NET Institute*

www.NETinst.org

Working Paper #17-08

September 2017

Online Auctions and Digital Marketing Agencies

Francesco Decarolis
Boston University

Gabriele Rovigatti
University of Chicago, Booth School of Business

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

Online Auctions and Digital Marketing Agencies*

Francesco Decarolis and Gabriele Rovigatti

October 2, 2017

Abstract

We present an empirical investigation of the role of marketing agencies in Google's online ad auctions. By combining data on advertisers' affiliation to marketing agencies with data on bidding in ad auctions, we analyze how changes in the concentration of clients in the same industry under the same ad network are associated with changes in keyword bidding in terms of entry, exit, and pricing strategies. Moreover, by exploiting the case of a recent merger between agencies, we estimate through a difference-in-differences strategy that an increase in concentration leads to reduction in the average cost-per-click of the keywords affected by the merger.

Keywords: Online Advertising, Internet Auctions, Marketing Agency, Ad Network, Agency Trading Desk.

JEL Classification: C72, D44, L81.

*Francesco Decarolis: Boston University, Department of Economics, and EIEF. E-mail address: *fdc@bu.edu..* Gabriele Rovigatti: University of Chicago Booth School of Business. E-mail address: *gabriele.rovigatti@uchicago.edu.* We are grateful for financial assistance from the NET Institute *www.netinst.org.*

I Introduction

Online advertising is the main source of revenues for important firms such as Google, Facebook, Twitter, etc., and it represents one of the largest and fastest growing industries in the US. Almost all online ads are sold through auctions, in which bidders compete for the adjudication of one of a given number of ‘slots’ available in various online venues, such as search engine result pages, social networks feeds, and so on.

Advertisers in these auctions are increasingly delegating their bidding campaigns to specialized “digital marketing agencies” (DMAs).¹ Many of which belong to a handful of networks (seven in the US) that conduct all bidding activities through centralized agency trading desks (ATDs). As a result, with increasing frequency, the same entity (be it DMA or ATD) bids in the same auction on behalf of different advertisers. But this clearly changes the strategic interaction, as these agencies have the opportunity to lower their payments by coordinating the bids of their clients. The current growth in “programmatic” or algorithmic media buying is making this phenomenon all the more prominent since all by delegating to DMAs or centralized ATDs can advertisers take advantage of the possibilities of effectively allocating their advertisement budgets to different media outlets in real time.

Understanding recent trends in bidding in online advertisement auctions is therefore a crucial element to analyze how revenues will be divided among the various players and how efficiency considerations might be influenced. This paper is an empirical assessment of the effects of bidding delegation to DMAs/ATDs and, more generally, of the effects of the increased concentration of bidding through the same intermediaries. Earlier research has documented that advertisers might dislike sharing marketing agencies with their competitors, likely due to the fear of information linkages (Day, 2015). However, specifically for the context of online auctions, Decarolis et al. (2017) has shown theoretically that advertisers can greatly improve their revenues by delegating to a common agent who bids in the auctions

¹A survey by the Association of National Advertisers (ANA) of 74 large U.S. advertisers indicates that about 77% of the respondents fully outsource their search engine marketing activities (and 16% partially outsource them) to specialized agencies, see ANA (2011). Analogously, a different survey of 325 mid-size advertisers by Econsultancy (EC) reveals that the fraction of companies not performing their paid-search marketing in house increased from 53% to 62% between 2010 and 2011, see EC (2011).

on behalf of them. This study will address three questions related to the quantification of these opposing forces.

The first question that we seek to address is hence whether advertisers o have a preference for joining ad networks that are also representing their competitors. By estimating a model of advertisers’ choice of which ad networks to get affiliated, we find that indeed the evidence is consisted with a preference toward sharing the ad network with competitors.

The second question regards the strategies followed in practice by the ad networks. In principle, these networks could benefit their clients in a number of distinct ways: from choosing the right set of keywords, to optimizing the individual bid to coordinate the auctions with rival advertisers interested in the same keywords to obtain mutually beneficial results. More in details, the theoretical relationship between delegated bidding and auctions cost-per-click is complex. It might involve a multiplicity of actions (choice of keywords, bids, targeting, ad content, etc.) and the same agency might follow different strategies for different advertisers (or for different campaigns by the same advertiser). What seems clear is that the increasing concentration in DMA and ATD mentioned above is bound to change the types of strategies available to agencies. The model in Decarolis et al. (2017) consider the situation in which an agency is jointly optimizing the bidding strategies of advertisers that have already chosen to bid on common keywords. However, a broader set of strategies is certainly possible as the agency might sustain bid rotation or market split schemes, such that the advertisers that it handles are never directly competing in the auctions. Furthermore, beside any collusive intent, a agencies can select different keywords than individual advertisers would.

By exploiting the richness of our data, the second part of the analysis offers a quantitative assessment of the strategies followed. In particular, the empirical analysis in this second part of the paper takes advantage of an acquisition that occurred in July 2016 when one of the largest independent DMA, Merkle, was purchased by one of the seven largest ad networks, Aegis-Dentsu. For a number of Merkle’s clients who were bidding on the same keywords on which also some Aegis-Dentsu’s clients were bidding, this implied the formation of what we refer to as a “common agency.” We explore a number of potential strategies that the

expanded ad network might follow to benefit its clients and find evidence suggestive of the network easing competition between its client through both market splits (i.e., avoiding having its clients simultaneously bidding on the same keywords) and lower bids in the auctions where more than one of the advertisers in the ad auction wins the auction.

Finally, in the third part of the paper, we further develop the analysis concerning the Aegis-Dentsu-Merkle merger to develop a causal identification strategy based on a difference-in-differences approach. The analysis conducted implements the difference-in-differences methodology by defining as “treated” the keywords for which the merger produced a situation of common agency and as “control” the keywords by the same Merkle’s advertisers involved in the merger but not interested by the formation of a common agency (as no Aegis-Dentsu’s clients was bidding on them before or after the merger). The (still preliminary) evidence indicates that the presence of shared DMAs causes the average cost-per-click to decline. Hence, the preliminary results indicate that, in line with the theoretical predictions, there is a major reduction in cost-per-click produced by the merger. Qualitatively the same result is confirmed when using alternative control groups constructed using non-treated keywords selected either from Aegis-Dentsu’s clients or from an algorithm identifying related keywords.

