our data, intermediaries are involved in about 75 percent of the slots sold.

Moreover, while thousands of intermediaries operate in the market, most of them belong to an agency network (henceforth *network*). Thus, essentially only seven networks are responsible for collecting data and optimizing bidding algorithms for most advertisers.[2] In the 2014–2017 period covered by our data, the four largest networks had a market share of approximately 70 percent of search volume and it was growing over time. The main goal of this study is to quantify the extent to which such increases in intermediary concentration affect platform revenue.

The research strategy is based on three ingredients. The first one is a novel dataset built by combining multiple sources. We obtained from Redbooks (2017)—the most comprehensive database on marketing agencies (MAs)—the list of the 6,000 largest US online advertisers. For these advertisers, we observe the MAs that they are affiliated with, together with the network that each individual agency belongs to. We combine this with data on Google's sponsored search auctions from SEMrush (2017), a major data provider for MAs.[3] For all of Redbooks advertisers, we know which keywords, if any, they bid on in Google's auctions. For each keyword, the data includes the position of the domain (and consequently of the advertiser) in the search results page, the volume of searches (i.e., the average number of search queries for the given keyword in the last 12 months) and the keyword-specific average price advertisers paid for a user's click on the ad (cost-per-click, or CPC).

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

[2] The seven networks are Interpublic Group of Comapnies (IPG), WPP, Publicis Groupe, Omnicom Group, Dentsu-Aegis, Havas, and MDC.

[3] Hereinafter, Redbooks (2017) shall be referred to as Redbooks, and the SEMrush (2017) shall be referred to as SEMrush.

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

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

We find with both ordinary least squares (OLS) and IV estimates that greater network concentration induces lower search engine revenue. Under our baseline IV model, a change in the Herfindahl-Hirschman index (HHI) of 245 points—the average HHI increase observed across the markets experiencing a merger event—leads to an 11.3 percent decrease in revenue. This quantitatively large estimate should be interpreted as a short run response, ignoring a series of changes to the auction environment that the selling platform has implemented in more recent years not covered in our data sample. In particular, we discuss a handful of recent trends in the market, from disintermediation to changes in auction reserve prices and the reduction in ad slots, that can be interpreted as a response from Google to the increased strength of intermediaries.

Furthermore, we analyze the mechanisms behind the baseline estimates, showing how the decline in revenue is driven by lower keyword ad prices. Indeed, we find that demand concentration is negatively associated with the average CPC, but not with the number of keywords or their search volume. We offer explanations of this effect based on both algorithm capabilities—in terms of bid coordination and keyword markets segmentation—and network bargaining strength.

Our findings represent a threefold contribution. First, they show the importance of countervailing power in the ongoing debate on concentration in digital markets and superstar firms.[4] Galbraith (1952) famously remarked that "the best and established answer to economic power is the building of countervailing power: the

[4] See, among others, Autor et al. (2017); De Loecker, Eeckhout, and Unger (2020); Werden and Froeb (2018); Gutierrez and Philippon (2017); and Weche and Wambach (2018); as well as the Obama administration's CEA (2016). Specifically on concentration in digital markets, see also the policy reports: Stigler Committee Report (2019), the Furman Review for the UK government (2019), the Competition Policy for the Digital Era report by the European Commission (2019), and the UK Competition and Markets Authority Report on Online Platforms and Digital Advertising (2019).

trade union remains an equalizing force in the labor markets, and the chain store is the best answer to the market power of big food companies." Our analysis illustrates how the market power of Google has been partially eroded by technological innovations and concentration among buyers. Although there is a vast literature on countervailing power with examples ranging from the US health insurance sector to the UK grocery market, this study offers its first, systematic application in the context of digital ad platforms.[5] Furthermore, since auctions play a key role in the mechanism through which we find buyer power operates, and since they are ubiquitous in digital platforms, the lessons learned from this study likely apply more broadly to digital markets working through auctions.[6] From a policy perspective, the evidence provided in this study supports the idea proposed by some observers that buyer power might serve as a remedy for the imbalance of bargaining power in favor of the digital platforms.[7] In the conclusions, we discuss the pros and cons of buyer power relative to the alternative policy interventions that are currently being debated.

Second, this study develops a novel measure of market definition for keyword ads. This is crucial to studying concentration and its effects due to the well known problem of the inadequacy of industry-level data (Berry, Gaynor, and Scott Morton 2019). In our setting, this problem emerges as a marked difference between industry-level and market-level estimates. The proposed approach is based on the use of machine learning algorithms in economics (Mullainathan and Spiess 2017; Agrawal, Gans, and Goldfarb 2019) and it is close to Hoberg and Phillips (2016) who pioneered this approach by employing a systematic text-based analysis of firm 10-K product descriptions to construct product similarity measures. Relative to that study, our clustering approach uses a different algorithm and can be implemented for all firms bidding in search auctions, regardless of whether they file 10-K forms or not.

The third and most direct contribution is to the understanding of online advertising. This is a particularly complex, economically relevant, and rapidly evolving market.[8] By focusing on the role of intermediaries, our study offers new insights into the firms that have practically taken over modern advertising markets, but whose role is not yet fully understood. In fact, we complement a small number of recent studies that have looked at these players (see the review in Choi et al. 2019). These works mostly emphasize the positive roles of intermediaries in improving the use of information to limit winners' curse risks (McAfee 2011), and in more

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

[6] In contrast to the existing buyer power literature—mostly centered around bargaining models—the focus on auctions makes our study close in spirit to the classic work of Snyder (1996).

[7] See Mullan and Timan (2018) for a discussion of the merits of this type of policy response. See also Loertscher and Marx (2019) for an analysis of the competitive effects of mergers in markets with buyer power.

[8] The existing studies on online ads have mostly focused on their effectiveness (see Goldfarb 2014; Blake, Nosko, and Tadelis 2015; Golden and Horton 2018; Johnson, Lewis, and Reiley 2017; Simonov, Nosko, and Rao 2018; and Simonov and Hill 2019), the functioning of the selling mechanisms (see Edelman, Ostrovsky, and Schwarz 2007; Varian 2007; Athey and Nekipelov 2014; Borgers et al. 2013; Balseiro, Besbes, and Weintraub 2015; and Celis et al. 2015) or platform competition (see Prat and Valletti 2019).

effectively administering client budgets in order to avoid the inefficiencies associated with budget constrained bidders (Balseiro and Candogan 2017). A handful of theoretical studies have, however, highlighted the downside of intermediary concentration: the vulnerability of online ad auctions to collusive bidding through common bidding agents (Bachrach 2010; Mansour, Muthukrishnan, and Nisan 2012; and Decarolis, Goldmanis, and Penta 2020). While only the latter is directly applicable to search advertising, all three studies focus on bidding equilibria within a one-shot auction. Our empirical analysis differs by looking more broadly at how the market works and, in this respect, allows us to account for both positive and negative effects of intermediary concentration and for the multiple mechanisms through which intermediaries operate—not only bid price coordination, but also ad targeting and keyword selection.

Finally, it should be remarked that the intermediary strategies that we describe below are proper from a legal perspective. They would not constitute a violation of the antitrust laws in the United States or the European Union because intermediaries are legal entities, independent from advertisers, operating unilaterally to maximize their profits. As such, they can freely decide how to arrange bidding strategies on behalf of their customers.[9]

**AQ1** The paper proceeds as follows: Section I presents a basic theoretical framework; Section II describes the data and market institutions; Section III reports a descriptive analysis of the data; Section IV illustrates the empirical methodology; Section V contains the results; Section VI finally concludes.

## I. Basic Framework

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

Now suppose that advertisers do not bid directly on the search auction. They submit their bid to an intermediary who internally runs a second price auction amongst its clients (we shall refer to this as the *intermediary auction*) and then bids on their behalf in the search auction. To see why this can affect the functioning of the search auction, consider the two cases illustrated in Figure 1. In panel A, each advertiser bids through a different intermediary, which we indicate as $\alpha$, $\beta$, and $\gamma$. In this case, intermediaries have no incentive to distort bids in the search auction. Hence,

[9] Indeed, outside specific cases, like those covered by the US Department of Agriculture (Coatney 2014), "common bidding agents" are not per se illegal. However, a caveat is that two situations might imply an antitrust infringement. The first case involves the discipline on "hub and spoke" cartels, (Harrington 2018), which would apply if it could be proved that advertisers had agreed to delegate their bidding to a common intermediary with the explicit intent of enforcing price coordination or market splits. The second case involves the discipline on purchasing agreements, or group purchasing organizations (GPOs). Although the intermediaries that we study are not GPOs, under the EU law, the limits to the activity of GPOs may be invoked if one could show that an intermediary controls such a large share of the market that its coordination activity could hurt Google's revenue to the point of leading to a worsening of the quality of its services.

