

November 2021

“Bid Coordination in Sponsored Search Auctions: Detection
Methodology and Empirical Analysis”

Francesco Decarolis, Maris Goldmanis, Antonio Penta
and Ksenia Shakhgildyan

Bid Coordination in Sponsored Search Auctions: Detection Methodology and Empirical Analysis *

Francesco Decarolis, Maris Goldmanis, Antonio Penta and Ksenia Shakhgildyan[†]

November 25, 2021

Abstract

Bid delegation to specialized intermediaries is common in the auction systems used to sell internet advertising. When the same intermediary concentrates the demand for ad space from competing advertisers, its incentive to coordinate client bids might alter the functioning of the auctions. This study develops a methodology to detect bid coordination, and presents a strategy to estimate a bound on the search engine revenue losses imposed by coordination relative to a counterfactual benchmark of competitive bidding. Using proprietary data from auctions held on a major search engine, coordination is detected in 55 percent of the cases of delegated bidding that we observed, and the associated upper bound on the search engine's revenue loss ranges between 5.3 and 10.4 percent.

JEL: C72, D44, L81.

Keywords: Online Advertising, Sponsored Search Auctions, Delegation, Common Agency.

*We are grateful for the comments received from Susan Athey, Jean-Pierre Benoit, Yeon-Koo Che, Kerry Gabrielson, Ken Hendricks, Jon Levin, Massimo Motta, Marco Ottaviani, Rob Porter, Marc Rysman, Andy Skrzypacz and Steve Tadelis and from seminar participants at Berkeley-Haas Business School, Boston University, Columbia University, CREST-Paris, European Commission DG Competition, Facebook Research Lab, Harvard-MIT joint IO seminar, London Business School, Microsoft Research Lab, Princeton University, Stanford University, University of British Columbia, University of California Davis, University of Chicago, University of Bologna, University of Toronto, University of Wisconsin Madison. The BSE acknowledges the financial support of the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). Francesco Decarolis acknowledges financial support from the PRIN 2017. Antonio Penta acknowledges financial support from the ERC Starting Grant #759424.

[†]Decarolis, Bocconi University and IGIER, francesco.decarolis@unibocconi.it. Goldmanis, Department of Economics, Royal Holloway, University of London, Maris.Goldmanis@rhul.ac.uk. Penta, ICREA-UPF, Barcelona School of Economics, and TSE, antonio.penta@upf.edu. Shakhgildyan, Bocconi University and IGIER, ksenia.shakhgildyan@unibocconi.it.

1 Introduction

Sponsored search auctions are those auction mechanisms used to allocate advertisement space on the results web page of search engines like Google, Microsoft Bing and Yahoo!. They represent one of the fastest-growing and most economically relevant forms of internet advertising, accounting for approximately half of the total revenues of this market, which in the United States alone totalled \$107.5 billion in 2018 (IAB, 2019). In recent years, advertisers have switched from individually managing their bidding campaigns to delegating them to specialized agencies known as Search Engine Marketing Agencies (SEMA).¹ Moreover, many of these SEMAs belong to a handful of agency networks (seven in the US) that conduct all bidding activities through centralized agency trading desks (ATDs). As a result, the same entity (be it a SEMA or ATD) often bids in the same auction on behalf of different advertisers.

The ultimate impact of this ongoing trend is difficult to predict. On the one hand, SEMAs and ATDs can help the functioning of this market by both fostering advertisers' participation and improving the quality of the ads consumers receive. But, on the other hand, the agencies' possibility to lower the payments of their clients by coordinating their bids changes the strategic interaction in these auctions, and hence their functioning. In this paper, we focus on bid coordination and abstract from the possible effect on advertising quality. Decarolis, Goldmanis and Penta (2020) (DGP hereafter) provides a theoretical analysis of price bid coordination, showing that the Generalized Second Price (GSP) auction – the most common auction format for this kind of auctions – is particularly vulnerable to this type of bidding coordination, even when agencies only control a small number of advertisers. This is due to the fact that agency bidding in the GSP auction may have both a *direct* and an *indirect effect* on revenues: the first is due to the lower payments associated with the lower bids of the agency bidders; the second is due to the equilibrium effects that manipulating the agency's bids may have on the bidding strategies of the independents, which – as it will be discussed below – typically operate side-by-side with agencies in this market.²

A question of obvious interest is to quantify the extent to which bid coordination can be a relevant channel through which SEMAs can hurt the search engine revenues. This is a crucial question since this revenue loss might negatively impact investments, thus lowering the service quality and, through it, consumers' welfare. Under the equilibrium structure of the GSP auction, even the relatively small coalitions (i.e., advertisers bidding through the same intermediary) observed in the data – the modal coalition size is 2 – might trigger large revenue losses, depending on intricate

¹As shown by the SEMrush data described later, 80 percent of the auctions held on Google in the US for popular keywords involve at least 1 bid submitted through an intermediary.

²Competitive bidding in the GSP was first studied by Varian (2007) and Edelman, Ostrovsky and Schwarz (2007) (EOS hereafter) in a complete information setting, and then by Borgers et al. (2013); an incomplete information model is studied by Gomes and Sweeney (2014). The role of marketing agencies in online ad auctions was first studied by DGP, who analyzed both the GSP and the VCG auction formats, maintaining Varian (2007)'s complete information assumption and allowing agencies to control arbitrary subsets of bidders.

features such as the position of the coalition advertisers in the ad ranking, the value that different advertisers assign to the ad slots up for sale, and more nuanced features.

In this study, we propose an easy-to-implement method to determine whether bid coordination is present in the data and if so, to quantify its revenue effects.³ This short paper thus supplements the theoretical analysis in DGP, by showing how that theoretical model can be applied to search auctions data. In particular, DGP characterize different types of coordination strategies, depending on the extent to which the coalition is willing to trade off collusive profits with higher chances of being identified as colluding by a monitor. The methodology we propose involves two steps: First, by exploiting repeated observations (i.e., auctions) for the same keyword, it determines which behavioral model between coordination and competition best fits the data; Then, the second step of the procedure uses bids, as well as DGP’s theoretical results on the bidding strategies identified in the first step, to back out the underlying bidders’ valuations. Under coordinated bidding, the true underlying valuations of coalition bidders are not point-identified from the data, only their bounds are. We thus use the upper bound, together with the point-identified values of non-coalition bidders, to quantify counterfactual revenues under competitive bidding. This way, we obtain an upper bound on the effect of coordinated bidding on the search engine’s revenues. We can also separately quantify the *direct* and *indirect effects* of coordination mentioned earlier.

We illustrate an application of the method to a proprietary dataset of search auctions held on a major search engine. The dataset consists of a large set of auctions for 71 different keywords: the search engine selected for us these auctions that involve popularly searched keywords, all having exactly 2 bidders acting under a common intermediary (i.e., 2-bidder coalitions).⁴ The application of the two-step method reveals that: (i) coordinated bidding is detected in 55% of the keywords analyzed, with most of the cases being classified as a relatively mild form of coordinated bidding. DGP argues that this form of coordination is indistinguishable from competitive bidding in a single auction, but our novel methodology is able to identify it by exploiting multiple observations from auctions on the same keyword; (ii) Despite its relative mildness, the effect of this form of coordinated bidding is not negligible, as the associated revenue losses may be as high as ranging between 5.3 and 10.4 percent of the revenues in the competitive benchmark; (iii) The findings also indicate that a large fraction of the revenue loss, about three quarters, is due to the *indirect effect* – the adjustment of the bids placed by independent advertisers in reaction to the reduced

³Under the antitrust laws in the US or the EU, the intermediary strategies that we describe below are proper from a legal perspective. They are not comparable to bidder collusion because intermediaries are legal entities, independent from advertisers, operating unilaterally to maximize their profits and free to arrange bidding strategies on behalf of their customers. Exceptions to this general rule might involve cases of “hub and spoke” cartels, where advertisers hire a common intermediary with the explicit intent of coordinating their bids, or cases where it might be applied the discipline on Purchasing Agreements (Group Purchasing Organizations, GPO).