II Literature

This is the first empirical study to focus on the impact of marketing agencies on online ad auctions, and in particular on their role in coordinating the bids of their clients. Earlier theoretical results in Decarolis et al. (2017) have uncovered a striking fragility of the GSP auction to bid coordination so that the empirical tests produced by this project will be of great relevance to understand how online ad auctions are changing in response to changes in bidding behaviors. Our preliminary results confirm the negative impact of common agency on keywords’ cost-per-click. This result is a first order finding since most of the online marketing is still passing through Generalized Second Price (GSP) auctions. Our findings may also provide a rationale for why Facebook has recently adopted the Vickrey-Clarke-Groves (VCG). Shifts between mechanisms are important both for the large stakes involved and be-

cause the proper functioning of this market is essential for both advertisers and consumers.

From a methodological perspective, we note that our dataset and approach will likely be usable for other researchers to explore related questions. This suggests that our approach, which combines large amount of novel data, a deep understanding of the theory behind optimal bidding and an extensive knowledge of the market institutions, may be able to produce new knowledge that will be not only relevant by itself, but also useful to open new avenues of research on online ad sales.

From a broader perspective, our findings complement other recent work the evolution of bidding behavior in online auctions. For instance, Blake, Nosko and Tadelis (2015) use large scale experiments to explore how eBay could benefit from a more nuanced bidding behavior that distinguishes between brand and non-brand keyword ads. Einav, Farronato and Sundaresan (2014) who document a decline in the importance of consumers' bidding in the eBay auctions, with a progressive shift towards purchasing at posted prices. Altogether, it emerges the picture that bidding behavior in online marketing platforms is undergoing important transformations that still need careful analysis.

Finally, as pointed out earlier, our findings are potentially relevant from an antitrust perspective. In particular, in many ways agency behavior in our framework is analogous to that of buying consortia, which have been sanctioned in the past. Nevertheless, the specificities of online ad market suggest a more nuanced view of the harm to the consumers. First, although our main analysis focuses on agencies' role to coordinate their clients' bids, agencies in this market have other roles which are expected to improve the efficiency of the system (e.g., in improving sellers' ability to reach new consumers, improve advertisers' campaigns, bringing new advertisers to the market, etc.). Second, it is likely that the degree of competition between different search engines is substantially less than that between most of advertisers. Since the lower auction prices due to agency bidding imply a reduction in the marginal cost advertisers pay to reach consumers, advertiser competition implies that some savings are passed on to consumers. Therefore, harm to consumers would result only if the agency engages in coordinating not only auction bids, but also the prices charged to consumers. Third, bid coordination can negatively affect the quality of the service received

by consumers by further exacerbating the advantage of dominant search engines relative to fringe ones. In Europe, for instance, where 90% of the searches pass via Google, agencies might be rather careful not to harm Google given the risk of being excluded from its results page. Smaller search engines cannot exert such a threat because agencies are essential to attract new customers. The shift of revenues from small search engines to marketing agencies could thus deprive the former of the essential resources needed for technology investments. Thus, to the extent that competing search engines exert pressure for quality improvements, bid coordination poses a threat to consumer welfare.

All these considerations represent potentially fruitful directions for future research on this important market that our project will likely contribute to open up.

III Industry Background

Online advertisement refers to all activities that use the Internet to deliver marketing messages to consumers. Internet advertising revenues have surpassed those based on traditional media in the US and this form of marketing is now the main channel through which most firms reach customers and customers learn about products. Online advertising has become the main source of revenues for superstar tech firms like Google, Facebook, Twitter, and it represents one of the largest and fastest growing industries in the US: in 2013, for instance, the value of advertising on search engines alone amounted to 50 billion dollars in the U.S., with an annual growth of 10% (PwC (2015)), and 96% of Google's global revenues in 2011 were attributed to advertisement (Blake, Nosko and Tadelis (2015)). Almost all online ads are sold through auctions, in which bidders compete for the adjudication of one of a given number of 'slots' available in various online venues, such as search engine result pages, social networks feeds, and so on. With the significant exception of Facebook, which recently adopted the VCG mechanism, for a long time this market has been dominated by the GSP auction (used, for instance, by Google, Microsoft-Bing, Yahoo! and Taobao).

The expansion of the online ad auctions has generated a tremendous amount of policy interest and academic research. However, the vast majority of the literature has been focused

on advertisers that bid individually.² In contrast, this project intends to examine the role of intermediaries known as Digital Marketing Agencies (DMAs), which are marketing agencies specialized in online markets and bidding on behalf of their clients in the online ad auctions. A major share of bidding and new account openings are currently occurring under DMAs.³

The bidding behavior in online ad auctions is vital to the operation of the market, as the system is centered around competitive bidding to allocate scarce display space to different advertisers. Search engines offer an essential tool for consumers to exploit the Internet and these search engines must make design choices of which selling mechanism to use.

Our main interest is on DMAs bidding in the same auction on behalf of different advertisers. As we explained in an earlier theoretical piece (Decarolis et al. (2017)), this increasingly common phenomenon has the potential to drastically change ad auctions as DMAs have the opportunity to lower their payments by coordinating the bids of their clients. In that study, we documented how bidding delegation to DMAs causes strategic manipulation of the most frequently used types of auctions, the Generalized Second Price Auction (GSP, used by Bing-Yahoo!, Google and Taobao) and the Vickrey-Clark-Groves auction (VCG, used by Facebook, and, recently, also by Google for a subset of its auctions). The analysis involved a theoretical characterization of equilibrium bidding in the auctions that accounted for DMAs' incentives: the agency is in a position to internalize the benefits of coordinated bidding by softening price competition between the advertisers it handles. In particular, our theoretical analysis uncovers a striking fragility of the GSP mechanism to the possibility of bid coordination and this is underscored by our finding that the GSP auction is outperformed by the VCG, a mechanism that is well known to perform poorly under coordinated bidding.⁴

² In two influential papers, Edelman, Ostrovsky and Schwarz (2007) and Varian (2007) pioneered the study of these auctions. For more recent work, see Athey and Nekipelov (2012).

³ Already in 2001, a survey by the Association of National Advertisers (ANA) among 74 large U.S. advertisers indicates that about 77% of the respondents fully outsource their search engine marketing activities (and 16% partially outsource them) to specialized agencies. For the same year, a different survey of 325 mid-size advertisers by Econsultancy reveals that the fraction of companies not performing their paid-search marketing in house increased from 53% to 62% between 2010 and 2011.