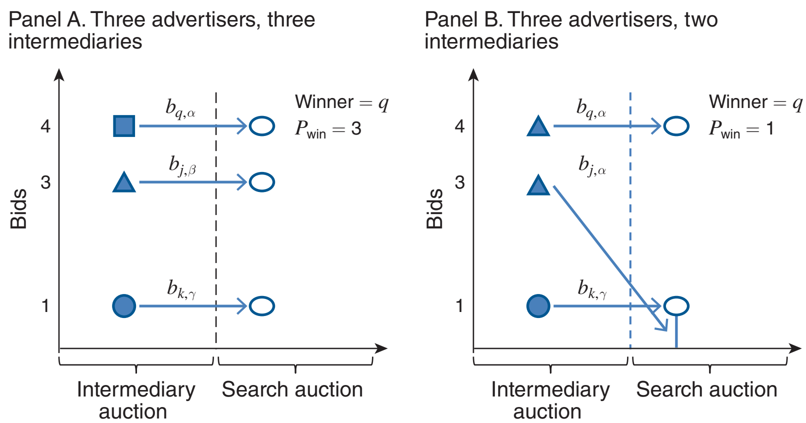


Figure 1. An Example of Bidding through Intermediaries

*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 ($\alpha$, $\beta$, or $\gamma$) to which they are affiliated. Intermediaries then bid in the search auction. In panel A, each advertiser has a different intermediary. In panel B, $q$ and $j$ share intermediary $\alpha$. The arrows indicate how the intermediary translates the bids in its own auction into the bids placed on the search auction. In panel A, bids are passed without distortions; in panel B, $j$'s bid is reduced. In both cases $q$ wins, paying the second highest bid which is either 3 (panel A) or 1 (panel B).

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 two intermediaries: both $q$ and $j$ bid through intermediary $\alpha$. This intermediary can now alter the search auction outcomes by retaining or amending the bids it places on behalf of its two clients: it can report just the highest bid among the two, $b_{q,\alpha} = 4$, or both bids, but setting $b_{j,\alpha} \in [0,1]$. In all these cases, $q$ wins and pays only 1 instead of 3, thanks to the reduction in $b_{j,\alpha}$. This example provides intuition on how intermediary concentration may lower the CPC in an auction, and consequently the search engine's revenue. Implementing this in practice would not be so simple for an intermediary handling many advertisers active over thousands of keywords, each with its own competitive structure dynamically evolving over time. Although algorithms for bid coordination in search have been proposed, keyword multiplicity allows for a simpler form of coordination: market split by keywords. Consider a modification to the example above where there are three "branded keywords" associated with the brands of each of the three advertisers. As in a prototypical prisoner's dilemma, all advertisers might be better off bidding only on their own brand, but—absent coordination—they bid on rival brands too. Explicit coordination by advertisers is illegal.[10] However, delegation to a common intermediary that autonomously implements the market split is a solution to the dilemma that would not be in breach of the law. But

[10] In 2019, the Federal Trade Commission (FTC) charged 1-800 Contacts Inc. for having entered into bidding agreements with at least 14 competing online contact lens retailers that eliminated competition on branded keywords search advertising. The FTC decision has been appealed and the appeal decision is currently pending.

for search engine revenue, the effect of advertisers splitting keywords in such a way can be rather dramatic: equilibrium bids in the generalized second price (GSP) auction are interlinked so that, if a bidder exits, this will typically cause the remaining bidders to drop their bids. Through this channel, even small changes in intermediary concentration might trigger large revenue losses. We illustrate this point further in online Appendix H by using a numerical example.

In addition to market splitting by keyword, intermediaries can exploit ad targeting to segment the markets for the same keyword. Consider a simple algorithm that targets ads on two dimensions: Google Ads allows geographical targeting (down to the ZIP code level) and schedule targeting (down to 15 minute intervals). For most keywords, however, the traffic volume is not so finely differentiated. This means that an algorithm that rotates bids between two advertisers so that they never meet in these ZIP code/quarter of hour intervals would expose these advertisers to the same audience, but without making them compete. Considering that many other targeting parameters are usually feasible, the number of possible market segmentations is nearly infinite.

What all of the strategies above have in common is that they would induce a lower CPC and, through that, lower revenue for the search engine. This negative effect, however, need not be the final outcome of intermediary concentration. Indeed, intermediaries can foster entry, by bringing to these auctions advertisers who would otherwise not enter. Moreover, thanks to their superior bidding technology, they can also bolster the number of keywords on which advertisers bid.[11] We will later show that the effect of lowering CPC prevails over the others. Before that, we illustrate below the data and the main industry facts.

## II. Industry Background and Data

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

The supply side, historically dominated by Google, has recently seen the emergence of new competitors (e.g., Amazon). Meanwhile, the demand side has experienced the emergence of new players—the intermediaries—which connect demand and supply of ads on several platforms, including search engines. There

[11] Furthermore, intermediaries can play other important roles. They can help internalize externalities (Jeziorski and Segal 2015; Gomes, Immorlica, and Markakis 2009): 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. In the closely related context of ad exchanges, the literature has identified further problems related to limited information leading to winners' curse (McAfee 2011) and budget constraints leading to inefficiencies (Balseiro and Candogan 2017) that a common intermediary might help to solve.

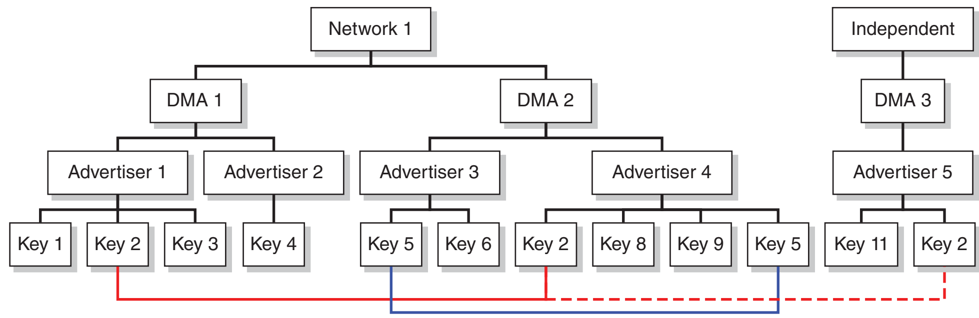


AQ2

Figure 2. Redbooks-SEMrush Data Structure

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

are two relevant levels of intermediation: (i) the marketing agencies, which are directly commissioned by advertisers to design, manage, and optimize marketing campaigns on their behalf; and (ii) the agency networks (or holding companies), which own most of the major marketing agencies in the United States and manage the bidding activities on behalf of their clients via centralized entities called "agency trading desks" (ATDs). The latter exploit automated bidding systems that allow for data-intensive targeting strategies with limited human intervention. ATDs represent the demand-side technological response to the incentive to improve bidding performance using better data and algorithms. For our analysis, the presence of ATDs implies that the most appropriate level at which we should analyze intermediary concentration is that of agency networks.

The dataset that we use allows us to observe a large set of keyword auctions, the advertisers bidding on them and their intermediaries, both at the marketing agencies and at the network level (when applicable). Indeed, the minimal data requirements to test the effects of intermediary concentration on the search engine's revenue are information on (i) the advertisers' affiliation to intermediaries, (ii) the set of keywords on which they bid, and (iii) the associated average CPC and search volume of these keywords. Our new dataset contains all this information, and more. Figure 2 shows the hierarchical structure of the data: the highest level (the networks, for nonindependent agencies) group the individual marketing agencies (layer 2). These, in turn, cluster the advertisers (layer 3), each bidding on a different set of keywords (layer 4). Solid lines indicate the cases of coalitions: in Figure 2, for example, Agency 2 participates in the auction for *key*5 on behalf of both Advertiser 3 and Advertiser 4. But we also consider *key*2 as having a coalition because Advertiser 1 and Advertiser 4 both bid through Network 1, although via different marketing agencies.