⁴The choice of that specific sample selection criterion was motivated by the fact that the median coalition size in the data is two, and it is rare to observe in the data more than one coalition. In general, it is not necessary to restrict attention to the two-bidders coalitions only. Indeed, the various models of coordinated bidding in DGP differ from the competitive benchmark in that the bids of all agency bidders, with the exception of the highest coalition member, are “too low.” Given that, the method can be applied to coalitions with any number of bidders.

competition among agency clients – which DGP’s theoretical results highlighted as the main source of the potential fragility of the GSP auction, vis-a-vis the strategic opportunities of the agencies.

2 Data

Our dataset is based on the internal records of one of the largest search engines.⁵ This company keeps track of all the search auctions taking place on its dedicated platform. Thus, every time a user queries the search engine for a keyword on which at least one advertiser had placed a bid, the system creates a record which reports: the keyword, the bids of all advertisers winning a position as well as their identity, ad, quality score, rank, clicks received, and, if present, the identity of their agency placing the bid.

The sample is based on a set of historical data from years 2010-2011, constructed as a representative sample of the search auctions involving some of the most frequently searched keywords. Within these “historical data,” a subset was selected by identifying those keywords for which *no more than one* agency was active in the auctions and this agency represented *exactly two* advertisers. This resulted in 71 keywords being selected. Then, the analysis sample was created by collecting all the search auctions involving these 71 keywords that were held during a randomly selected set of 12 days within a three-month time window around the end of 2010 and the beginning of 2011.⁶ These keywords are from different industries and involve different sets of advertisers and intermediaries. Although they obviously cannot span the vast and diverse market of search advertising, they are a useful dataset to illustrate our method.

Working with search engine data is a rare opportunity, but necessarily comes with limitations to the reporting freedom of researchers imposed by confidentiality agreements.⁷ Hence, to help assessing external validity of the proprietary data, about which we are not allowed to disclose further information, we describe some stylized features of the market through publicly available data on Google sponsored search. Thanks to the availability of these new data, this study presents a rare opportunity to analyze a phenomenon – bid coordination via intermediaries – that so far has only been possible to study in the theoretical literature. Specifically, we combine two datasets offering a snapshot of the Google search ads in the US market as of January 2017. The first dataset is Redbook, which links advertisers to intermediaries; the second is SEMrush, which links advertisers to search auctions. In particular, the Redbook data allow linking approximately 6,000 among the largest US advertisers to their marketing agencies and these agencies to their agency

⁵The company name cannot be disclosed, due to a confidentiality agreement.

⁶Specifically, one day was selected at random in each of the 12 sample weeks.

⁷In fact, despite the stunning economic importance of the sponsored search auction, the confidential nature of the data has hindered their empirical analysis. Important exceptions are those in Varian (2007), Ghose and Yang (2009), Athey and Nekipelov (2014), Borgers et al. (2013), Lewis and Rao (2015), Goldman and Rao (2015), and Hsieh, Shum and Yang (2018), as well as those based on the Microsoft’s Beyond Search initiative, Gomes, Immorlica and Markakis (2009), Jeziorski and Segal (2015), and Jeziorski and Moorthy (2018). None of which, however, considers the case of intermediaries.

networks. We thus consider two advertisers as bidding under a common intermediary when they are linked either to the same SEMA or to different agencies but belonging to the same ATD. For all the advertisers in Redbook, we use SEMrush data to create a link with search ad: we combine the list of keywords on which the 6,000 Redbook advertisers appear among the Google search ads, with that containing the 10,000 most frequently searched keywords of 2017 in the US (from SEMrush).⁸ Table 1 presents summary statistics for the resulting sample. Due to the confidentiality agreement we are bound to, we can neither confirm nor deny that the statistics in Table 1 resemble the proprietary sample. Nevertheless, we shall stress the similarities in how the two samples were constructed, that is by looking at very frequently searched keywords on major search engines.

In Table 1, we separate outcomes for the subsample of 1,102 keywords with at least one ad placed by an intermediary (panel B) from those for the subsample of 1,011 keywords with no intermediary bidding (panel D). Panel A corresponds to the full sample. Intermediated bidding is clearly very common and it involves keywords that, in terms of the median outcomes, are close to those in the full sample. For both groups, the median *cost-per-click* (CPC) is about 80 cents, but the mean exceeds \$1.5. Next, *search volume* indicates the monthly number (in millions) of search queries for the given keyword, averaged over the last 12 months. Similarly to the case of the CPC, the average values far exceed the median ones, whereas the median values are similar across the groups. Keywords tend to have substantial variability in their composition in terms of number of words, characters, and whether they are “long tail” (i.e., involving at least 4 words) or not. Finally and most crucially, within the subsample of keywords with delegated bidding, Table 1 reveals that:

- (i) Delegation to a *shared* intermediary is widespread.
- (ii) A coalition of at least two bidders is present in 41% of the keywords.
- (iii) When a coalition is present, its median size is 2.
- (iv) There is never a case of competing coalitions: that is, there is not an auction with two (or more) agencies representing at least two bidders each.

The evidence from the public dataset thus clarifies how certain features of the proprietary dataset are not the result of an arbitrary selection, but typical elements of the market. In particular, we refer to the use in the proprietary data of keywords with 2-bidder coalitions only and with no instances of competing coalitions. In panel C of Table 1, we separate outcomes for the subsample of keywords with the 2-bidder coalitions. It shows that the keywords with 2-bidder coalitions represent well the full subsample of keywords with some intermediary bidding. In the following,

⁸Further details on the data are presented in Decarolis and Rovigatti (2021). Our study combines their dataset with the list of the top 10,000 keywords on Google US (in terms of the number of searches) in 2017. We restrict the attention to these popular keywords to enhance the comparability with the proprietary data.

Table 1: Summary Statistics: *Google Search Auctions - US, 2017*

| | A | | | B | | | C | | | D | | |
|----------------------|-------|--------|-------|-------|--------|------|-------|--------|------|-------|--------|-------|
| | Mean | Median | SD | Mean | Median | SD | Mean | Median | SD | Mean | Median | SD |
| Cost-per-click (CPC) | 1.58 | 0.74 | 3.23 | 1.53 | 0.81 | 2.66 | 1.55 | 0.80 | 1.94 | 1.63 | 0.68 | 3.76 |
| Search Volume | 0.99 | 0.25 | 10.17 | 0.65 | 0.25 | 2.50 | 0.64 | 0.25 | 1.52 | 1.35 | 0.25 | 14.47 |
| # of Words | 1.86 | 2.00 | 0.83 | 1.85 | 2.00 | 0.82 | 1.81 | 2.00 | 0.79 | 1.87 | 2.00 | 0.83 |
| # of Characters | 11.06 | 11.00 | 5.02 | 10.86 | 10.00 | 4.96 | 10.48 | 10.00 | 4.42 | 11.28 | 11.00 | 5.08 |
| Long Tail | 0.03 | 0.00 | 0.18 | 0.03 | 0.00 | 0.18 | 0.04 | 0.00 | 0.19 | 0.04 | 0.00 | 0.19 |
| Coalition | 0.21 | 0.00 | 0.41 | 0.41 | 0.00 | 0.49 | 1.00 | 1.00 | 0.00 | 1.00 | 1.00 | 2.48 |
| Coalition Size | 2.79 | 2.00 | 1.17 | 2.79 | 2.00 | 1.17 | 2.00 | 2.00 | 0.00 | 2.00 | 2.00 | 2.48 |

Notes: statistics at the keyword level. The last four columns are for the full sample, while the first four are for the subset of keywords with at least one ad coming from an intermediary. *Cost-per-click* is in USD; *Search Volume* is the (average) monthly number of searches (in millions); the next three variables measure features of the keywords' length; *Long Tail* is an indicator variable for keywords composed by at least 4 words; *Coalition* is an indicator for the presence among the keyword ads of multiple advertisers affiliated with the same intermediary; *Coalition size* is the number of advertisers under the coalition, calculated exclusively for those keywords with coalitions.

we will focus exclusively on the proprietary data, since they contain the essential information (namely, individual bids and quality scores) needed to apply our methodology.