⁴ This theoretical work received crucial financial support from a NET Institute grant. Although it has not yet been published, we have disseminated its findings through a series of invited seminars held at a multiplicity of academic and institutional settings. These include invited seminars at Berkeley-Haas Business School, Boston University, Columbia University, CREST-Paris, European Commission DG Competition, Facebook Research Lab, Harvard-MIT joint IO seminar, HEC-Lausanne, London Business School, Microsoft Research Lab, Princeton University, Stanford University, University of British Columbia, University of Cali-

The current project will complement the earlier theoretical analysis by offering an in-depth empirical investigation of how auctions’ prices and allocations are affected by bidding through a common agency. In the early days of online ad auctions, advertisers bid through their own individual accounts, often managed separately across platforms. But already back in 2011 (see footnote 3), a large share of advertisers in the US delegated their bidding activities to specialised digital marketing agencies (DMAs), whose diffusion quickly led to the issue of common agency discussed in the introduction. The case of Merkle, one of the major agencies in the both Europe and the U.S., provides a clear example of this phenomenon. Crucially for our purposes, many of Merkle’s clients operate in the same industries, and are therefore likely to bid on the same keywords.⁵ For instance, data from Redbooks (the leading public database to link advertisers to their agencies) confirm that Merkle managed the campaigns of two leading charities in 2016, *Habitat for Humanity* and *Salvation Army*, both of which were bidding in the same auctions for hundreds of keywords.⁶ Table 1 reports the top six of these keywords, in terms of their average cost-per-click (CPC).

Table 1: Top 6 common keywords between *Salvation Army* and *Habitat for Humanity*

Keyword	CPC	Volume	Position	
			<i>Habitat</i>	<i>Salv.Army</i>
habitat for humanity donations pick up	4.01	40	1	4
charities to donate furniture	1.08	20	3	9
donate online charity	0.93	20	11	10
website for charity donations	0.90	19	11	6
salvation army disaster relief fund	0.03	20	2	1
giving to charities	0.05	30	8	5

Notes: CPC is the average cost-per-click in \$US. Volume is the number of monthly searches, in thousands. Position refers to rank among paid search links on Google’s results page for the relevant keyword. Source: 2016 US Google sponsored search data from SEMrush.

While several hundred DMAs are active in the US, most of them belong to one of the seven main agency networks (Aegis-Dentsu, Publicis Groupe, IPG, Omnicom Group, WPP/Group

fornia Davis, University of Chicago, University of Mannheim, University of Toronto, University of Wisconsin Madison and at the conferences of the EIEF-University of Bologna-Bocconi joint IO workshop, Society for Economic Dynamics and Toulouse Network for Information Technology Meeting.

⁵See: <https://www.merkleinc.com/who-we-are-performance-marketing-agency/our-clients>.

⁶Similar examples can be identified for nearly every industry: for clothing, *Urban Outfitters* and *Eddie Bauer* use Rimm-Kaufman; for pharmaceuticals, *Pfizer* and *Sanofi* use Digitas; etc. (Source: Redbook.)

M, Havas, MDC), which operate through their corresponding ATDs (respectively: Amnet, Vivaki, Cadreon, Accuen, Xaxis, Affiperf and Varick Media). ATDs' importance is growing alongside another trend in this industry, in which DMAs also play a central role.

That is, the ongoing shift towards the so called 'programmatic' or 'algorithmic' real time bidding: the algorithmic purchase of ad space in real time over all biddable platforms through specialized software. ATDs are the units that centralize all bidding activities within a network for 'biddable' media like Google, Bing, Twitter, iAd, and Facebook. Hence, while DMAs were originally not much more sophisticated than individual advertisers, over time they evolved into more and more sophisticated players, and their diffusion and integration through ATDs has made the issue of common agency increasingly frequent. Table 2 lists all the main networks, the associated ATD and the the date of their first creation.

Table 2: Main AD Networks and their Agency Trading Desks

Network	ATD	Year
Dentsu-Aegis	Amnet	2011
	Accordant Media	2016
Publicis Groupe	Vivaki (Audience on Demand, AOD)	2008-2014
	RUN	2014-2014
	Precision	2017
IPG	Cadreon (Mediabrand's Audience Platform)	2009
Omnicom Group	Accuen	2009
WPP/Group M	Xaxis	2011
Havas	Adnetik (spun off as an independent company in 2010) Affiperf	-2010
MDC	Varick Media	2008

Notes: the table reports the networks - column 1 - and their controlled ATD - column 2 - along with their years of activity - column 3.

Our interest focuses on one specific consequence of these phenomena: agencies' ability

Figure 1: For each of Merkle’s advertisers in footnote ??, the figure represents the number of keywords on which it bid alongside at least one member of the Aegis-Dentsu network (and as a share of the total number of keywords on which it bid, in parenthesis) between June 2015 and January 2017 (Merkle’s acquisition by Aegis-Dentsu was in July 2016). The graph shows whether bids on these ‘shared’ keywords occurred only pre-acquisition (dark blue: all keywords appearing only before July 2016), only post-acquisition (turquoise: all keywords appearing only after July 2016), or both pre- and post-acquisition (blue: all keywords appearing both before and after July 2016.) Source: keyword-level data provided by SEMrush.

to lower the payments of their clients by coordinating their bids. But this need not be the only way in which agencies implement collusion. One alternative could be to split the keywords among an agency’s clients, so that they do not compete in the same auctions. This ‘bid retention’ strategy is obviously advantageous in single-unit auctions, but in principle it might be used in multi-unit auctions too. A recent episode, also part of the trend towards concentrated bidding outlined above, may help us illustrate the significance of the potential for bid coordination which our model focuses on.

In July 2016, Aegis-Dentsu acquired Merkle, which was not previously affiliated to any network. At that time, many of Merkle’s clients were bidding on the same keywords as some of Aegis-Dentsu’s advertisers: for instance, in the electronics sector, *Dell* and *Samsung* were in Merkle’s portfolio, placing bids on keywords also targeted by Aegis-Dentsu’s clients

Apple, HP, IBM/Lenovo and Intel. Other examples include: in the financial sector, Merkle’s *Lending Tree* and *Metlife* were bidding in auctions alongside Aegis-Dentsu’s *Capitalone, Discover, Fidelity, Equifax, JP Morgan-Chase*; for car manufacturers, Merkle’s *FIAT-Chrysler* and *Mercedes-Benz USA* bid alongside Aegis-Dentsu’s *Toyota, Volkswagen, Subaru*; in phone services, Merkle’s *Vonage* bid alongside Aegis-Dentsu’s *T-Mobile*.⁷ This acquisition therefore further increased the potential for coordinated bidding. Figure 1 reports, for each of Merkle’s advertisers listed above, the fraction of the total keywords on which they were bidding at the same time as some of Aegis-Dentsu’s clients, and whether joint targeting of such keywords happened only pre-acquisition, only post-acquisition, or both pre- and post-acquisition. Although there is some variation among these advertisers, we clearly see that shared keywords are a quantitatively large phenomenon also post-acquisition (interestingly, a small fraction of keywords are shared *only* post-acquisition). Hence, coordinated bidding through a common agency in the same auction – the focus of our analysis – is clearly relevant.