From Redbooks, a comprehensive database on marketing agencies, we obtained a list of advertisers representing all the major US firms active in online marketing (see Dai 2014 for an application of these data to the pharmaceutical sector and for a review of other studies using Redbooks data). For each of these advertisers, the

Redbooks data give us the full list of marketing agencies. The data are 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. Each advertiser is associated with one of the 23 industries in which Redbooks classifies advertisers. Starting in 2014, we also have access to a linkage variable that relates each individual agency to its agency network, if any. Overall, there are seven networks and about a thousand independent agencies.[12]

We combine the data on intermediaries with sponsored search data from the most comprehensive provider of digital ad data, *semrush.com* (henceforth, SEMrush). For keywords searched on Google, it collects the identity and website of advertisers appearing in the sponsored ad slots. Moreover, it gathers information on the keyword-specific average CPC; the position of the ad in the search outcome page; the volume of searches associated with the keyword; the visible URL of the ad; the traffic (that is, the share of the advertiser's traffic associated with the specific keyword); and the organic links. Thanks to the visible URL and the advertiser name, we are able to link Redbooks and SEMrush data for the years 2012–2017. Although the SEMrush data are available at a relatively high frequency (up to daily for certain keywords), we use the yearly average to match the frequency in the Redbooks data. CPC, volume, and traffic are monthly averages, calculated over the past 12 months.[13] Although these averages are calculated through proprietary algorithms that we could not inspect, they are considered reliable (and widely used) by marketing agencies and individual advertisers (see the online Appendix for a more extensive discussion of the data). Whilst the use of yearly averages implies foregoing some of the richness in the geographic and time dynamics in keyword bidding, this is necessary to match the two data sources and it is adequate to address our research question involving aggregate impacts at the level of markets (i.e., groups of keywords, as discussed below).

Table 1 presents summary statistics, by keyword and advertiser type (panel A) and by network (panel B). In the left columns of panel A, we report the statistics for keywords with at least one network advertiser; in the right columns, those for keywords with at least one independent advertiser (i.e., an advertiser bidding either autonomously without any marketing agency or through an agency not affiliated to any network). The two groups are thus not mutually exclusive. For both of them, we see a similar CPC. In terms of volume, for both groups the substantially lower value of the median relative to the mean indicates a tendency to bid on keywords that are infrequently searched. The lower value of *Traffic* (1 percent) observed for the network advertisers relative to the 6 percent for the non-network advertisers is compatible with the former placing ads over more keywords. *Coalition* measures the cases of keywords where more than one of the ads shown belongs to different advertisers represented by the same agency network. Within this subset of cases, *Coalition size* shows that the average number of advertisers bidding in a coalition is 2.36 and, indeed, the vast majority of cases involve coalitions of size 2. In essentially all

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

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

Table 1—Summary Statistics: Keywords, Networks, and Markets

| | Keywords with at least one network, years 2014–2017 | | | | Keywords with at least one independent, years 2014–2017 | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean | SD | Median | Observations | Mean | SD | Median | Observations |
| *Panel A. Statistics by keywords and advertiser type* | | | | | | | | |
| CPC | 2.33 | 5.76 | 0.90 | 15,140,935 | 2.36 | 5.97 | 0.89 | 21,683,735 |
| Volume (000) | 503 | 35,198 | 40 | 15,140,935 | 360 | 99,677 | 40 | 21,683,735 |
| Traffic | 0.01 | 0.53 | 0.00 | 15,140,935 | 0.06 | 1.27 | 0.00 | 21,683,735 |
| Number of advertisers | 1.30 | 0.69 | 1.00 | 15,140,935 | 1.22 | 0.54 | 1.00 | 21,683,735 |
| Organic results | 47.16 | 257.44 | 1.78 | 15,140,935 | 38.30 | 193.31 | 1.57 | 21,683,735 |
| Number of characters | 22.78 | 7.74 | 22.00 | 15,140,935 | 22.85 | 7.58 | 22.00 | 21,683,735 |
| Number of words | 3.71 | 1.35 | 4.00 | 15,140,935 | 3.66 | 1.30 | 3.00 | 21,683,735 |
| Long tail | 0.50 | 0.50 | 1.00 | 15,140,935 | 0.48 | 0.50 | 0.00 | 21,683,735 |
| Branded | 0.10 | 0.30 | 0.00 | 15,140,935 | 0.07 | 0.25 | 0.00 | 21,683,735 |
| Coalition | 0.15 | 0.36 | 0.00 | 15,140,935 | 0.00 | 0.00 | 0.00 | 21,683,735 |
| Coalition size | 2.36 | 0.68 | 2.00 | 339.779 | | | | |
| | Market share (search volume share) | | | | Presence across keywords | | | |
| | 2014 | 2015 | 2016 | 2017 | 2014 | 2015 | 2016 | 2017 |
| *Panel B. Statistics by network* | | | | | | | | |
| IPG | 0.21 | 0.19 | 0.21 | 0.19 | 0.26 | 0.33 | 0.34 | 0.39 |
| WPP | 0.17 | 0.20 | 0.17 | 0.23 | 0.29 | 0.29 | 0.34 | 0.43 |
| Omnicom | 0.17 | 0.16 | 0.18 | 0.14 | 0.39 | 0.38 | 0.38 | 0.38 |
| Publicis | 0.14 | 0.13 | 0.13 | 0.18 | 0.30 | 0.30 | 0.30 | 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 | 0.13 | 0.06 |
| Dentsu-Aegis | 0.05 | 0.08 | 0.11 | 0.09 | 0.14 | 0.17 | 0.21 | 0.25 |
| Ind agency | 0.13 | 0.09 | 0.06 | 0.06 | 0.42 | 0.38 | 0.32 | 0.22 |

*Notes:* Panel A: statistics at the keyword level, separately for keywords where at least one ad comes from either a network bid (columns 1 to 4) or a non-network marketing agency bid (columns 5 to 8). The variables included are *CPC*, reported in US$; *coalition*, an indicator function for the presence of multiple advertisers affiliated with the same network participating to the keyword auction; *coalition size*, which is populated for keywords with coalitions only. Both *long tail* and *branded* refer to the type of keyword: the first indicates those composed by at least four terms and the latter those including the name of a brand. *Organic results* reports the number of nonsponsored search results returned by the search engine (rescaled by one million). Panel B: on the left (columns 1 to 4) we report for the years 2014–2017 the market share (in terms of *search volume*) of the seven network and non-network marketing agencies; on the right (columns 5 to 8), we report the presence of the networks across all keywords in our data—the sum of these values within columns does not add up to one since the same keyword can display ads from multiple network and non-network bidders.

cases where there is a coalition, there is exactly one coalition, suggesting that different networks are specialized in different segments of the keyword markets.

Panel B shows the relative size of each of the seven networks, both in terms of the volume of searches covered and in terms of their presence across keywords. If we consider just the largest four networks—the "big four" as they are often referred to (WPP, Omnicom, Publicis, and IPG)—their combined market share (in terms of search volume) reaches 74 percent of the total volume in 2017. The situation is similar across years and concentration tends to increase over time. The situation is also similar if we look at the network presence across keywords. The sheer prominence

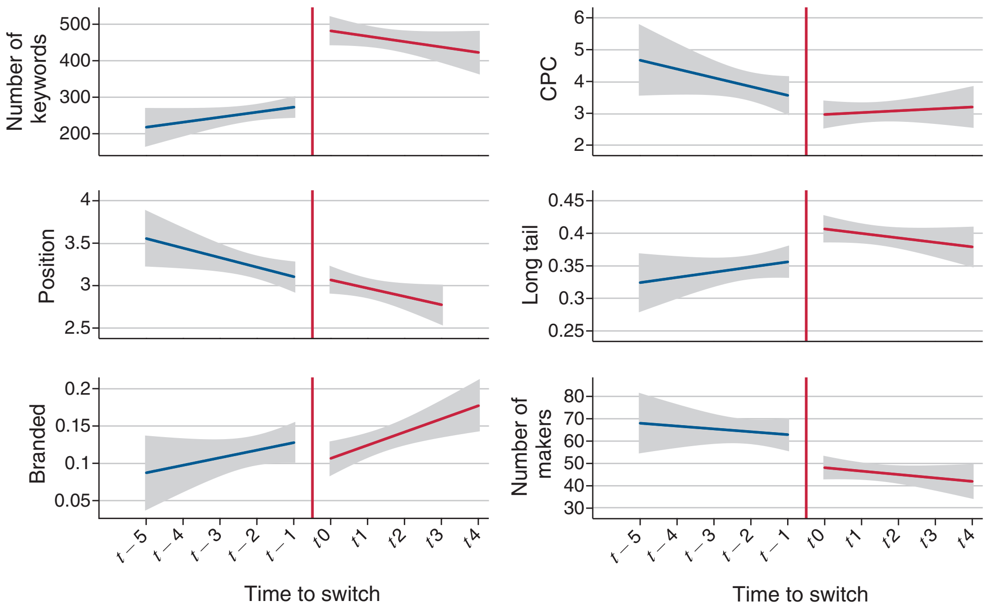


FIGURE 3. INDIVIDUAL ADVERTISERS JOINING MAS

*Notes:* Blue (maroon) lines are linear fits of average values before (after) joining an agency at *t*0 (red vertical line). The reported variables are (left to right and top to bottom) *number of keywords*, *CPC*, *position*, *long-tail keywords*, *branded*, and *number of markets*. *CPC* value is reported in US$; the shaded area corresponds to the standard deviation of the mean. The sample covers the 2012–2017 period.

of networks in the data, together with bidding centralization at the network ATD level leads us to consider networks as the key players in our analysis.