3 Theoretical Background

Online ad auctions are mechanisms to assign agents $i \in I = \{1, \dots, n\}$ to slots $s = 1, \dots, S$, $n \geq S$ where for simplicity we assume $n = S + 1$ (the extension to $n \geq S$ is straightforward). In our case, agents are advertisers, and slots are positions for ads on a webpage (e.g., on a social media’s newsfeed for a certain set of cookies, on a search-engine result page for a given keyword, etc.). Slot $s = 1$ corresponds to the highest (i.e., best) position, $s = 2$ to the second-highest, and so on until $s = S$, which is the slot in the lowest (i.e., worst) position. Following Varian (2007), the ‘click-through-rate’ (CTR) of slot s – that is, the number of clicks that an ad in position s is expected to receive – is equal to the product of a ‘quality effect’, $e_i \in \mathbb{R}_+$, associated with the advertiser who obtains the slot, and a ‘position effect’, x^s : if bidder i gets slot s , then the expected number of clicks is $e_i x^s$. We assume that $x^1 > x^2 > \dots > x^S > 0$, and let $x^t = 0$ for all $t > S$. Finally, we let v_i denote the per-click-valuation of advertiser i .

In the GSP auction of the search-engine that we analyze, advertisers submit bids $b_i \in \mathbb{R}_+$, which are then adjusted by the *quality scores*. The search-engine’s rationale for using quality scores is to favor advertisers with idiosyncratically higher CTRs. We thus follow Varian (2007) in assuming that they coincide with advertisers’ quality effects, $(e_i)_{i \in I}$. Hence, slots are assigned according to the ranking of the adjusted bids, $\bar{b}_i = e_i b_i$: the first slot to the bidder who submitted the highest adjusted bid, the second slot to the second-highest adjusted bidder, and so on.⁹ A bidder who obtains the s -th highest slot pays a price-per-click equal to the minimum bid he would need to pay to retain the s -th position. We denote bid profiles by $b = (b_i)_{i \in I}$ and $b_{-i} = (b_j)_{j \neq i}$; vectors of adjusted bids (for a given $e = (e_i)_{i \in I}$ profile) are denoted by $\bar{b} = (e_i b_i)_{i \in I}$ and $\bar{b}_{-i} = (e_j b_j)_{j \neq i}$. Finally, we relabel bidders, if necessary, according to the slot they occupy: hence, given a profile of bids $b = (b_i, b_{-i})$, b_i denotes the bid placed by the advertiser in position i , that is the one who placed the i -th highest adjusted bid $\bar{b}_i = b_i e_i$. With this notation, the payoffs which result from bid profile b , given a vector of quality scores e , can be written as $u_i(b; e) = \left(v_i - \frac{e_{i+1}}{e_i} b_{i+1} \right) e_i x^i$.

Competitive Bidding: Varian (2007) and Edelman, Ostrovsky and Schwarz (2007) identified a specific refinement in this auction, the lowest-revenue locally envy-free equilibrium, which in this

⁹As in Varian (2007) and EOS, we maintain that quality scores, valuations and CTRs are common knowledge (EOS actually abstracted from quality scores). This complete information environment is the main benchmark for the literature on the GSP auction. A notable exception is Gomes and Sweeney (2014). Borgers et al. (2013) maintain the complete information assumption, but consider a more general model of CTRs and valuations. Athey and Nekipelov (2014) introduce uncertainty over quality scores in a model with competitive bids.

setting is characterized by the following recursion: $b_i^{EOS} = v_i$ for all $i > S$, and for all $i = 2, \dots, S$,

$$b_i^{EOS} = v_i - \frac{x^i}{x^{i-1}} \left(v_i - \frac{e_{i+1}}{e_i} b_{i+1}^{EOS} \right). \quad (1)$$

In turn, this characterization implies that the resulting allocation is efficient, in the sense that positions are assigned so that $v_1 e_1 \geq v_2 e_2 \geq \dots \geq v_n e_n$. This characterization will represent our competitive benchmark. It is particularly convenient because, as shown by Varian (2007), a simple rearrangement of the equilibrium characterization delivers testable predictions, based on the observables of the model and our dataset (namely, all variables except bidders' valuations). In particular, a bid profile is compatible with an EOS equilibrium (for some profile of valuations $(v_i)_{i \in I}$) if and only if for all $j = 2, \dots, S$:

$$\frac{e_j b_j x^{j-1} - e_{j+1} b_{j+1} x^j}{x^{j-1} - x^j} \geq \frac{e_{j+1} b_{j+1} x^j - e_{j+2} b_{j+2} x^{j+1}}{x^j - x^{j+1}}. \quad (2)$$

Coordinated Bidding: Decarolis, Goldmanis and Penta (2020) (DGP) provided a theoretical analysis of the GSP auction when some of the advertisers' bids are placed by a common agency. The agency is modeled as a subset of bidders $C \subseteq I$, which places bids jointly for their members in order to maximize their joint surplus, subject to participation and stability constraints (note that this approach abstracts from the possibility of misaligned incentives between the intermediary and the bidders). In particular, DGP put forward the notion of ‘‘Recursively-stable Agency Equilibrium’’ (RAE), which can be used to study agency bidding in general mechanisms for online ad auctions, and general agency configurations. Crucially, the RAE's framework enables to accommodate the case of partial cartels, i.e. situations in which agencies operate side-by-side with independent bidders, which is the most relevant case in the data.¹⁰

DGP provided several models of agency bidding under RAE, which correspond to progressively weaker constraints on the behavior of the agency. In the first, most restrictive model, it is assumed that the agency is constrained to placing bids that could not be distinguished from a competitive EOS equilibrium by an external observer, within a single auction, even if the independents had revealed their own valuations to the external observer. DGP's characterization of the resulting equilibrium, dubbed ‘‘undistinguishable (from EOS) coordination RAE’’ (UC-RAE), shows that the UC-RAE is efficient and essentially unique, and it is such that: (i) all independent bidders bid according to the same recursion as in equation (1); (ii) agency members place bids which are consistent with the same recursion, except that they replace the true valuations with a *feigned valuation*, $(v_i^f)_{i \in I}$, optimally chosen in order to maximize the agency's surplus, subject

¹⁰DGP's formulation of the agency problem is closely related to the notion of ‘Equilibrium Binding Agreements’ introduced by Ray and Vohra (1997) and applied in several studies, including Aghion, Antras and Helpman (2007) and Ray and Vohra (2014). It is also related to the literature on mediators in games, Monderer and Tennenholtz (2009), Ashlagi, Monderer and Tennenholtz (2009), Kalai (2010) and Roth and Shorrer (2018).

to RAE’s stability constraints. DGP show that, in equilibrium, such feigned valuations are set at the lowest possible value which ensures that agency clients maintain their efficient position.¹¹ Similar to Varian’s characterization of the competitive equilibrium (2), a re-arrangement of DGP’s characterization of the UC-RAE yields clear testable predictions: for any coalition member (other than the highest-placed) $j \in C \setminus \min \{i : i \in C\}$ who is placed immediately above an independent (i.e., such that $j + 1 \notin C$), the condition in (2) must hold with equality.