IV Data

The minimal data requirements to test these hypotheses are information on: *i)* the advertisers’ affiliation to DMAs, *ii)* the set of keywords on which they bid and *iii)* the associated average CPC of these keywords. We have put together for this project a new dataset that contains all this information, and more. We have obtained from Redbooks - the most comprehensive database on marketing agencies - a list of advertisers representing nearly the universe of major US firms active in online marketing. For each of these advertisers, the Redbooks data gives us the full list of marketing agencies (both DMAs and traditional marketing agencies) affiliated with them. We have obtained these data for the years 2011-2016 and they contain on average 6,000 advertisers (i.e., web domains) per year, covering all sectors of the economy. For the years 2016 and 2017 only, we also have access to a variable that relates each individual agency to the ad network to which it belongs. These ad networks will play a particularly salient role in our analysis and several of our results are thus based on

⁷Source: Redbooks data from January 2016.

these two years of data. We have combined these data on agencies with data on the Google’s sponsored search auctions from SEMrush, the most important provider of online advertising data. For all Redbooks advertisers, we know which keywords, if any, they bid on, via Bing or Google. For each keyword and year, we know the position of the domain in the search outcome page, the volume of searches associated with the keyword (i.e., the average number of search queries for the given keyword in the last 12 months); the visible URL of the ad; the content of the ad; and, most importantly, the keyword-specific average CPC (i.e., the price advertisers pay for a user’s click on an ad triggered by the given keyword).

The preliminary aspect that these data allow us to address is whether it is common for advertisers bidding on the same keywords to use the same DMA.⁸ Indeed, not only data indicates that this is the case, but also they show that this phenomenon has increased over time. Combining together the Bing and Google data, the number of keywords where at least two bidders shared the same marketing agency went from 13 thousand in 2011 to 56 thousand in 2016. Table 3 reports summary statistics on these keywords for 2016, separately for Bing and Google.

Table 3: Summary Statistics: Keywords with at Least One Common DMA

	Bing			Google		
	Mean	Median	N	Mean	Median	N
Coalition size	2.00	2	6,367	2.04	2	50,108
Num. coalitions per keyword	1.04	1	6,367	1.05	1	50,108
Ag.Type: DMA	0.15	0	6,367	0.19	0	50,108
Ag.Type: Traditional	0.95	1	6,367	0.77	1	50,108

The coalition size variable measures the number of advertisers that are under the same DMA, for the given keyword. Although the vast majority of cases involves coalitions of size 2, there are a few examples of larger coalitions. For instance, for the keyword “online banking” there are four advertisers (Bank of America, Travelers, Geico and State Farm) all affiliated with a single DMA (The Martin Agency, a major marketing agency that is also part of the Interpublic Group of Companies). The next row in Table 3 reports how many coalitions are present within the same auction (i.e., how many different DMAs represent more than one

⁸All the empirical analysis is still ongoing and the numbers reported below are subject to change.

advertiser bidding on the same keyword). The statistics indicate that there is typically a single multi-advertiser DMA per keyword. This is consistent with the agencies’ specialization and helps explaining why the phenomenon of common DMA appears to be expanding over time. Finally, the last two rows of the table report agency characteristics and, in particular, whether the agency is a DMA or whether it is a traditional marketing agency. While DMAs are agencies specialized in bidding in the paid search auctions, traditional agencies are firms operating on multiple media that might or might not include paid search. Using the names of the agencies, we plan to further refine these categories to achieve a more precise classification of agencies between those active in paid search auctions and those not.

In principle, our dataset allows us to test the effect of having a common DMA on the average keyword cost-per-click by exploiting the variation over time and across keywords in common DMAs. In particular, a difference-in-differences strategy could be used by looking at all those keywords for which the switch from having no more than one advertiser under the same DMA to having at least two advertisers under it occurred during the sample period. These will be our *treated* keywords. Variation in advertisers’ affiliations to agencies guarantees that this strategy is feasible. Moreover, we have available as potential *control* keywords a large number of keywords that are similar to the treated ones. For instance, since for each advertiser we observe all the keywords on which it bids, it is likely that advertisers who are part of the common DMA also bid on non-treated keywords that are nevertheless closely related to the treated ones. This is due to the nature of the ad auctions where, for instance, bidding on “online banking” is not the same as bidding on “online banking account.” Therefore, by exploiting the variation over time and across keywords in the common-DMA phenomenon, we could identify whether having a common DMA reduces the cost-per-click. We could also be able to assess potential heterogeneity of the treatment across advertiser’s industries and agency types. Clearly, however, this type of approach faces a series of challenges and below we explain how we plan to address them. Before doing so, however, we shall list the main project’s objectives and expected outcomes.

V Advertisers' Choice of Ad Network

The first question that we seek to address is hence whether advertisers o have a preference for joining ad networks that are also representing their competitors. We analyze this aspect by estimating a model of advertisers' choice of which ad networks to get affiliated, we find that indeed the evidence is consisted with a preference toward sharing the ad network with competitors. In particular, we estimate a discrete choice model where, for each advertiser, the choice set is represented by 8 options: the 7 ad networks or the outside option of an independent agency.

To estimate this model we use the dataset from Redbooks on the affiliation of advertisers to marketing agencies and of marketing agencies to ad networks. We estimate the model using the data for 2016 and 2017, which are the only two years of data for which we observe the agencies' affiliation to ad networks. Table 4 reports some summary statistics for this sample. Each row reports the data for one of the 7 ad networks, or -the last row -for the outside option. The following table, instead, Table 5, reports the summary statistics at advertiser level.

Table 4: Descriptive Statistics: Networks and Independents

	2016				2017			
	#agencies	Online	#adv	ave#Adv	#agencies	Online	#adv	ave#Adv
Dentsu-Aegis	47	0.21	185	3.94	51	0.20	217	4.25
Publicis	171	0.058	677	3.96	148	0.095	562	3.80
IPG	174	0.11	735	4.22	167	0.14	640	3.83
Omnicom	248	0.081	942	3.80	229	0.092	848	3.70
WPP	293	0.089	850	2.90	291	0.079	858	2.95
Havas	46	0.043	167	3.63	36	0.083	138	3.83
MDC	35	0.23	219	6.26	41	0.24	254	6.20
Independents	2500	0.22	5688	2.28	2113	0.22	4890	2.31

Notes: table reports the number of agencies, the share of SEO agencies, the number of unique advertisers (i.e., clients) and the average number of advertisers *per agency* of the 7 main networks - first 7 rows - and the Independent agencies aggregated (last row). Columns (1)-(4) refer to 2016, columns (5)-(8) report the same values for 2017.