## III. Keyword-Level Descriptive Evidence

In this section, we use keyword-level data to perform a descriptive analysis of whether the choices of intermediaries are indicative of the types of strategies adopted.

### A. *Individual Advertisers Joining a Marketing Agency*

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

However, it is risky to analyze the effects of intermediary concentration by looking at concentration increases driven by the incorporation of formerly independent bidders joining MAs. Some advertisers might join an agency due to their inability to optimize bidding. But then the lower CPC after joining might be explained by excessively high bids in the previous period, rather than by bid coordination by the intermediary. In the analysis below, we therefore rely on a different type of variation in demand concentration: the one produced by ad networks incorporating previously independent MAs. In these situations, it is reasonable to assume that bids are already optimized from an individual bidder's perspective and that any strategy change is driven by the intermediaries' incentives to coordinate their advertisers' actions, as described earlier.

### B. *Network Expansions via M&As*

A second dimension along which keyword-level data is informative regards market splits by keyword. For this, we analyze changes in the composition of advertisers' keywords after their agency is acquired by a network and, in particular, we ask whether there is any change in the overlap in the sets of keywords of the clients of either the network or the acquired MA.

The data reveal substantial heterogeneity across networks. For instance, Figure 4 shows two polar cases. When the MDC network acquired the Forsman MA, most of the keywords that used to be shared by clients of both MDC and Forsman before the merger stopped being shared afterward (pink area), and only a few new common keywords were introduced (green area). Instead, when WPP acquired *SHIFT Communications* most of the keywords that were shared before the merger continued to be shared after it (brown area) and new shared keywords were also introduced (green area). Overall, the great variety of possible strategies and the heterogeneity across networks in their usage make it difficult to quantify their impacts. Thus, in the next section we propose an empirical strategy using market-level data to quantify the effects of intermediary concentration on Google's revenue.[14]

[14] Due to the availability in the Redbooks data of the link between MAs and their network only from 2014, the analysis of agency networks presented here (and in the next sections) is limited to the 2014–2017 period.

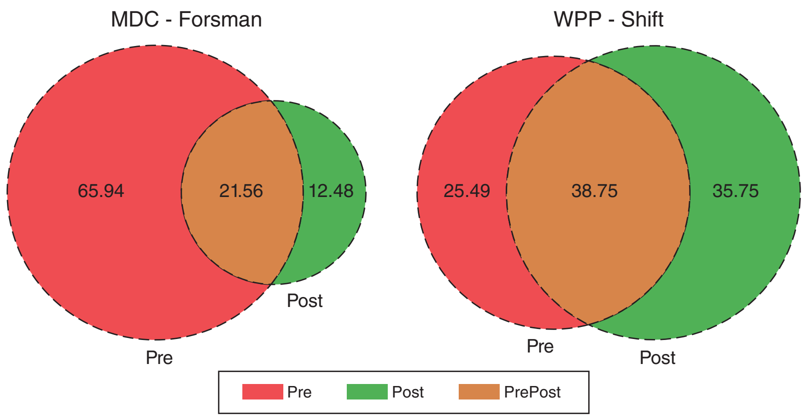


FIGURE 4. NETWORK EXPANSIONS VIA AGENCY ACQUISITIONS

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

## IV. Market-Level Empirical Strategy

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

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

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

**Step 1:** Market Definition

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

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

*Keyword Vectorization*.—For each keyword appearing in SEMrush data, we need a vector representation. The reason is straightforward: "red car," "blue car," and "automobile" are three keywords that we would like to see grouped together, but using keyword match approaches (e.g., using matches between single words), only "red car" and "blue car" would be pooled together. The vector representation systems developed in the natural language processing literature are meant to directly address the issues related to synonyms and antonyms in text clustering or semantic similarity exercises. We use an unsupervised learning algorithm (GloVe, developed by Pennington, Socher, and Manning 2014) to obtain vector representations for each term within the keywords. The GloVe model is a word embedding system which builds on the classical matrix of word co-occurrences in a corpus—i.e., a sparse matrix with one row per document in the corpus, and one column per word, populated with the number of occurrences (see details in the online Appendix). We use a GloVe dataset pretrained on 840 billion documents, corresponding to approximately 2.2 million unique terms, from Common Crawl in English, featuring 300 dimensions. Such an extensive corpus originating from mimicking the web surfing behavior of typical internet users makes the resulting vectorization analogous to surveying people about the proximity between keywords.[15] Similarly, when applied to the sponsored search keywords in our data, the vectorization should reflect the proximity between products and services identified by the semantic similarity between keywords. Once every keyword is split into its constituent *terms*, we proceed by merging every term with the corresponding GloVe vector. Finally, we obtain the vector representation of each keyword by summing together the vectors relative to all its constituent terms.

[15] The dataset and GloVe code are available at https://nlp.stanford.edu/projects/glove/.

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

*Layer 2: Competitive Clustering*.—To incorporate supply-side information into the clusters, we exploit the competitive structure within each thematic cluster to further subdivide them into what we will refer to as "markets." The basic idea is to pool together keywords that are close in terms of the set of advertisers bidding for them. This is implemented by constructing, separately for each thematic cluster, a sparse matrix whose rows correspond to the keywords in the cluster and whose columns match the advertisers that bid on at least one of these keywords. The resulting row vectors are projections of the keywords in the space spanned by the advertisers (which we consider, to all extent and purposes, the competitive structure space). Through such vectors, we build for each pair of keywords a measure of similarity (the Euclidean distance between the corresponding row vectors).[16] Finally, we feed the similarity matrix describing the proximity of each pair of keywords into a hierarchical clustering algorithm to partition the keywords into "markets."

Table 2 reports summary statistics for the subset of thematic clusters and markets.[17] In the top-right panel of Table 2, the summary statistics indicate that an average market has 37 keywords and 4 competing advertisers, with the number of competing advertisers within single keywords (not reported) being on average 1.62. The statistics in the top-left panel further show that there are on average five markets within a thematic cluster. The bottom panel of this table reports summary statistics for the market-level variables that we describe below. Before moving

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

[17] Since many clusters are composed of keywords that contribute either very little or nothing to the search engine's revenue, and are never involved in any of the changes in agency ownership that we exploit for the IV strategy, we keep in the analysis sample only markets that either experience variation in the instrument at least once during the sample period or, for the remaining ones, that are in the top quartile of revenue. This leads us to drop markets that represent between 1 percent and 2 percent of the total yearly revenue. In the online Appendix, we report robustness checks regarding this sample selection.