In the second model, the UC-restriction is lifted but the agency is restrained to preserving the allocative efficiency (so-called Eff-RAE). DGP show that bids in this case can be further lowered compared to the UC-RAE, which in turn directly implies that the condition in (2) is violated, in that the inequality holds with the reversed sign. Therefore, the DGP models of coordinated bidding entail the following restrictions on the observable data:¹²

$$\begin{cases} \frac{e_j b_j x^{j-1} - e_{j+1} b_{j+1} x^j}{x^{j-1} - x^j} \geq \frac{e_{j+1} b_{j+1} x^j - e_{j+2} b_{j+2} x^{j+1}}{x^j - x^{j+1}} & \text{if } j \notin C, \\ \frac{e_j b_j x^{j-1} - e_{j+1} b_{j+1} x^j}{x^{j-1} - x^j} \leq \frac{e_{j+1} b_{j+1} x^j - e_{j+2} b_{j+2} x^{j+1}}{x^j - x^{j+1}} & \text{if } j \in C \setminus \min \{i : i \in C\}. \end{cases} \quad (3)$$

4 Empirical Analysis

The system of inequalities (3) can be verified using sponsored search data on $(b_i, e_i, x^s, C)_{i \in I, s \in S}$. This is the logic of the empirical method that we develop in Step 1 (Section 4.1).

As for the next step (the revenue-loss quantification, Section 4.2), first note that the model thus far does not include any *noise* or error term, and hence it is not yet amenable to a standard empirical application. The usual econometric approach entails asking what perturbation in the data would be sufficient to make the data compatible with the (noise-free) model. In his seminal work on the sponsored search auctions, Varian (2007) suggested that the most natural variable to perturb is the ad quality, e_i . This is because the quality score is the most difficult variable for advertisers to observe and, indeed, the exact quality score of each ad is known to the search engine, but revealed to the advertisers only ex post.¹³ Because of this uncertainty about quality scores, Varian (2007) proposed the formulation of an empirical model where $d_i e_i$ is the value of the perturbed ad quality. In this model, d_i is a set of multipliers indicating how much each ad

¹¹See DGP for the recursive characterization analogous to EOS’ recursion in (1). (DGP’s main result refers to a model without quality scores, which corresponds to the case in which $e_i = 1$ for all i . However, similar to Varian (2007)’s analysis of the competitive case, it is straightforward to extend DGP’s results to the case with quality scores simply replacing the plain bids b_i in DGP with the adjusted bids $\tilde{b}_i = e_i b_i$.)

¹²In DGP, a third type of coordinated bid equilibrium is discussed and dubbed “unconstrained” RAE. Since, however, this alternative equilibrium is observationally equivalent to the Eff-RAE in our data, we ignore it in the discussion that follows and only consider the UC-RAE and Eff-RAE.

¹³Similarly, in Athey and Nekipelov (2014)’s model of bidding under uncertainty, the source of uncertainty is the quality score, along with the composition of the set of rival bidders. Interestingly, when Athey and Nekipelov (2014) use their model to infer from bid data the underlying bidders’ values, they obtain estimates that are nearly identical to those implied by applying to the same data the full information model of EOS-Varian. This underscores that the complete information model is a good approximation of bidders’ behaviour in the data. Further discussion of the complete information assumption is provided in Section 5.

quality e_i needs to be perturbed in order for the inequalities characterizing his model’s equilibrium to be satisfied. Therefore, bidders consider $d_i e_i$ to be the ad quality at the time of bidding, while the econometrician only observes the value of e_i . Step 2 of our methodology follows exactly this approach, and hence we introduce a stochastic element in our context by perturbing ad quality, $d_i e_i$, using in each keyword the model of coordinated bidding identified in Step 1 of our procedure to estimate bounds on the revenue losses compared to the competitive benchmark.

4.1 Step 1: Bid Coordination Detection

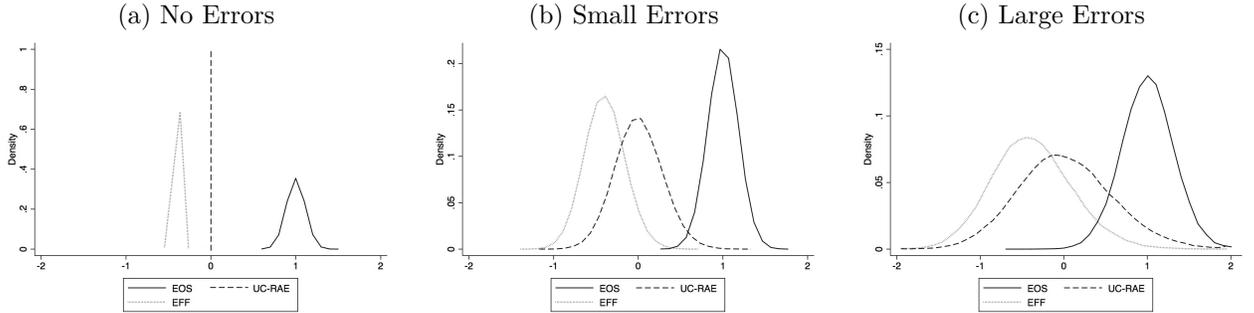
The extent to which coordination can be detected hinges on the type of data available. With *one* auction only, the UC-RAE and EOS equilibria are indistinguishable as they both satisfy inequalities (2). By construction, under UC-RAE, the conditions in (2) must hold with equality for all coalition bidders except the one with the highest valuation. This, however, suggests that when *multiple* auctions are observed for the same keyword, the UC-RAE and EOS equilibria might be distinguished. In particular, since quality scores are changed nearly in real-time by the search engine, while valuations and CTRs are likely more persistent, let us suppose that we observe a dataset with T auctions that are identical in terms of the number of bidders and their valuations and the number of slots and their CTRs, but differ for the quality scores. Then, equilibrium strategies require bids to vary across auctions to ensure that the (observable) restrictions in either (2) or (3) hold, both of which are functions of quality scores. Hence, the specific way in which bids change across auctions differs depending on which equilibrium is played. For all b_j , with $j \in \mathcal{C} \setminus \min \{i : i \in \mathcal{C}\}$, it must be that (2) holds with equality under UC-RAE, it holds with weak inequalities under competitive EOS bidding, and it is violated under Eff-RAE. This is the key idea behind our method to detect bid coordination, which we illustrate through a simple example.

Illustrative Example - Consider a keyword with four available ad slots and five bidders. Their valuations are $v = (5, 4, 3, 2, 1)$. The CTRs for the five positions are the following: $x = (20, 10, 5, 2, 0)$. If the quality scores equal 1, the bids in EOS lowest envy-free equilibrium are: $b_5 = 1$, $b_4 = 1.6$, $b_3 = 2.3$ and $b_2 = 3.15$. Now suppose that a coalition exists which comprises the first and third highest value bidders. Under coordinated bidding, $b_3 = 1.8$ under UC-RAE, and $b_3 = 1.6$ under Eff-RAE.¹⁴ We introduce variation in the quality scores across auctions by simulating 100,000 repetitions of this auction game, drawing each time an i.i.d. quality score for each bidder from a (truncated) Normal distribution with mean 1 and s.d. 0.03.

For each of these 100,000 replicas, we compute the equilibrium bid vectors under the three equilibrium models of EOS, UC-RAE and Eff-RAE. For this example with a 2-bidder coalition,

¹⁴In all cases, the highest bid $b_1 > b_2$ is not uniquely determined, but it does not affect the revenues, which equal 96 under competition, 86 under UC-RAE, and 82 under Eff-RAE. See details in DGP, Table 1.