Table 5: Descriptive Statistics - Advertisers

	Mean	SD	Min	Max	N
Num Agencies	2.65	3.72	1.00	25.00	3,577
Num Networks	0.65	1.06	0.00	6.00	3,577
Sales (000000)	10,381.03	25,539.88	0.09	259488.00	818
Num Employees (000)	12.23	38.04	0.00	663.24	1,898
US based (%)	99.86	3.73	0.00	100.00	2,159
Media (%)	4.50	20.74	0.00	100.00	3,577
Financial (%)	5.70	23.19	0.00	100.00	3,577
Food & Beverage (%)	2.15	14.52	0.00	100.00	3,577

Notes:

Figure 2: Number of advertisers per industry - 2016

Table 5 reports the characteristics of the advertisers that are observable in the Redbooks data. We have about 3,500 advertisers representing all the main advertisers active in the US

market. On average, they affiliate with 0.65 networks (underscoring the importance of the independent agencies) and with 2.65 agencies (either part of ad networks or independent). For all these advertisers, we know their main sector of business (media, financial services and food-and-beverages being the three largest). Figure 2 reports the number of advertisers for each one of the 24 industries categorized in the Redbooks data. Finally, for a subset of advertisers, we also know information on revenues and the number of employees.

Figure 3: Network Industry Specialization

The discrete choice model is estimated following Train (2003) via maximum likelihood. Advertiser i choice of which network to join, or whether to take an independent agency, is modeled as a function of the characteristics of each one of these options, the characteristics of the advertiser himself and a logit error. The choice model that we estimate is derived from a linear indirect utility formulation like the one that follows:

$$u_{ijt} = \alpha_i(Industry_{it} * NetworkIndustry_{jt}) + x_{jt}\beta_i + \epsilon_{ijt}$$

The coefficient of interest is α , the effect of the interaction between the advertiser's i

industry and the exposure of ad network j to that industry (measured as the share of clients that the network has in that industry relative to its overall client portfolio). Figure 3 visually describes the evidence by reporting the degree of specialization between the 7 ad networks.

Table 6 reports the result of the choice model. Various specifications are reported, depending on the different covariates for both i and j are included in x . The result of a positive and significant effect is clearly indicative of a preference to join network more exposed to one’s own industry. This is an interesting result as the sign of the estimated coefficient seems a rejection of the conventional “sleeping with the enemy” story. Indeed, advertisers tend to strictly prefer agencies handling accounts in similar markets. We consider that this might be a direct effect of how advertisement over the internet works and we will explore this channel in the next section.

Table 6: Choice Logit

choice					
Industry * Network Industry	3.456*** (0.604)	3.892*** (0.632)	3.899*** (0.633)	3.899*** (0.633)	4.550*** (1.102)
Digital Net 2016			-3.251** (1.229)		
Share SEO				-10.19* (5.124)	-10.25* (5.129)
Share Multimedia				5.268 (9.409)	5.566 (9.431)
Share Paidsearch				-0.546 (5.257)	-0.628 (5.288)
Share Web				-3.728 (4.127)	-3.782 (4.163)
N	57,232	57,232	57,232	57,232	57,232
Network FE	✓	✓			
Industry Net Level		✓	✓	✓	✓
Interactions					✓

Before, we concluding this section, it is useful to stress that the flux of advertisers leaving and joining the agencies is large and continuous. Therefore, in the next section we will exploit the merger between Aegis-Dentsu and Merkle - already discussed in section 3 - to analyze the effects of increased agency concentration on the strategies used in online auctions. In this respect, it is useful to illustrate through Table 7 that networks grow both through the

acquisition of existing agencies (like Merkle) or the formation of the new ones. The former case is nearly ideal to study the effects of changes in the agency structure to auction-level outcomes, like CPC and entry. This is the reason why all the following analysis is based on the Aegis-Dentsu-Merkle merger.

Table 7: Entrant agencies *per network* in 2017

Panel a): Previously Independent						
Dentsu-Aegis	Publicis	IPG	Omnicom	WPP	Havas	MDC
Grip Gyro Happy Creative Services Merkle	North Strategic	BNP Worldwide ReviveHealth StickyEyes		Cavalry Agency Deeplocal Essence Digital Mirum Global Muh-Tay-Zik Hof-fer SHIFT Communications Zubi Advertising Services iStrategyLabs		Forsman & Bodenfors Laird+Partners
Panel b): Brand New						
Dentsu-Aegis	Publicis	IPG	Omnicom	WPP	Havas	MDC
Band Pte Barnes Catmur & Friends C2C Outdoor IMPAQT Perfect Relations		Flipside Group Healix Rapport Worldwide SociedAD Trilia Media	Hearts & Science United State of Fans	Code Computerlove Conrad Caine GmbH Famous nv/sa Quirk m/SIX Tank	Ignition Holdings	

Notes: Previously Independent (panel a) and brand new (panel b) agencies merged and acquired by the 7 networks during 2016.

VI Ad Auction Strategies: Descriptive Evidence

In this section we intend to explore what are the effects produced by a change in concentration on the agency side. Our descriptive analysis takes advantage of Aegis-Dentsu’s acquisition of Merkle in the summer 2016. As discussed in section 3, Aegis-Dentsu and Merkle share many clients active on the same keywords. Therefore this merger is ideal to study the various potential strategies that the merged entity can implement, for those clients that decide to remain with it post merger. For all such advertisers, we collect all the keywords they have been bidding on Google according to the SEMrush data. Summary statistics for this dataset are reported in Table 8.

Our interest in this section is to describe a few basic strategies that an agency like Aegis-

Table 8: Dentsu-Aegis-Merkle: Descriptive Statistics

	Avg # adv	Avg # Keys - DA	Avg # Keys - Merkle	Avg # Key per adv	Tot bids	Unique adv	Tot bids - DA
2015m10	1.17	36590.5	1796347	98	1462432	67959	1530391
2015m11	1.21	36962.0	1892307	97	1497229	68338	1565567
2015m12	1.19	37135.7	1839558	97	1474382	68814	1543196
2015m7	1.24	33241.0	1319261	93	1042726	18991	1061717
2015m8	1.25	34047.9	1369240	94	1068823	27303	1096126
2015m9	1.24	33972.3	1404477	97	1090650	44211	1134861
2016m1	1.17	35980.7	1644409	96	1360718	41087	1401805
2016m10	1.09	28653.8	613932	92	532806	29354	562160
2016m11	1.14	25940.7	448827	92	374235	20140	394375
2016m12	1.10	29175.1	699357	92	579550	55861	635411
2016m2	1.17	35625.2	1536291	94	1272156	44805	1316961
2016m3	1.17	35462.1	1520696	95	1252334	51906	1304240
2016m4	1.17	35299.1	1498427	97	1227925	54876	1282801
2016m5	1.13	33284.6	1259161	97	1061696	54771	1116467
2016m6	1.13	32290.6	1175706	93	989126	55586	1044712
2016m7	1.10	31014.7	1095491	95	944247	47490	991737
2016m8	1.09	30595.8	1035128	91	895644	55682	951326
2016m9	1.10	29076.2	894873	92	771842	43217	815059
2017m1	1.09	32211.7	507478	53	461692	2924	464616

Notes: Descriptive Statistics on AegisDentsu and Merkle keywords and advertisers - 2016.