Table 2—Market-Level Descriptives, Thematic and Competitive: Analysis Sample

| | Thematic clusters | | | | Competitive clusters (markets) | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean | SD | Median | Observations | Mean | SD | Median | Observations |
| *Market characteristics* | | | | | | | | |
| Number of advertisers | 6.70 | 10.50 | 3.00 | 8,324 | 4.00 | 4.80 | 3.00 | 25,947 |
| Number of keywords | 116.10 | 180.30 | 55.00 | 8,324 | 37.20 | 104.90 | 4.00 | 25,947 |
| Number of networks | 2.79 | 1.77 | 2.00 | 8,324 | 2.23 | 1.27 | 2.00 | 25,947 |
| Competitive clusters | 5.00 | 5.00 | 3.00 | 8,324 | - | - | - | - |
| *Market variables* | | | | | | | | |
| $\log(R_{m,t})$ | 10.89 | 2.27 | 10.92 | 29,796 | 10.41 | 1.96 | 10.37 | 52,476 |
| $HHI_{m,t}$ | 2,765 | 2,311 | 2,000 | 29,899 | 2,740 | 2,257 | 2,000 | 52,476 |
| Long tail | 0.32 | 0.35 | 0.18 | 29,899 | 0.27 | 0.37 | 0.01 | 52,476 |
| $\Delta R_{m,t}$ | −0.05 | 1.78 | 0.00 | 21,256 | 0.40 | 1.53 | 0.28 | 43,973 |
| Number of results (mil) | 76.93 | 269.19 | 21.52 | 29,899 | 75.98 | 231.28 | 19.70 | 52,476 |
| Number of clusters | | | 8,324 | | | | 25,947 | |

*Notes:* Top panel (*market characteristics*) reports the features of the thematic (left) and competitive clusters (right). The first three rows are the number of advertisers, keywords, and networks. *Competitive clusters* are the number of clusters identified by the hierarchical clustering algorithm in the second layer. In lower panels (*market variables*) we report relevant outcome and explanatory variables relative to the estimation sample: $\log(R_{m,t})$ stands for search engine's market revenue, $HHI_{m,t}$ is our demand concentration proxy, *long tail* is an indicator for keyword with four or more terms, $\Delta R_{m,t}$ is the yearly change in revenue and *number of results* is the number of organic results—in millions.

to that, however, we stress that we cannot directly test the quality of the clusters obtained as that would require a reference sample where keywords and markets are correctly associated. Nevertheless, lacking this type of sample, we resorted to random inspection of the cluster quality. Overall, we find very satisfactory results with our initial motivating concern of related but different keywords (like "car" and "automobile") systematically pooled together. Moreover, we designed and implemented a simple task aimed at testing the reliability of the clusters, and we ran it through *Amazon Mechanical Turk* (see the online Appendix for a description of the test design with some examples and the results). With the exception of the residual industry that pools together many heterogeneous advertisers (*miscellaneous*), for all other industries the share of correctly classified keywords is between 80 percent and 90 percent (see online Appendix D).

**Step 2:** Measurement of the Main Variables

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

*Outcome Variable*.—Suppose that the clustering procedure has identified $M$ markets, $m = 1, \ldots, M$. 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 revenue $(R)$ in market $m$ in period $t$ by aggregating revenue over keywords:

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

where $CPC_{kmt}$ is the average CPC of keyword $k$ (belonging to the set $K_m$ in market $m$ ) at time $t$, $Volume_{kmt}$ is its overall number of searches, and $CTR_{kmt}$ is the cumulative click-through rate (CTR) of all the sponsored ad slots shown for keyword $k$.[18] There is substantial heterogeneity in the levels of revenue across markets, mostly driven by heterogeneity in volume and CPC. To perform a meaningful analysis of the association of the revenue's level and the level of concentration, we thus work with $\log(R)$.

*Concentration Measure*.—Suppose we have a market $m$ defined by the set of keywords $K_m$. For each keyword $k \in K_m$, there are $J_k$ sponsored ad slots, each occupied by an advertiser $a$. Each of these slots brings a certain number of clicks, which are ultimately the advertisers' object of interest. We therefore measure the "market size" $(S_{mt})$ as the sum of all the clicks of all the ad slots allocated in all the keywords in market $m$. That is, $S_{mt} = \sum_{k \in K_m} Volume_{kmt} \times CTR_{kmt}$. The market share of intermediaries 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

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

Thus, our concentration measure for market $m$ at time $t$ is the squared sum of each intermediary's market share: $HHI_{mt} = \sum_{i=1}^{I} \left(s^i_{mt}\right)^2$.[19] As stated earlier, the intermediary is the network, or the independent MAs.

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

[18] For each $k$, the overall $CTR_{kmt}$ is the cumulative sum of the number of clicks across all $j$ ad slots appearing on the search outcomes page of keyword $k$: $CTR_{kmt} = \sum_{j \in J_k} CTR_{jkmt}$. Since CTRs are not part of our dataset, we supplement this information from Advanced Web Ranking (2017). As discussed in the online Appendix, although the CTR data are likely to involve measurement error, our baseline findings are qualitatively robust to two sets of robustness checks. First, we exclude entirely the CTR from the analysis by setting all CTRs to one (see online Appendix Table F.3 and F.4) and, second, we randomly rematch CTRs to keywords (see online Appendix Figure F.1).

[19] Despite several theoretical and practical shortcomings of the HHI (see O'Brien 2017), it is commonly used in both academia and competition policy as a proxy for concentration (see Hastings and Gilbert 2005; Dafny, Duggan, and Ramanarayanan 2012; and the US Horizontal Merger Guidelines). In our setting, the use of the HHI as a proxy for demand concentration has a theoretical foundation in the results of Decarolis, Goldmanis, and Penta (2020) and, moreover, it will be empirically implemented through an IV strategy to control for measurement error problems.

**Step 3:** Identification Strategy

There are two main reasons why the OLS estimation of equation (1) might lead to biased estimates of $\beta$. The first is the measurement error problem associated with the HHI being only an imperfect proxy of demand concentration. The second is the risk of an omitted variable bias. For instance, a keyword *k* might have become suddenly fashionable for some exogenous reasons, such as changes in consumer taste; advertisers that were previously not interested in *k* now hire an intermediary to bid for it. Moreover, they all hire the same intermediary as it is the one specialized in the market to which *k* belongs. This situation would likely induce the observation of a positive association between intermediary concentration and the growth of search engine revenue, but it does not imply the existence of a causal relationship between the two phenomena. In practice, the available data allow us to reduce the risk of an omitted variable bias in two ways. First, we can include among the set of covariates market-time varying observables (like the number of organic links) that can likely control for phenomena such as the sudden change in appeal of a keyword, as mentioned above. Second, we can include fixed effects for the thematic clusters, thus exploiting the cross-sectional variation across markets within a cluster. This clearly reduces the extent to which relevant factors might be omitted since, for instance, omitted demand factors should be controlled through the thematic cluster fixed effects.

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

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

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

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

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

where the subscript 0 denotes the year before the merger year $t^*$. We use, for each market-year, the variable $sim\Delta HHI_{mt}$ as instrument for $HHI_{mt}$. In total, there are 21 mergers in our sample (details on each merger are in online Appendix Table A.2).[21] Across networks, there is heterogeneity both in the number and the size of the MAs acquired. While Dentsu-Aegis appears to be the most "active" network with eight acquisitions (including the one with most clients, Merkle), WPP secured the largest acquisition in terms of presence in the markets (*SHIFT Communications* with clients active across 1,049 different markets). Some acquisitions take the form of hostile takeovers, with subsequent attempts to buy back independence and, as mentioned above, we observe two cases of divestitures. The effects of these M&As on the HHI measure described above are substantial: across markets affected by mergers, the average HHI increase between the year of the merger and the preceding year is 245 points.[22] For our baseline estimates, we will use an IV that exploits the variation from the whole set of M&A episodes. Clearly, the instrument's validity would be violated if the M&A operations were driven by expectations about revenue performance in the search auctions. In the online Appendix, we look in isolation at the larger merger episodes involving several clients active in many markets, as they are the least likely to be endogenously driven by revenue in local markets. Furthermore, the larger the merger, the more likely the companies interested will do advertisement activities outside Google's search auctions, thus making less likely their endogenous determination within our empirical framework.

Using $sim\Delta HHI_{mt}$ as instrument thus entails the following first stage:

$$(5)\qquad HHI_{mt} = \beta^{FS} sim\Delta HHI_{mt} + \phi X_{mt} + \tau_t + \gamma_z + \epsilon_{mt}.$$

## V. Results

We begin the presentation of our results from the first-stage and reduced-form estimates in Table 3. It reports the estimates for four different model specifications, gradually expanding the set of covariates. Model 1 includes demand concentration only, while model 2 adds thematic clusters fixed effects. Model 3 adds year fixed effects and a control for the number of organic results, which captures the "popularity" of the keywords in the market, thus reflecting the appeal to customers. This latter model is our baseline. Indeed, while model 4 includes further

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

[22] To put this number in perspective, consider that, according to the US Horizontal Merger Guidelines, when a merger results in an HHI increase of more than 200 points and a highly concentrated market, it will be "*presumed to be likely to enhance market power*."