Figure 1: Simulation of the J -statistic for one keyword under different modes of coordination and varying size of the belief errors.



the only observable difference between these equilibria involves b_3 , the bid of the lowest-value coalition bidder. We thus calculate for each auction $t = 1, \dots, T$ the following quantity J_t :

$$J_t = \frac{e_{t3}b_{t3}x^2 - e_{t4}b_{t4}x^3}{x^2 - x^3} - \frac{e_{t4}b_{t4}x^3 - e_{t5}b_{t5}x^4}{x^3 - x^4},$$

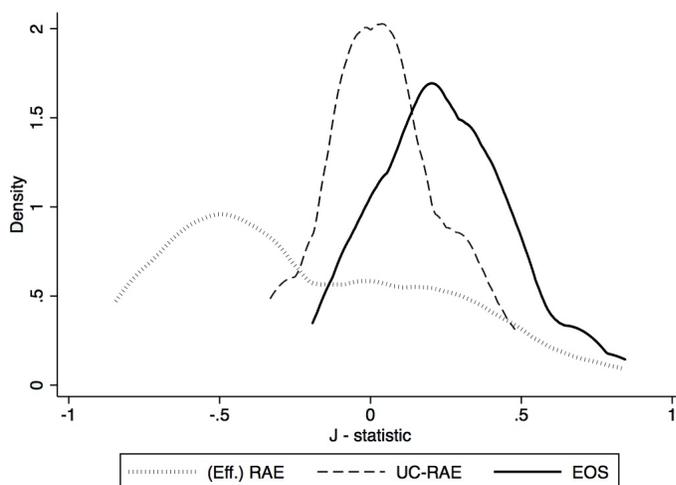
where the CTRs are not indexed by t as they are assumed to stay fixed across auctions. In each auction t , for given values of $(e_{t3}, e_{t4}, e_{t5}, b_{t4}, b_{t5})$, the value taken by b_{t3} determines whether J_t is positive, negative, or equal to zero. Under competitive bidding, b_{t3} must be high enough that $J_t \geq 0$, under UC-RAE it is as low as to make $J_t = 0$, and under the Eff-RAE it is even lower so that $J_t < 0$. Clearly, observing the distribution of J_t in the data reveals the type of equilibrium being played, and this is exactly what we see in panel (a) of Figure 1. The three curves report the distribution of J_t across the 100,000 auctions in the example above: under the UC-RAE the distribution is degenerate with a mass point at zero, while it is a non-degenerate distribution with positive support (in the case of EOS) or with negative support (in the case of Eff-RAE). \square

As discussed above, however, a useful empirical model of the search auctions must entail perturbations in the quality scores. This will clearly impact the detectability of bid coordination. In particular, suppose that for each bidder i and auction t the quality score is e_{it} , but bidders believe it to be \tilde{e}_{it} , where $\tilde{e}_{it} = d_{it} \cdot e_{it}$. Continuing from our previous example, consider two cases with different magnitudes of the belief error. For instance, a ‘small error’ case, with $d_{it} \sim \text{truncated } \mathcal{N}(1, 0.05^2)$, and a ‘larger error’ case, with $d_{it} \sim \text{truncated } \mathcal{N}(1, 0.1^2)$. Panels (b) and (c) in Figure 1 illustrate how J_t is distributed in these two cases. Not surprisingly, the presence of belief errors makes the detection harder as now the observable data used to calculate J_t differ from the information upon which bidders based their bidding choices. In panel (b), visual inspection of the distribution of J_t is still quite revealing of the type of equilibrium being played. Indeed, while there is overlap in the support of all the three distributions, it is still the case that most of the

mass lies in the positive realm in case of competition, in the negative realm in the case of Eff-RAE, and it is approximately centered around zero in the case of the UC-RAE. Although the accuracy of this approach worsens as the magnitude of belief errors grows – as shown by the comparison of panels (b) and (c) – detecting coordination should still be feasible under the moderate size of belief errors measured by the earlier literature on EOS bidding.¹⁵

Empirical Results - We now turn to the proprietary data to apply the detection method. Separately for each one of the keywords, $k = 1, \dots, 71$, we use all the available auctions to calculate the value of J_{kt} for the lowest ranked member of the 2-bidder coalitions in these data. As Figure 2 shows, the results are quite comparable to those produced by the simulation exercise in Figure 1. In particular, in Figure 2 we report the distributions of J_{kt} for three different keywords k . These are selected to be illustrative of the different types of behavior present in the data: the solid distribution is located mostly to the right of zero, thus supporting the case for EOS; the dotted distribution lies mostly to the left of zero, supporting the case for the Eff-RAE; finally, the dashed distribution is concentrated around zero, indicating that the equilibrium is the UC-RAE.

Figure 2: Three Example Keywords



Distribution of J_t for three keywords exemplifying the different equilibria.

The different behavior observed across keywords is likely associated with differences in both strategies of the bidders (and their intermediaries) and in the structural features of their markets. For instance, for a given keyword, the observed behavior might change due to the launch of a new aggressive advertising campaign by one of the bidders. Since we do not observe the essential details needed to perform such an in-depth analysis, we propose a simple classification that partitions at the keyword level the instances of coordination and of competition. Hence, for each of the keywords,

¹⁵Varian (2007) and Athey and Nekipelov (2014) both found that typically a small (on average 5 percent) belief perturbation suffices to rationalize EOS bids.

we calculate a 95% confidence interval for the median of J_{kt} , and we classify k as competitive if the lower bound of the confidence interval is positive, as Eff-RAE if its upper bound is negative, and as UC-RAE if it includes zero. We find that most keywords are classified as collusive, with 36 of them being classified as UC-RAE, and 3 as Eff-RAE. The remaining 32 are classified as competitive.

We note that there are several important caveats to this method. For instance, contrary to the simulation where the quality scores were independent identically distributed draws, this is unlikely to be the case in our data, thus implying that a different weighting of the observations might be needed to calculate the distributions of the J_{kt} . In Section 4.3 we explore a further check of the validity of our approach (that is, to detect coordination based on the patterns in the distribution of J_t), and we also discuss some of the caveats and extensions in Section 5.

4.2 Step 2: Revenue Loss Quantification

The second step consists of evaluating the revenue effects of coordination. It exploits the canonical approach in the structural estimation of auction games: given the observables used in step 1, and an equilibrium that maps the unobservable (to the econometrician) valuations to the bids, the equilibrium mapping can be inverted to back out valuations from bids. In this sense, the first step of the procedure above is important to guarantee that what is imposed on the data is a sensible equilibrium model.¹⁶

Note, however, that not all valuations can be point identified under a bid coordination equilibrium.¹⁷ In particular, point-identification of valuations will be possible for the independent bidders, for whom there is a one-to-one mapping between their bid and value (except for the overall highest ranked bidder). But for coalition bidders, their incentive to shade their bids implies that there is a range of valuations that would all be compatible with the observables. Nevertheless, by exploiting the equilibrium property that coalition bidders are ranked efficiently, relative to both other coalition members and to the independents, we can pin down bounds on the valuations of coalition bidders. To see this, consider again our earlier example with 5 bidders. The first and third highest bidders are part of a coalition. If we observe the full bid vector and assume that it is the outcome of a UC-RAE, then inversion of the equilibrium mapping implies that $v = (z_1, 4, z_3, 2, 1)$ and, by efficiency, also that $z_3 \in [2, 4]$. Although no bound can be derived when the coalition that occupies the top two slots or when its lowest valued member has no bidder below it, in all other cases this approach is informative and allows us to construct counterfactual bounds on revenues under competitive bidding.

¹⁶Since intermediaries are many and heterogeneous, assuming that they all play a specific type of coordination equilibrium would have clearly been more restrictive compared to proceeding in steps as we do, i.e. by first detecting whether coordination is present or not and, if present, whether it is more likely to be a case of UC-RAE or Eff-RAE.