Dentsu-Merkle can take and to illustrate their quantitative relevance. Therefore, we cluster all the keywords of Table 8 in 6 different classes, according to the strategies followed by the advertisers and of the cost-per-click dynamics after the merge. In particular, these six keyword strategies are: *exiter* (keywords which appear before but not after the merge), *entrant* (keywords appearing only after the merge) and *stayer* (keywords appearing in both periods).

A further distinction is made according to the number of advertisers bidding on the keywords: these are either *shared* - i.e., more than one advertiser bids on it - or *non-shared*. Since our interest lies in the strategies followed by advertisers, hence couples highlighted in table 9 are divided either in [entry/exit, shared/nonshared] or, in the stay case, in keywords whose cpc went up or down after the merge. Table 9 summarizes the number of keyword present in each of the six, mutually exclusive, basic strategies.

Below we display a series of time trend graphs of various quantities, by month, on different keyword classes and sample sizes. In particular, samples are clustered according to the presence of the keyword in our data before (these are the *exit* keywords), after (these are the *entry*), or both (*stay*) in SEM data regarding DentsuAegisMerkle advertisers.⁹

⁹Note that *exit* keyword are present also after July, 2016, and *entry* keywords are present even before the

Table 9: Keyword classification - descriptives

	Shared	NonShared	Cpc - up	Cpc - down
Exit	2,900,158	707,740		
Entry	1,073,845	119,054		
Stay			852,454	302,028
N	3,974,003	826,794	852,454	302,028

Notes: frequency of keywords by classes. Exit and entry keywords are divided in shared and non-shared - columns (1) and (2) - while Stay keywords are divided according to the cpc dynamics - either increasing or decreasing.

Figure 4 highlights the CPC dynamics of *stay* keywords from July, 2015 until January, 2017 for various subsamples - 100, 200, 500 and all top volume keywords. In figure 5 a similar picture is reported for *exit* and *entry* keywords in left and right panels, respectively. Interestingly, the average CPC shows an increasing trend for exit shared keywords in the aftermath of the merge - i.e., when Dentsu-Aegis/Merkle agencies stopped bidding on them - whereas the opposite occurs for *entry* shared keywords: the presence of new bidders led the average CPC down. Finally, figures 6 and 7 show the effects of the merger on the average number of bidders and on the total number of bids for *stay* keywords.

merger date because these keyword are bid by other advertisers and we can thus have information on them.

Figure 4: Average CPC dynamics in *stay* keywords

Notes: Panel a): average cpc in best 100 *stay* keywords; panel b) average cpc in best 200 *stay* keywords; panel c) average cpc in best 500 *stay* keywords; panel d) average cpc in all *stay* keywords.

Figure 5: Entry and Exit keywords dynamics

Notes: Panel a): weighted average cpc in best 500 *exit* keywords; panel b) weighted average cpc in best 500 *entry* keywords; panel c) weighted average cpc in all *exit* keywords; panel d) weighted average cpc in all *entry* keywords.

Figure 6: Stay keywords: Average Number of Bidders

Notes: Panel a): weighted average number of bidders in best 100 *stay* keywords; panel b) weighted average number of bidders in best 200 *stay* keywords; panel c) weighted average number of bidders in best 500 *stay* keywords; panel d) weighted average number of bidders in all *stay* keywords.

Figure 7: Stay keywords: Total Number of Bids

Notes: Panel a): total number of bids in best 100 *stay* keywords; panel b) total number of bids in best 200 *stay* keywords; panel c) total number of bids in best 500 *stay* keywords; panel d) total number of bids in all *stay* keywords.

VII Agency Behavior in Shared Keywords

In this section, we zoom in the case of shared keywords as those are the ones for which we have clean hypotheses to test. The discussion in the previous section spelled out in a generic way the differences-in-differences method that we plan to follow. Mapping that general idea to concrete steps capable of effectively delivering reliable estimates of the effects of common agency requires substantial care. Two features are particularly salient. The first involves when can we exploit the emergence of a case of common agency to study bidding behavior. The second involves the construction of the control group.

Regarding the first issue, not all instances of common agency that arises in our data are appropriate to address the question of interest. For instance, if an advertiser bidding independently were to join an agency already handling the bids of some of this advertiser's competitors, we would be clearly observing the formation of a case of common agency with well defined pre/post periods. However, questions regarding why did such an advertiser join the DMA would naturally arise. The complexities of bidding in the ad auctions have pushed over time the advertisers that were considering their ad campaigns not optimally run to delegate bidding to DMAs more expert in optimizing bids. But then if the source of common agency is a failure to optimally bid of one of the members of the agency in the pre-common agency period, then our difference-in-differences analysis would produce estimates that convolute together bid changes arising from both collusion and failure to optimize. Therefore, the ideal setting for our analysis requires looking at advertisers that are most likely already bidding optimally before the creation of the common agency. This condition is satisfied for advertisers that are already delegating their bids to DMAs. However, even in this case the switch from one DMA to another operated by an advertiser might signal the presence of elements that would limit the validity of our empirical strategy. This would be the case, for instance, if the advertiser was considering the bids of its DMAs too high/low relative to its perceived optimal bids.

Our solution is therefore to exploit once again the acquisition of Merkle by Aegis-Dentsu in July 2016 as an event of formation of a common agency formation that is exogenous to

the choices of any individual advertiser. As discussed above, both agencies are highly sophisticated players and, therefore, the bids submitted for their clients can reasonably be considered as optimal bids. This merger produced instances of common agencies in a broad range of extremely relevant markets. Indeed, by indicating in red the Dentsu-Aegis advertisers at the time of the merger and in blue those of Merkle, we have the following cases of pools of advertisers belonging to the same industries and that, according to the SEMrush data, bid on the same sponsored keywords:

- Electronics: Dell, Samsung → Apple, HP, IBM/Lenovo, Intel (also, but for fewer keywords in common: eBay, HomeDepot, Target, Wal-Mart)
- Financial: LendingTree, MetLife → CapitalOne, Discover, Fidelity, Equifax, JP Morgan-Chase
- Car manufacturers: Fiat Chrysler Automotive, Mercedes-Benz USA → Toyota, Volkswagen, Subaru (also, but for fewer keywords in common: Autotrader, KBB, eBay)
- Phone services: Vonage → T-Mobile

All the cases above can therefore be studied as nearly ideal instances for testing whether the keywords involved in the merger experienced a decline in their cost-per-click and a variation in the ranking of their advertisers.