AQ3

Table 3—Reduced-Form and First-Stage Estimates

| | (1) | | (2) | | (3) | | (4) | |
|---|---|---|---|---|---|---|---|---|
| | RF | FS | RF | FS | RF | FS | RF | FS |
| $sim\Delta\widehat{HHI}$ | −6.761 | 0.618 | −4.070 | 0.957 | −3.831 | 0.829 | −3.723 | 0.831 |
| | (1.110) | (0.170) | (1.133) | (0.0790) | (1.165) | (0.0915) | (1.165) | (0.0913) |
| Weak id. $F$-test | | 13.21 | | 146.99 | | 82.18 | | 82.94 |
| Underid. $F$-test | | 4.56 | | 13.67 | | 11.01 | | 11.02 |
| Observations | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 |
| Cluster FE | | | ✓ | | ✓ | | ✓ | |
| Year FE | | | | | ✓ | | ✓ | |
| Organic results | | | | | ✓ | | ✓ | |
| Keyword characteristics | | | | | | | ✓ | |

*Notes:* The dependent variable is the (log) revenue, $R_{mt}$. *RF* columns report the reduced-form estimates, *FS* columns the first-stage ones; the models 1 to 4 have an increasing number of controls and fixed effects. Model 1 includes industry fixed effects (FE). In the baseline model, reported in column 3, we control for the average number of organic results, thematic cluster, and year FE. Model 4, in which we add keyword characteristics such as the share of long tail and branded keywords, is likely to suffer from an additional source of endogeneity. In all models the standard errors are clustered at the thematic clusters level.

controls for the types of keyword composing the market (i.e., the average number of *long-tail* and *branded* keywords), we know from the earlier discussion that these might be endogenously determined by the strategies of intermediaries. Nevertheless, by way of comparison it is useful to report the estimates of model 4 as they offer a way to check whether these keyword choices affect revenue through increases in concentration.

Both the first-stage and reduced-form estimates in Table 3 are rather stable across model specifications. As expected, magnitudes are impacted the most by the addition of thematic cluster fixed effects between model 1 and model 2. We consider the latter level of clustering quite useful to control for most of the potential omitted variable bias and, therefore, rely on this cross-sectional variation within clusters as a main source of causal identification. In terms of the results, the positive sign of the $sim\Delta HHI$ estimate in the first-stage regression indicates that the HHI increases in the markets where the simulated HHI grows the most. This implies that the clients of an agency acquired by a network tend to remain within the acquired network. This result was not obvious ex ante. In fact, to the extent that there is persistency in the market shares, we would expect a positive sign, but a negative sign could reveal that advertisers prefer avoiding sharing MAs with rivals (i.e., "sleeping with the enemy," Villas-Boas 1994). Although the estimated coefficient of 0.829 falls short of 1, its large magnitude indicates that the "sleeping with the enemy" concern does not appear to drive a reshuffling of clients among acquired MAs.[23] The large value of the $F$-statistics also confirms the relevance of the proposed instrument.

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

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

| | OLS | | | | IV | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| $\widehat{HHI}$ | −2.217 | −2.120 | −2.122 | −2.130 | −10.93 | −4.252 | −4.620 | −4.479 |
| | (0.0718) | (0.0567) | (0.0572) | (0.0569) | (2.902) | (1.068) | (1.204) | (1.201) |
| Organic results (million) | | | 0.252 | 0.263 | | | 0.206 | 0.225 |
| | | | (0.0437) | (0.0458) | | | (0.0463) | (0.0477) |
| Keywords characteristics | | | | | | | | |
| Branded keyword | | | | 0.396 | | | | 0.458 |
| | | | | (0.0537) | | | | (0.0639) |
| Long-tail keywords | | | | −0.0908 | | | | −0.0491 |
| | | | | (0.0367) | | | | (0.0423) |
| Observations | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 | 52,476 |
| Cluster FE | | ✓ | ✓ | ✓ | | ✓ | ✓ | ✓ |
| Year FE | | | ✓ | ✓ | | | ✓ | ✓ |

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

On the other hand, the reduced-form estimates indicate a negative and statistically significant relationship between (log) revenue and the simulated change in HHI.[24]

Table 4 reports OLS (columns 1 to 4) and IV (columns 5 to 8) estimates. Both sets of coefficients are negative and statistically significant. IV coefficients are larger, being about twice the corresponding OLS ones. This is compatible with both measurement error in the demand concentration proxy and with residual omitted variable bias. As expected from the estimates in Table 3, there is a large drop in the magnitude of the coefficient of the IV estimates when controlling for thematic cluster fixed effects. With these fixed effects, the estimates are remarkably stable across all models, in terms of both magnitude and significance. Controlling for either organic results or keyword characteristics has quantitatively no impact on the findings. In the online Appendix, we present an extensive set of robustness checks of these baseline estimates.[25]

To ease the economic interpretation of the estimates, it is useful to recall that the average HHI increase induced by mergers of 245 points. Under the baseline estimate (column 7), such an HHI increase implies a decrease in revenue of 11.3 percent (that is $4.62 \times 100 \times 0.02451$).[26] While this magnitude might seem large, we recall from

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

[25] These robustness checks involve both restricting the analysis to the largest mergers where the IV assumptions are more likely to be satisfied and addressing measurement errors problems. Among the latter, it's worth mentioning that, if we use as an alternative definition of "markets," the advertisers' industries, the sign of $\beta$ flips and the magnitude grows to unreasonable levels; see column 1 in online Appendix Table F.2.

[26] Similarly, if instead of using 245 points, which is the average HHI change across all markets affected by a merger, we use 120 points, which is the average across all merger events of the merger-specific average HHI change, the implied effect is a decline in revenues of 5.5 percent.

Table 5—Revenues Components: IV Estimates

| | $\log(cpc)$ (1) | $\log(vol)$ (2) | $\log(\#keys)$ (3) |
|---|---|---|---|
| $\widehat{HHI}$ | −1.271 | −0.669 | −0.842 |
| | (0.427) | (0.983) | (0.741) |
| Observations | 52,476 | 52,476 | 52,476 |
| Cluster fixed effects | ✓ | ✓ | ✓ |
| Year fixed effects | ✓ | ✓ | ✓ |

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

the discussion in Section I that the GSP auction is particularly weak with respect to advertiser coordination: its lack of strategy proofness implies that the effect of any bid coordination by an intermediary is amplified by its equilibrium effect on advertisers that are not directly part of the concentration.[27] Furthermore, we should also stress that our estimates are best interpreted as a static response ignoring all those dynamic responses that Google might undertake to reverse its loss of revenue. As we discuss in the conclusions, ongoing changes like the reduction in the number of slots up for sale, the increase in reserve prices and the promotion of disintermediation services are examples of these dynamic responses that might effectively limit the intermediary ability to pull revenue out of the platform.

*Channels*.—The findings above indicate that the effects of increased buyer power dominate the efficiency gains from which the search engine might benefit. To better understand our findings, here we analyze the channels through which competition impacts revenue. In Table 5 we explore the relationship between market concentration and changes in the average CPC (column 1), search volume (column 2) and number of keywords (column 3). The estimates are noisy and not statistically significant for the latter two, but they are negative and strongly significant for the CPC. This latter estimate is in accordance with the theoretical predictions about the incentive to coordinate prices.

In interpreting this evidence, it is interesting to recall the descriptive evidence presented in Section III. While the estimates in Table 5 exploit variation driven by network M&A activity, the graphical evidence in Figure 3 is based on what happens when individual advertisers join agencies. As that figure shows, multiple changes occur and some, such as the expansion in the number of keywords, are clearly beneficial for the search engine. But when an intermediary is acquired by a network, the changes in the types of bidding behavior are more subtle and pertain exclusively to what is allowed by greater concentration within an intermediary. Thus, the lack of effects on the number and volume of keywords is indicative of demand concentration by itself which has mostly negative effects on the search

[27] For instance, in the numerical example of online Appendix H, a merger affecting two advertisers that entails a mere 33-point HHI increase causes platform revenue to drop by 18 percent.

engine revenue, whereas the activity of MAs more generally has both positive and negative effects on revenue.