¹⁷Valuations are also not point identified in Hortaçsu and McAdams (2010) and Kastl (2011)'s analysis of multi-unit auctions. There, partial identification comes from the optimality of the individual bids. In our paper, instead, the agency bidders' valuations are bounded by others' valuations due to the efficiency of the equilibrium allocation.

The presence of unobserved variability in the quality scores introduces some nuances into this approach. What is needed in this case is a method to infer valuations from bids, allowing for bids to be based on the perturbed quality scores. Hence, let $\tilde{b}_i = d_i e_i b_i$. Valuations are recovered by first estimating the perturbations d_i , and then by including these estimates into the bids' inversion procedure. In particular, first note that $(d_i)_{i \in I}$ in our approach is a vector of multipliers, one for each bidder, indicating how much each ad quality needs to be perturbed to satisfy all of the equilibrium inequalities. As in Varian (2007), the method works by measuring the magnitude of the departure from the underlying model of optimizing behavior. Ideally, we would like each of the individual perturbations to be as close to one as possible: if some data fail the conditions summarized by the equilibrium inequalities, but only by a small amount, we might attribute this failure to uncertainty about the quality. For instance, suppose that the data are assumed to be generated under a UC-RAE, then, the perturbations $(d_i)_{i \in I}$ are chosen in order to solve the following program:

$$\begin{aligned} \min_d \sum_{i>1} (d_i - 1)^2 \quad \text{subject to:} \\ \left\{ \begin{array}{ll} \frac{\tilde{b}_i x^{i-1} - \tilde{b}_{i+1} x^i}{x^{i-1} - x^i} \geq \frac{\tilde{b}_{i+1} x^i - \tilde{b}_{i+2} x^{i+1}}{x^i - x^{i+1}}, & \text{if } i \notin \mathcal{C} \text{ or } i \in \{\min(\mathcal{C})\}; \\ \tilde{b}_i x^{i-1} = \frac{x^i - x^{i+1}}{x^{i+1} - x^{i+2}} [\tilde{b}_{i+2} x^{i+1} - \tilde{b}_{i+3} x^{i+2}] + \gamma d_i e_i [x^{i-1} - x^i] + \tilde{b}_{i+1} x^i, & \text{if } i \in \mathcal{C} \setminus \{\min(\mathcal{C})\}; \end{array} \right. \end{aligned}$$

where γ is the minimum bid increment (5 cents in the data). In words, the solution to the program above finds the smallest d such that the UC-RAE restrictions are satisfied. Varian (1985) provides an extensive discussion of the statistical interpretation of the minimal perturbation quantification, and of its good statistical properties compared to several other methods, including parametric approaches. To further clarify the approach, note that even if there might exist other perturbations of qualities $(d_i)_{i \in I}$ that would satisfy the set of equilibrium inequalities and equalities, it is not the case that we can rationalize the data with any collection of d_i that are larger than those that are currently found, due to the interdependencies that the program generates across different i 's.¹⁸

Provided with the estimated $(\hat{d}_i)_{i \in I}$, we can proceed as in the example above and point identify, for each bid placed by an independent, the corresponding valuation. From that, we then obtain bounds on the coalition bidders' values. Finally, assigning to each independent bidder its estimated value and to the coalition bidders their estimated upper bound valuations, we can solve for the competitive EOS equilibrium and compare its revenues to the ones under coordinated bidding.¹⁹

¹⁸That is most clearly seen by noting that each of the perturbations enters several equations. For example, let's consider independent bidder i , and increase the perturbation of the player $i + 1$ significantly. In that case, the inequality $\frac{\tilde{b}_i x^{i-1} - \tilde{b}_{i+1} x^i}{x^{i-1} - x^i} \geq \frac{\tilde{b}_{i+1} x^i - \tilde{b}_{i+2} x^{i+1}}{x^i - x^{i+1}}$ will no longer be satisfied. A similar formula applies in the case of the Eff-RAE. It can be obtained by replacing the constraint for $i \in \mathcal{C} \setminus \{\min(\mathcal{C})\}$ in the formula above with the corresponding ones in DGP.

¹⁹We only quantify upper bounds, since the lower bounds must coincide with those valuations that entail no revenue losses. This scenario, in turn, corresponds to the case in which the agency sets the same bids that its clients would have placed as independents.

Empirical Results - We apply this revenue-loss quantification method to the subset of keywords classified above as UC-RAE. Separately for each of these 36 sets of keyword auctions, we obtain bounds on the revenues for each keyword. In Table 2, we report the mean revenue across the 36 keywords. In the first column, we report the observed revenues. They are normalized to 100, while all other revenue figures are expressed as a percentage of the total observed revenues. The top row reports the total revenues across all bidders, while the following two rows offer a breakdown between revenues originating from payments by the coalition members and by independents. The second column reports the counterfactual upper bound revenues under competitive bidding, while the last column reports the difference between the previous two columns and, in parenthesis, the 95 percent confidence interval of a matched-pairs t-test that the means difference equals zero. This test is valid under the assumption that keywords are independent. Importantly, we estimate the revenues for each of the 36 keywords separately, whereas the matched-pairs t-test is applied to the mean revenues across 36 keywords.

Table 2: Revenue Effects for the 36 UC-RAE Keywords

| | Observed | Counterfactual Upper Bound | Difference $\Delta = UpperB. - Obs.$ |
|---------------------------------------|----------|-------------------------------|---|
| Normalized Total Revenues | 100 | 107.90 | 7.9 [5.32; 10.44] |
| Payments from Agency Advertisers | 33.20 | 35.28 | 2.08 [1.49; 2.68] |
| Payments from Independent Advertisers | 66.80 | 72.62 | 5.82 [3.73; 7.91] |

Separately for each of the 36 keywords, the normalized revenues set total observed revenues (i.e., the sum of all payments across all auctions for the same keyword) equal to 100. The three rows report: total revenues, revenues originating from the payments by agency advertisers; revenues originating from payments by independent advertisers. The three columns report the observed (normalized) revenues, the upper bound of the counterfactual revenues and the difference between the two. The values in the squared bracket are the endpoints of a 95% confidence interval for matched differences in the average revenues.

Overall, we find a statistically significant revenue loss due to bid coordination of up to 7.9 percent. Interestingly, most of the loss originates not from the direct effect of the reduced bids by coalition bidders, but from the *indirect* effect of reduced bids by the independents. The presence of this *indirect* effect taking place through equilibrium bidding underscores why even small coalitions with members occupying low-ranked slots can significantly hurt revenues. Once again, however, the correlation across auctions requires caution in interpreting the findings, as discussed next.

4.3 The J_t distribution: a separate test of validity

In this subsection, we explore a further check of the validity of our approach (that is, to detect coordination based on the patterns in the distribution of J_t). One may wonder how the distributions for the keywords for which we detected coordination compare with a possible *reference* distribution

of J_t , in auctions which are known to involve no bid coordination.²⁰ The most obvious way of doing this would be to consider a dataset of “comparable keywords” where no coalition is present and then compare the J_t distributions for keywords with and without coalitions. Unfortunately, such a direct solution is not viable within our dataset, since the provider had specifically constructed it in order to have exactly two bidders under a common intermediary for all keywords. We thus developed an alternative solution to explore this idea, which we believe to be the second-best.