Second, a crucial element deserving close attention is the construction of a reliable control group of keywords. We shall ensure not only the comparability of the keywords, but also avoid the risk that too closely related keywords are “contaminated” by changes in the cost-per-click of the treated keywords. Furthermore, we shall consider with care all the possible cases of keywords that were shared only for a subset of the months analyzed (either before or after the merger). However, the large number of available keywords, as well as specific market features like the one described above, make us confident that this task is feasible. If, for instance, we focus on the case of Dell, we see in Figure 8 that the number of keywords shared (and not shared) with Dentsu-Aegis’ advertisers is large both before and after the merger. In this figure, we consider a keyword as shared if there is common bidding for at least one month, but clearly this definition could be made more or less stringent and it will be interesting to explore to what extent the results will be affected by this margin of choice.

Figure 8: The Case of Dell: Keywords Split

M refers to Dell, while DA refers to the set of all relevant Aegis-Dentsu advertisers. The bar chart illustrates for all the keywords on which Dell bid before/after the merger all their situations (split or not split) both before and after the merger.

Continuing with the Dell example, the next margin of choice is whether to use all the shared keywords (several thousands in this specific case), or to focus on a more restricted set. Given the risk of excessive noise by selecting marginal keywords, we see as ideal to perform our analysis on a more restricted set of keywords. More specifically, we plan to focus on four different subgroups: the set of top 30, 50, 100, 500 keywords by traffic volume. The set of top 10 keywords and their relevant summary statistics are reported in Table 10. From this table we can immediately see that for several keywords their CPC goes down in the six months after the merger, relative to the six months before the merger. However, for several other keywords the movement of the CPC is in the opposite direction. This underscores the

need for a control group and a difference-in-differences strategy.

Table 10: The Case of Dell: Keyword Example

keyword	N. Obs.	CPC	Pre Merger		CPC	Post Merger	
			Position	N. Bidders		Position	N. Bidders
build your computer	17	1.87 (0.17)	3.80 (0.84)	2.17 (0.75)	1.23 (0.31)	1.00 (0.70)	2.00 (0.12)
cloud computing	16	71.90 (30.75)	3.10 (2.08)	2.10 (0.74)	35.64 (0.00)	2.00 (1.00)	1.67 (0.82)
computer deals	17	1.73 (0.11)	2.63 (1.60)	2.33 (0.89)	1.71 (0.00)	1.75 (0.50)	2.20 (1.10)
dell 2 in 1	17	0.87 (0.28)	1.09 (0.30)	2.17 (0.83)	1.04 (0.00)	1.50 (1.00)	2.00 (0.00)
desktop computer	16	1.35 (0.24)	2.25 (1.36)	3.67 (1.30)	1.63 (0.00)	1.75 (0.96)	2.25 (1.26)
desktop computers	18	2.52 (0.47)	2.08 (1.00)	3.67 (1.30)	1.99 (0.00)	2.00 (1.41)	2.33 (0.52)
laptops	17	3.65 (0.99)	3.20 (2.20)	4.25 (1.06)	5.26 (0.00)	3.00 (1.73)	2.00 (0.71)
laptops on sale	15	1.93 (0.41)	4.56 (1.81)	4.00 (1.28)	2.59 (0.00)	5.00 (2.83)	4.00 (0.00)
small laptop	16	1.86 (0.84)	4.22 (2.11)	4.45 (1.21)	3.55 (0.00)	3.00 (1.41)	2.20 (1.10)
windows laptops	13	2.60 (0.43)	2.27 (1.35)	3.33 (1.15)	3.01 (0.00)	3.00 (0.00)	3.00 (0.00)

The most basic version of the model that we will estimate for our difference-in-differences analysis is as follows:

$$CPC_{kt} = a_k + b_t + \beta(PostMerger_t * SharedKeyword_k) + \varepsilon_{kt}, \quad (1)$$

where k is a keyword, t is a moth, CPC is the average cost-per-click of the keyword in the month, a is a keyword fixed effect, b is a month fixed effect, $PostMerger$ is a dummy equal to one for the months after July 2016 and zero before then, and $SharedKeyword$ is a dummy equal to one for the Dell's keywords that are shared both before and after the merger with at least one Dentsu-Aegis' advertiser and zero for the Dell's keywords that are never shared (neither before nor after) with any of the Dentsu-Aegis' advertisers.

Table 11: Diff-in-diff Estimates (Still Preliminary)

	30key	50key	100key	500key
Dell	-2.84*** (0.16)	-1.82*** (0.16)	-1.33*** (0.08)	-0.22 (0.59)
Samsung	-0.04 (0.88)	0.22 (0.52)	0.14 (0.67)	-0.10 (0.38)
LendingTree	-0.75** (0.466)	-0.82 (0.31)	-0.25*** (0.62)	-0.36*** (0.07)
MetLife	-1.57*** (0.22)	-1.27*** (0.38)	-0.36 (0.71)	0.74* (0.39)
FCA	-2.05*** (0.02)	-1.28*** (0.07)	-0.99*** (0.06)	-0.54*** (0.14)
MBauto	0.22 (0.45)	0.08 (0.73)	0.47*** (0.00)	0.65*** (0.00)
Vonage	3.37*** (0.17)	3.13*** (0.22)	2.97*** (0.17)	2.03*** (0.11)
Observations	120	200	400	2,000

We perform this analysis separately for each one of the Merkle advertisers involved in the merger. The **still preliminary** estimates that we obtain are reported in Table 11. While several robustness checks have still to be performed, the above estimates offer a clear illustration of both the feasibility of our approach and how we intend to accomplish. In this table, we perform the analysis separately for top 30, 50, 100, 500 keywords by traffic volume. What we find is that for most of the advertisers interested in the merger, the effect on the average CPC is negative as predicted by the theory (these are the advertisers marked in green in the table); for two of them the effect is not statistically significant (advertisers in blue); and for one of them it goes in the opposite direction (advertiser in red). Interestingly, this latter case involving keywords shared by Vonage and T-mobile (example: “business landlines”), is characterized by the fact that the shared keywords are the main keywords in terms of traffic for Vonage, but of very low (relative) volume for T-mobile. This might indicate a potential channel for the failure of our theoretical characterization to hold in this case and suggests interesting directions for extending our analysis.