The capacity of concentrated networks to lower the CPC helps to explain why advertisers use them instead of replacing them with their own bidding algorithms, despite the ease of developing such algorithms and the hefty fees of the networks (of the order of 17 percent of ad spending).[28] But what are the means through which networks can lower the CPC? We discussed this question with industry experts. Some experts highlighted a mechanical effect linked to the quality scores: demand concentration allows the larger intermediaries to pool together relevant data from rival advertisers and this allows them to attain improvements in the client quality scores, which mechanically implies lowering their CPC. The other answers that we got can be grouped into two broad strategies: easing competition among network clients and bolstering competition between ad selling platforms.[29] The first type of strategy involves employing bidding algorithms that optimize joint bidding within an auction[30] or that exploit the targeting features of Google Ads to reduce (or even eliminate) competition among clients of the same network.[31] Recall the example earlier about splitting the market for the same keyword by targeting two dimensions (geography and timing). On Google Ads, the set of targeting dimensions is extensive and has expanded over time. Currently, it includes demographics (six groups for age, six for income, and two for gender), device (computer, tablet, or mobile phone) and audiences (i.e., groups of people with specific interests, as estimated by Google).[32]

Market segmentation might also be implemented by splitting keywords. For instance, significant shares of marketing budgets are spent for ads' own brands and those of rivals (Blake, Nosko, and Tadelis 2015). In Figure 5, we apply the method by Dobkin et al. (2018) to describe the change in probability for both the other-brand (left panel) and own-brand (right panel) bidding before and after the merger, indicated by the dashed vertical line at $t^*$. We also add the linear fit, estimated in the period *before* the merger and projected in the post period. Other-brand bidding is clearly impacted negatively by the M&A event, with advertisers significantly bidding less on the brands of rivals after the merger; on the other hand, the own-brand bidding appears not to be negatively impacted, and instead records an upward jump at $t^* + 2$. The effectiveness of this type of brand splitting strategy is suggested by two recent studies, Simonov, Nosko, and Rao (2018) and Simonov

[28] The 17 percent figure is obtained as the sum of the fees for the agency of record (5 percent) and of the trading desk (12 percent) reported in Figure 6 in Adshead et al. (2019). ISBA (2020) also finds similarly large fees, as well as reporting the presence of large hidden fees. Both studies are based on display advertising.

[29] Selected quotes from the interviews are reported in an ad hoc online Appendix.

[30] A glimpse of what might be happening in practice can be grasped by looking at the case of iProspect—a leading independent MA, later acquired by the Dentsu-Aegis network. This company is credited with having developed one of the earliest automated bidding systems for search auctions. It is thus intriguing that the scientist who developed this algorithm is also the leading author of a computer science paper, Kitts, Laxminarayan, and Leblanc (2005), on cooperative strategies for search auctions that proposes "a coordination algorithm that optimally distributes profit on the auction between participating players" and shows its implementation in real data.

[31] Other features of the intermediary bidding process might also drive a reduction in market competition. A germane explanation might be increased experimentation, which intermediaries could use to evaluate and optimize bids. Randomizing two advertisers into 50/50 treatment and holdout groups implies that advertisers would only compete directly in a quarter of the markets. This explanation might, however, overstate the extent to which agencies resort to experimentation.

[32] See https://support.google.com/google-ads/answer/2732132?hl=en.

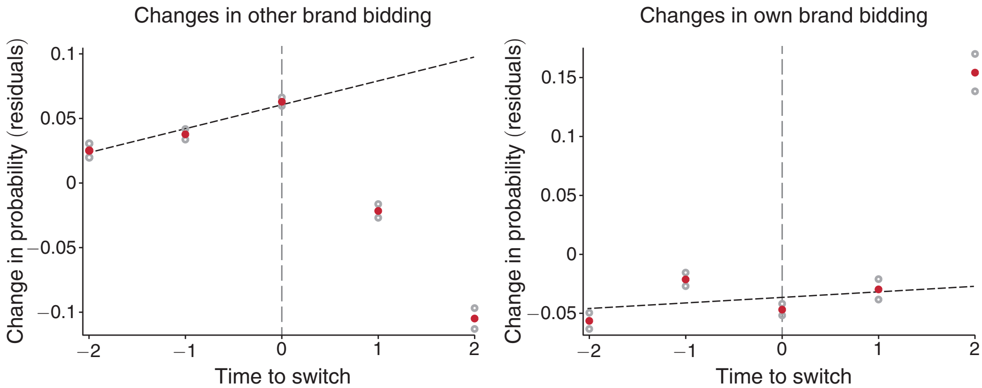


Figure 5. Changes in Own and Other Brand Keywords

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

and Hill (2019): by experimentally manipulating the number, order, and identities of paid links on the Bing search results page, their findings indicate that competitors steal traffic from the brand owner, but that this stolen traffic is often too costly given the low conversion rate on these clicks.

The second type of strategy involves increasing the competition level between ad selling platforms. The most straightforward way to do this is by splitting marketing budgets across more digital ad platforms. This form of market segmentation differs from those described earlier because its efficacy hinges on the availability of alternatives to search ads on Google. These alternatives mainly involve other search platforms (primarily Amazon and Bing), ad platforms in display advertising (where there are a handful of competitors to Google) and social media advertising (mostly Facebook). There is also a second way through which large networks can exploit the presence of competing platforms to reduce the cost of search ads, which is bargaining. Within Google's rigid auction system, there seems to be no room for bargaining, but this is a naive view, according to the experts we spoke to. There could be simple tweaks to the auction algorithm that may implement side deals with networks, for instance, by bolstering the quality scores of selected advertisers.[33] There is, however, no guarantee that deals negotiated by the networks will benefit advertisers, as we discuss below in the conclusion.

## VI. Conclusions

The findings we present indicate that concentration among the intermediaries bidding on behalf of advertisers in sponsored search auctions negatively and significantly impacts search engine revenue. Despite the potential benefits for the

[33] See the discussion of Google's "US Media Rebate Program" in the online Appendix.

search engine from the increased efficiency and market expansion that intermediaries bring, the negative revenue result is indicative of the capability of intermediaries to reduce average prices. This is a novel insight into what is currently one of the largest advertising markets and underscores the relevant role of intermediaries. The three key elements of our analysis are first, a novel dataset linking together keywords, advertisers, and intermediaries; second, a new approach to defining markets by aggregating keywords through a two-layer machine learning algorithm incorporating both demand and supply information; and third, the application of an IV strategy based on intermediary mergers.

Several questions are left open for future research and we conclude by briefly exploring two questions, the answers to which are particularly important in interpreting the broader impacts of our findings. The first question is about the internal or external factors that could slow down, or even revert the processes discussed here. Internal factors would involve advertisers choosing to forego the benefits of joint bidding in order to avoid sharing intermediaries (and data) with rivals. But this type of friction does not appear to be salient according to our analysis. Instead, external factors can derive from the actions of either antitrust authorities or the platform. The former are limited to the very specific cases mentioned in the introduction, while the latter could involve a large spectrum of actions initiated by the ad selling platform. Four industry trends might reveal what the selling platforms are doing to reduce their loss of market power: increasing the auction reserve price, reducing the number of ad slots offered, promoting disintermediation services and lastly—as done most notably by Facebook—changing the auction format. Among these four changes, market efficiency is more likely endangered by the first two. In May 2017, Google introduced higher reserve prices differentiated by keyword. In a market dominated by concentrated intermediaries, however, substantial reserve price increases might be required to increase the average CPC. But this would likely hurt the "wrong" advertisers (i.e., those not sharing a common intermediary). Small advertisers placing low bids near the reserve price might find themselves either paying substantially higher prices or being outright excluded from the set of ads that is displayed, thus undermining market efficiency. Over the last few years, Google also started reducing the available ad slots (by eliminating the sidebar and adding a bottom bar with fewer ads). But clearly this approach to increasing competition, by creating slot scarcity, might have the same perverse effect of hurting the "wrong" advertisers mentioned above in relation to the reserve price.[34]

The second question is the extent to which the drop in Google's revenue may be passed on to consumers and, hence, positively contribute to consumer welfare. Since most advertisers operate in markets more competitive than internet search, a transfer of revenue from Google to the advertisers should induce a drop in their costs and, consequently, in consumer prices. If that were the case, increasing buyer power would represent a particularly desirable policy to address the concerns associated with platform concentration. In particular, it might reduce the platform market

[34] Regarding disintermediation—the practice by the selling platform of offering services in direct competition with those of the intermediaries—since it entails a choice by advertisers, we might expect the platform to offer valuable options to induce the advertisers to abandon their MA. But trusting Google to bid on its own auctions, as well as on rival ad selling platforms, might seem problematic to some advertisers. The growth of Google's *smart bidding*, the suite of AI-bidding options, might nevertheless bolster disintermediation.

power without affecting market shares. This is important for search as the market size mirrors the extent of the within-group network effects (Belleflamme and Peitz 2018): the quality of search outcomes depends on the size of the user base. Hence, there is an evident risk with the alternative policies currently debated which involve either helping consumers switch between search engines or improving the quality of smaller search engines through mandated access to Google's data.[35]

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

The final concern worth mentioning regards dynamic implications. Increased buyer power may lead to reduced incentives to innovate by the selling platforms. Moreover, increased buyer power by the merged networks may increase costs for other competing intermediaries, for instance due to a relative worsening of the quality scores of their clients. This would lead to a worsening in choice (or service) for advertisers and, through their higher costs, would harm consumers. Regarding these dynamic considerations, however, more than 30 years after the breakup of the Bell System in 1982, how an economist should look at the long run effects of the loss of power by dominant firms, like Google or the Bell System, is still an open question.