Specifically, we used publicly available data from Yahoo!, easily accessible from the website <https://webscope.sandbox.yahoo.com/catalog.php?datatype=a&did=21>, which contain detailed information on the positions, bids, impressions, and clicks for a large number of advertiser-keywords pairs over 123 days, from January 2008 to April 2008.²¹ For a better comparison with the data used in the previous sections of this paper, we focused on the auctions in this dataset with exactly nine bidders. Unfortunately, the same dataset does not also contain information about agencies but based on our general knowledge and given what is observed in the SEMRush data, agency bidding was not a common phenomenon at the time of these data. These observations, as well as the fact (which will be shown below) that the J_t distribution for all of the bidders in these auctions looks a lot like the EOS distribution, makes us confident that the distributions of J_t in these auctions can be safely taken as a meaningful *reference* against which to compare the ones that we obtained from our original dataset.²²

To calculate all J_t for this dataset, we first need to estimate the quality scores. We follow a similar method as in Agarwal and Mukhopadhyay (2016), and infer quality scores from the available ad positions and bids for every keyword. Finally, we calculate the J_t across all keywords and create the reference distribution (a description is included in the Appendix, where we also provide a “one-click” code that creates the reference distribution). As we already mentioned, it is important to note that the resulting distribution matches closely the EOS distribution, and that the 95% confidence interval for the median is [0.07004, 0.0706], which is significantly above zero. Hence, the distribution thus obtained can safely be taken to play the role of benchmark that we explained above, and in the following, it will be referred to as the *reference distribution*.

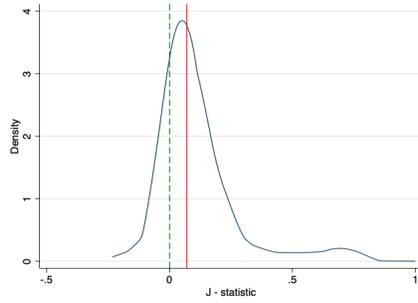
Importantly, once we create a joint distribution of J_t based on all the keywords combined, the median is statistically different from the median of the reference distribution (See panel (a) of Figure 4 below). The comparison between the distributions of the J_t under the reference distribution and for the three keywords from Figure 2 is presented in panels (b)-(d) of Figure 4 below. We also consider an alternative classification criterion to demarcate competitive versus coordinated bidding. In particular, we classify the keyword as competitive if its median is higher than the

²⁰We are thankful to an anonymous referee for spurring us to explore this direction.

²¹Yahoo!’s search market share in the United State was about 20% in 2008 <https://www.statista.com/statistics/269668/market-share-of-search-engines-in-the-united-states/>.

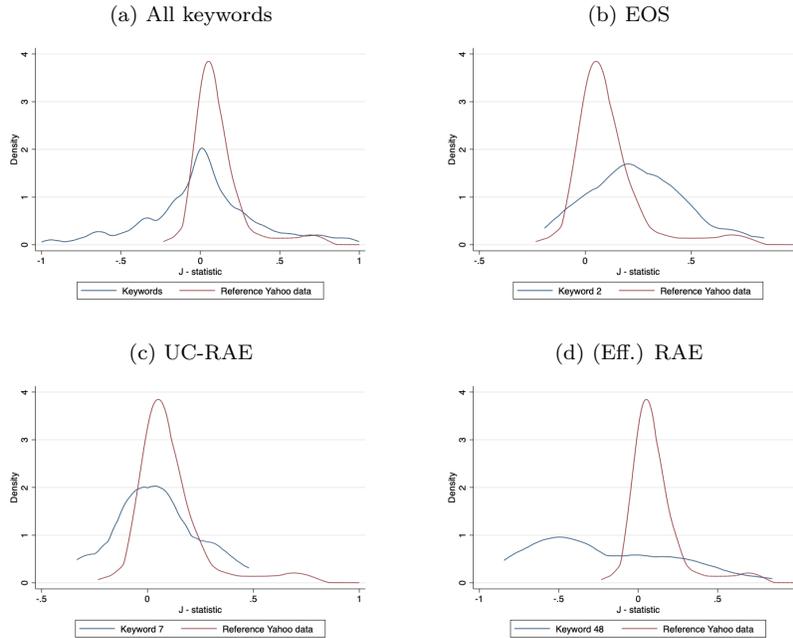
²²Moreover, note that should agency bidding be nonetheless present in such a dataset, then it would mean that our second-best approach would be more conservative (i.e., detect coordination in fewer cases) than the method that uses the distribution with no agency bidding.

Figure 3: The *reference distribution*.



lower bound of the 95 percent confidence interval of the median of the reference distribution, and collusive otherwise. This is a less conservative method (i.e., more likely to detect coordination) compared to the main one we developed in the previous sections of this paper because the lower bound of the confidence interval of the median of the reference distribution is substantially greater than zero. As a result, compared to the 55% of the keywords for which coordinated bidding was detected, with this alternative method we detect it for 70% of the keywords.

Figure 4: J_t for the keywords in the paper vs reference distribution



5 Discussion

The analysis above illustrates the potential of using search auction data to detect bid coordination and quantify its effects on revenues. We conclude with a discussion of the possible limitations of the proposed methodology and of their potential solutions.

First, as we already mentioned, our bid detection method does not account for the possible serial correlation across auctions.²³ For instance, in the simulations in Figure 1, the i.i.d. draws for quality score allow for a clear interpretation of the distribution of the J_t statistic. Serial correlation across auctions might require a different weighting of the observations since some of them might be more informative than others about the type of equilibrium being played. Choi and Varian (2012), however, studied the time series structure of a large sample of sponsored search auctions and concluded that there is not a unique time series model that can be applied generally across different keywords. Correcting for the possible serial correlation across auctions would thus require a preliminary in-depth analysis of each single market, as the correct specification of the time series would require detailed information on the process through which bidders' valuations evolve for a specific keyword. Our approach therefore can still be seen as a faster (if perhaps rougher) way to detect agency bidding strategies across keywords, in the absence of detailed information on the keyword-specific time series structure.

Second, the classification criterion that we adopted in Section 4.1 is based on the location of a confidence interval around the median. Clearly enough, different criteria could also be used (e.g., criteria based on the mode, on the concentration of the mass of the distribution, etc.) Our results are robust to several of these changes. For instance, we obtain an identical classification if we calculate the smallest interval of their support including 80% of the mass: we classify k as competitive if the lower bound of this support is positive, as Eff-RAE if the upper bound of the support is negative, and as UC-RAE if the interval includes zero. The classification changes substantially if instead we use the mean of J_t . In this case, the presence of outliers produces less interpretable results.

Another important assumption that we maintain from DGP is that of complete information. The strategy we adopted in the previous sections, that is to introduce perturbations to quality scores in a baseline model with complete information, is supported by the findings in both the empirical and experimental literature, which agree that estimates of bidders' valuations based on our approach (first introduced by Varian (2007)) are very close to those based on more sophisticated models that account for different sources of incomplete information.²⁴ The first case in point is

²³Such serial correlation could be the result of a constraint optimization problem that some advertisers might follow given a fixed budget, which we do not consider. Importantly, Athey and Nekipelov (2014) implicitly considers the case in which bidders specify budgets which the ad platforms respect by spreading out the advertiser's participation in auctions over time, withholding participation in a fraction of auctions. The fact that the estimated valuations in Athey and Nekipelov (2014) are close to those implied by the EOS-Varian complete information model suggests that the existence of a budget constraint is not a primary concern.

²⁴The complete information model of competitive bidding that we adopt also has its own theoretical appeal, as

provided by Varian (2007), who shows that the inequalities we introduced in Section 3, which characterize our baseline model of competitive bidding under complete information, are largely consistent with the data. Another influential example is provided by Athey and Nekipelov (2014), who specify a model with uncertainty over both quality scores and the set of bidders and obtain estimates that are very close to those implied by the EOS-Varian complete information model. Further evidence in this sense can be gathered by the experimental results in Che, Choi and Kim (2017) and McLaughlin and Friedman (2016), which confirm that the static complete information model closely approximates the dynamic incomplete information setting, in that they entail similar estimates of the platform revenues.