In future drafts of this study, we plan to sharpen the analysis of the CPC and then to

expand the same methodological approach to look at additional outcomes, like the ranking of the advertisers involved in the merger as for them our theoretical analysis in Decarolis et al. (2017) delivers sharp predictions on when their rank should be preserved by the common agency or not. Furthermore, we will also descriptively analyze a broader set of effects of DMAs on other features of sponsored search auctions. We will look at the effect on advertisers' positions, as this observable variable is a useful proxy to understand whether DMAs might influence allocative efficiency, as predicted by one of our collusive models. Furthermore, in order to assess the prominence of the type of bid coordination, we plan to expand our assessment beyond the effects taking place within a given keyword and to look for strategies involving multiple keywords. In particular, we will assess whether the number and type of keywords an advertiser bids on is affected by the DMA and, if so, whether a DMA handling competing customers might induce a collusive outcome by sharing the set of keywords on which its clients bid. We will also look at the potential benefits induced by the DMAs in terms of click volumes and try to quantify both the sources of these changes (for instance, via changes in ads content).

VIII References

1. Aghion, P., P. Antràs, E. Helpman, (2007). “Negotiating free trade,” *Journal of International Economics*, 73, 1–30.
2. Association of National Advertisers (ANA), (2011). *Trends in Digital Agency Compensation*, 4th Edition, ANA publishing.
3. Athey, S., D. Nekipelov, (2012). “A Structural Model of Sponsored Search Advertising Auctions,” mimeo.
4. Bachrach, Y., (2010), “Honor Among Thieves - Collusion in Multi-Unit Auctions,” *Proc. of 9th int. Conf. on Autonomous Agents and Multiagent Systems* van der Hoek, Kaminka, Luck and Sen (eds.), Toronto Canada.

5. Balseiro, S. r., O. Besbes, G. Weintraub, (2015). “Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design,” *Management Science*, *forthcoming*.
6. Bernheim, B. D., M. D. Whinston, (1985). “Common Marketing Agency as a Device for Facilitating Collusion,” *The RAND Journal of Economics*, 16(2), 269–281.
7. Bernheim, B. D., M. D. Whinston, (1986). “Common Agency,” *Econometrica*, 54(4), 923–942.
8. Blake, T., C. Nosko, S. Tadelis, (2015). “Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment,” *Econometrica*, 83(1):155-174.
9. Borgers, T., I. Cox, M. Pesendorfer, V. Petricek, (2013). “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” *American Economic Journal: Microeconomics*, 5, 163-187
10. Che, Y.K., Choi, S. and J. Kim, (2013). “An Experimental Study of Sponsored-Search Auctions,” *mimeo*.
11. Che, Y.K., Condorelli, D. and J. Kim, (2016). “Weak Cartels and Optimal Auctions,” *mimeo*.
12. Che, Y.K., and J. Kim, (2006). ““Robustly Collusion-Proof Implementation,” *Econometrica*, 74, 1063–1107.
13. Che, Y.K., and J. Kim, (2009). “Optimal Collusion-Proof Auctions,” *Journal of Economic Theory*, 144, 565–603.
14. Choi, H. and H. Varian, (2012). “Predicting the Present with Google Trends,” *Economic Record*, 88 (s1), 2-9.
15. Celis, L. E., G. Lewis. M. M. Mobius, H. Nazerzadeh, (2015). “Buy-it-now or Take-a-chance: A New Pricing Mechanism for Online Advertising,” *Management Science*, *forthcoming*.
16. Decarolis, F., M. Goldmanis, A. Penta, (2017). “Marketing Agencies and Collusive Bidding in Online Ad Auctions,” *mimeo*.

17. Econconsultancy (EC), (2011). "State of Search Marketing Report 2011," accessed on January 2012 at: <http://econsultancy.com/us/reports/sempo-state-of-search>
18. Edelman, B., M. Ostrovsky, M. Schwarz, (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242-259.
19. Edelman, B., M. Ostrovsky, (2007). "Strategic Bidder Behavior in Sponsored Search Auctions," *Journal of Decision Support Systems*, v.43(1), pp.192-198.
20. Edelman, B., M. Schwarz, (2010). "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100(2), 597–602.
21. Einav, L., C. Farronato, J. Levin, N. Sundaresan, (2014). "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?," *mimeo*.
22. Gomes, R., K. Sweeney, (2014). "Bayes-Nash Equilibria of the Generalized Second-Price Auction," *Games and Economic Behavior*, 86, 421-437
23. Graham, D. and R. Marshall, (1987). "Collusive Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, 95, 1217-1239.
24. Harrington, J. E., A. Sckypacz, (2007). "Collusion under Monitoring of Sales," *RAND Journal of Economics*, 38(2), pp. 314-331.
25. Harrington, J. E., A. Sckypacz, (2011). "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101(6), 2425-49.
26. Hendricks, K., R. Porter and G. Tan, (2008). "Bidding Rings and the Winner's Curse," *RAND Journal of Economics*, 39, 1018-1041.
27. Levin, J., A. Sckypacz, (2016). "Properties of the Combinatorial Clock Auction" (2016) with Jonathan Levin," *American Economic Review* 106(9), pp. 2528-2551.
28. Mailath, G., P. Zemski, (1991). "Collusion in Second Price Auctions with Heterogenous Bidders," *Games and Economic Behavior*, 3, 467-420.

29. McAfee, P. and J. MacMillan, (1992). "Bidding Rings," *American Economic Review*, 82, 579-599.
30. Milgrom, P., J. Mollner, (2014). "Equilibrium Selection in Multi-Player Games with Auction Applications," mimeo.
31. Ottaviani, M., (2003). "Overture and Google: Internet Pay-Per-Click (PPC) Advertising Auctions," London Business School - Case Study-03-022
32. Pricewaterhouse Coopers, (2015). "Global Entertainment Media Outlook," accessed on April 2015 at: <http://www.pwc.com/gx/en/global-entertainment-media-outlook/segment-insights/internet-advertising.jhtml>
33. Ray, D., R. Vohra, (1997). "Equilibrium Binding Agreements," *Journal of Economic Theory*, 73, 30–78.
34. Ray, D., R. Vohra, (2014). "The Farsighted Stable Set," *Econometrica*, forthcoming.
35. Varian, H., (2007). "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163–1178.
36. Wired, (2015) <https://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/>.
37. Wu, T. (2016). *The attention merchants: the epic scramble to get inside our heads*," Knopf, New York.