**AQ4**

## REFERENCES

**Adshead, Stephen, Grant Forsyth, Sam Wood, and Laura Wilkinson.** 2019. *Online Advertising in the UK*. London: PLUM Consulting.

**Advanced Web Ranking.** 2017. "Google Organic CTR History." https://www.advancedwebranking.com/ctrstudy/.

**Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, eds.** 2019. *The Economics of Artificial Intelligence: An Agenda.* NBER Conference Report.

[35] For an overview of the policy proposals currently being debated to deal with market power by online platforms see the Stigler Report (2019), Furman Review (2019), Competition Policy for the Digital Era (2019), Competition and Markets Authority Interim Report on Online Platforms and Digital Advertising (2019).

**Angrist, Joshua D., and Guido W. Imbens.** 1995. "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity." *Journal of the American Statistical Association* 90 (430): 431–42.

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

**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 2010*).

**Balseiro, Santiago R., Omar Besbes, and Gabriel Y. Weintraub.** 2015. "Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design." *Management Science* 61 (4): 864–84.

**Balseiro, Santiago R., and Ozan Candogan.** 2017. "Optimal Contracts for Intermediaries in Online Advertising." *Operations Research* 65 (4): 878–96.

**Belleflamme, Paul, and Martin Peitz.** 2018. "Platforms and Network Effects." In *Handbook of Game Theory and Industrial Organization*, Vol. 2, edited by Luis C. Corchón and Marco A. Marini, 286–317. Cheltenham: Edward Elgar.

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

**Blake, Thomas, Chris Nosko, and Steven Tadelis.** 2015. "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment." *Econometrica* 83 (1): 155–74.

**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): 163–87.

**Cao, Xinyu, and T. Tony Ke.** 2019. "Cooperative Search Advertising." *Marketing Science* 38 (1): 44–67.

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

**Council of Economic Advisers (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.** 2014. "Buy-It-Now or Take-a-Chance: Price Discrimination through Randomized Auctions." *Management Science* 60 (12): 2927–48.

**Chipty, Tasneem, and Christopher M. Snyder.** 1999. "The Role of Firm Size in Bilateral Bargaining: A Study of the Cable Television Industry." *Review of Economics and Statistics* 81 (2): 326–40.

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

**Coatney, Kalyn T., and Jesse B. Tack.** 2014. "The Impacts of an Antitrust Investigation: A Case Study in Agriculture." *Review of Industrial Organization* 44 (4): 423–41.

**Craig, Stuart, Matthew Grennan, and Ashley Swanson.** 2018. "Mergers and Marginal Costs: New Evidence on Hospital Buyer Power." NBER Working Paper 24926.

**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.

**Dafny, Leemore, Kate Ho, and Robin S. Lee.** 2019. "The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry." *RAND Journal of Economics* 50 (2): 286–325.

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

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

**Decarolis, Francesco, Maris Goldmanis, and Antonio Penta.** 2020. "Marketing Agencies and Collusive Bidding in Online Ad Auctions." *Management Science*, 66(10): 4359–4919.

**Decarolis, Francesco, and Gabriele Rovigatti.** 2021. "Replication Data for: From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E130502V1.

**Dhillon, Inderjit S., and Dharmendra S. Modha.** 2001. "Concept Decompositions for Large Sparse Text Data Using Clustering." *Machine Learning* 42 (1–2): 143–75.

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

**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): 242–59.

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

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

**Galbraith, John Kenneth.** 1952. *American Capitalism: The Concept of Countervailing Power.* Boston, MA: Houghton Mifflin.

**Gaynor, Martin, and Robert J. Town.** 2012. "Competition in Health Care Markets." *Handbook of Health Economics*, Vol. 2, edited by Mark V. Pauly, Thomas G. McGuire, and Pedro Pita Barros, 499–637. Amsterdam: Elsevier.

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

**Goldfarb, Avi.** 2014. "What Is Different about Online Advertising?" *Review of Industrial Organization* 44 (2): 115–29.

**Gomes, Renato, Nicole Immorlica, and Evangelos Markakis.** 2009. In *Internet and Network Economics*, edited by Stefano Leonardi, 172–83. New York: Springer Nature.

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

**Gutierrez, German, and Thomas Philippon.** 2017. "Declining Competition and Investment in the U.S." NBER Working Paper 23583.

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

**Hastings, Justine S., and Richard J. Gilbert.** 2005. "Market Power, Vertical Integration and the Wholesale Price of Gasoline." *Journal of Industrial Economics* 53 (4): 469–92.

**Hoberg, Gerard, and Gordon Phillips.** 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy* 124 (5): 1423–65.

**Interactive Advertising Bureau (IAB).** 2018. *IAB Internet Advertising Revenue Report: 2017 Full Year Results*. PwC report for the Interactive Advertising Bureau.

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

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

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

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

**Loertscher, Simon, and Leslie M. Marx.** 2019. "Merger Review for Markets with Buyer Power." *Journal of Political Economy* 127 (6): 2967–3017.

**Mansour, Yishay, S. Muthukrishnan, and Noam Nisan.** 2012. "Doubleclick Ad Exchange Auction." arXiv:1204.0535.

**McAfee, R. Preston.** 2011. "The Design of Advertising Exchanges." *Review of Industrial Organization* 39 (3): 169–85.

**Mullainathan, Sendhil, and Jann Spiess.** 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives* 31 (2): 87–106.

**Mullan, Hugh, and Natalie Timan.** 2018. "Strengthening Buyer Power as a Solution to Platform Market Power." *CPI Antitrust Chronicle*, September 20.

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

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

**Pennington, Jeffrey, Richard Socher, and Christopher D. Manning.** 2014. "GloVe: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–43.

**Prat, Andrea, and Tommaso M. Valletti.** 2019. "Attention Oligopoly." Unpublished.

**Redbooks.** 2017. "Redbooks: Agency and Advertising Databases." https://www.winmo.com/redbooks-agency-and-advertising-database.

**Schmitt, Matt.** 2017. "Do Hospital Mergers Reduce Costs?" *Journal of Health Economics* 52: 74–94.

**SEMrush.** 2017. "Search Engine Data: Domain Rankings and Keywords." https://www.semrush.com/kb/998-where-does-semrush-data-come-from#4.

**Silk, Alvin J., and Charles King III.** 2013. "How Concentrated Is the U.S. Advertising and Marketing Services Industry? Myth vs. Reality." *Journal of Current Issues & Research in Advertising* 34 (1): 166–93.

**Simonov, Andrey, and Shawndra Hill.** 2019. "Competitive Advertising on Brand Search: Traffic Stealing and Customer Selection." Unpublished.

**Simonov, Andrey, Chris Nosko, and Justin M. Rao.** 2018. "Competition and Crowd-Out for Brand Keywords in Sponsored Search." *Marketing Science* 37 (2): 200–215.

**Snyder, Christopher M.** 1996. "A Dynamic Theory of Countervailing Power." *RAND Journal of Economics* 27 (4): 747–69.

**Varian, Hal.** 2007. "Position Auctions." *International Journal of Industrial Organization* 25 (6): 1163–78.

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

**Villas-Boas, Sofia Berto.** 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies* 74 (2): 625–52.

**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*, Fall.

**PLEASE ANSWER ALL AUTHOR QUERIES (numbered with "AQ" in the margin of the page). Please disregard all Editor Queries (numbered with "EQ" in the margins). They are reminders for the editorial staff.**

| AQ# | Question | Response |
|---|---|---|
| **1.** | To comply with our publication style your sections have been renumbered. Please check to insure references within the text correspond to the intended section. | |
| **2.** | Please define DMA in Figure 2. | |
| **3.** | Please provide page numbers for directly quoted material in footnotes 22 and 30 | |
| **4.** | In a word doc we can easily edit please provide these missing references, https://www.aeaweb.org/journals/policies/sample-references:<br>Stigler Report 2019;<br>Furman Review 2019;<br>Competition Policy for the Digital Era 2019;<br>Competition Markets Authority Report on Online Platforms and Digital Advertising 2019;<br>US Horizontal Merger Guidelines | |