An interesting extension for future work would be to identify whether the coalition’s bids are too low to be competitive, based on the J_t statistics computed over the independents, instead of the lowest coalition members. This alternative strategy – an application of the randomization inference of Rosenbaum (2002) – would be particularly helpful to test for a tendency of coalition bidders to place suspiciously low bids, without imposing the equilibrium assumptions of the DGP coordination models. Taking this route has both pros and cons, the main downside being losing the ability to connect the results of the detection step to the calculation of the counterfactual revenues in the second step. This is the main reason why it was not pursued in this study.

Our analysis abstracted from the possibility that intermediaries enforce more complex forms of coordination that go beyond bid coordination. For instance, agencies might split the markets by allocating their clients to different keywords, or to the same keyword but split the targeted audiences (targeting options are abundant in sponsored search auctions, and algorithmic bidding makes it easy to arrange bidding strategies aimed at reducing direct competition between an agencies’ clients in the same auction.) This kind of coordinated strategies would entail even stronger downward pressures on the cost-per-click. Hence, if agencies engage in this kind of coordination, the actual revenue losses might be even larger than those identified by our methodology. It should be pointed out, however, that market splitting is not as profitable in multi-item auctions as they normally are in standard single-item auctions, since they may come at the cost of forgoing the surplus that the agencies’ clients may still make by obtaining a different slot, and crowding out non-agency bidders. In fact, coalition bidding in the same auctions is indeed frequent in the data, as confirmed in both our proprietary data and the SEMrush data discussed in Section 2. Hence, at least for some keywords, agencies’ strategy does not entail a complete market split. Nevertheless, an empirical analysis that looks more comprehensively at the effects of intermediaries would be a highly valuable extension of our analysis.

Our analysis also abstracted from the possibility of externalities among bidders. In particular, shown by EOS, Varian (2007) and (Milgrom and Mollner, 2018). Moreover, the model of individual bidding on which we base our analysis conforms with the search engines’ tutorials on how to bid in these auctions (see, e.g., the Google AdWord tutorial in which Hal Varian teaches how to maximize profits: <http://www.youtube.com/watch?v=jRx7AMb6rZ0>. For further discussions, see Decarolis, Goldmanis and Penta (2020).

both theoretical and empirical studies have considered CTRs that depend on the identities and positions of advertisers. The literature refers to this case as CTR-externalities (see, among others, Gomes, Immorlica and Markakis (2009) and Jeziorski and Segal (2015)). An interesting avenue for future research would be to extend the theoretical model of DGP to accommodate the case of externalities, similarly to what was done for the competitive bidding case by Gomes, Immorlica and Markakis (2009). Provided with a theoretical characterization of coordinated bidding with CTR externalities, it should be possible to extend the empirical method described in this paper to detect coordination under such a richer environment.

Finally, the estimated valuations might be helpful to evaluate the impact of possible changes in the auction design. For instance, Google has increased the reserve price in its auctions for the first time in May 2017. While this change might help limit the revenue losses caused by bid coordination, it might end up hurting also non-coordinating bidders. Hence, monitoring evolutions in the market is certainly worthwhile to better understand who are winners and losers in this market.

References

- Agarwal, Ashish, and Tridas Mukhopadhyay.** 2016. “The impact of competing ads on click performance in sponsored search.” *Information Systems Research*, 27(3): 538–557.
- Aghion, Philippe, Pol Antras, and Elhanan Helpman.** 2007. “Negotiating free trade.” *Journal of International Economics*, 73(1): 1–30.
- Ashlagi, Itai, Dov Monderer, and Moshe Tennenholtz.** 2009. “Mediators in position auctions.” *Games and Economic Behavior*, 67(1): 2–21.
- Athey, Susan, and Denis Nekipelov.** 2014. “A Structural Model of Sponsored Search Advertising Auctions.” *mimeo*.
- Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek.** 2013. “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence.” *American Economic Journal: Microeconomics*, 5(4): 163–187.
- Che, Yeon-Koo, Syngjoo Choi, and Jinwoo Kim.** 2017. “An experimental study of sponsored-search auctions.” *Games and Economic Behavior*, 102: 20 – 43.
- Choi, Hyunyoung, and Hal Varian.** 2012. “Predicting the Present with Google Trends.” *Economic Record*, 88(s1): 2–9.
- Decarolis, Francesco, and Gabriele Rovigatti.** 2021. “From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising.” *American Economic Review*, forthcoming.

- Decarolis, Francesco, Maris Goldmanis, and Antonio Penta.** 2020. “Marketing Agencies and Collusive Bidding in Online Ad Auctions.” *Management Science*, Vol.66, No 10, 4433-4454.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz.** 2007. “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords.” *American Economic Review*, 97(1): 242–259.
- Ghose, A, and S Yang.** 2009. “An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets.” *Management Science*.
- Goldman, Matthew, and Justin Rao.** 2015. “Experiments as Instruments: Heterogeneous Position Effects in Sponsored Search Auctions.” *EAI Endorsed Transactions on Serious Games*, 3.
- Gomes, Renato, and Kane Sweeney.** 2014. “Bayes-Nash Equilibria of the Generalized Second Price Auction.” *Games and Economic Behavior*, 86: 421–437.
- Gomes, Renato, Nicole Immorlica, and Evangelos Markakis.** 2009. “Externalities in Keyword Auctions: An Empirical and Theoretical Assessment.” International Workshop on Internet and Network Economics, Springer.
- Hortaçsu, Ali, and David McAdams.** 2010. “Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market.” *Journal of Political Economy*, 118(5): 833–865.
- Hsieh, Yu-Wei, Matthew Shum, and Sha Yang.** 2018. “To Score or Not to Score? Estimates of a Sponsored Search Auction Model.” *USC-INET Research Paper No. 15-09*.
- IAB.** 2019. “Internet Advertising Revenue Report, Full Year 2018.” *PwC Report for the Internet Advertising Board*.
- 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.
- Jeziorski, Przemyslaw, and Sridhar Moorthy.** 2018. “Advertiser Prominence Effects in Search Advertising.” *Management Science*, 64(3): 1365–1383.
- Kalai, Adam Tauman.** 2010. “A commitment folk theorem.” *Games and Economic Behavior*, 69(1): 127 – 137. Special Issue In Honor of Robert Aumann.
- Kastl, Jakub.** 2011. “Discrete bids and empirical inference in divisible good auctions.” *The Review of Economic Studies*, 78(3): 974–1014.

- Lewis, R. A., and J. M. Rao.** 2015. “The Unfavorable Economics of Measuring the Returns to Advertising.” *Quarterly Journal of Economics*.
- McLaughlin, Kevin, and Daniel Friedman.** 2016. “Online Ad Auctions: An Experiment.”
- Milgrom, Paul, and Joshua Mollner.** 2018. “Equilibrium Selection in Auctions and High Stakes Games.” *Econometrica*, 86(1): 219–261.
- Monderer, Dov, and Moshe Tennenholtz.** 2009. “Strong mediated equilibrium.” *Artificial Intelligence*, 173: 180–195.
- Ray, Debraj, and Rajiv Vohra.** 1997. “Equilibrium Binding Agreements.” *Journal of Economic Theory*, 73(1): 30 – 78.
- Ray, Debraj, and Rajiv Vohra.** 2014. “The Farsighted Stable Set.” *Econometrica*, 83(3).
- Rosenbaum, Paul.** 2002. *Observational Studies*. Springer.
- Roth, Benjamin, and Ran I. Shorrer.** 2018. “Making it Safe to Use Centralized Marketplaces: Dominant Individual Rationality and Applications to Market Design.” *WP, SSRN 3073027*, 173.
- Varian, Hal.** 2007. “Position auctions.” *International Journal of Industrial Organization*, 25(6): 1163–1178.
- Varian, Hal R.** 1985. “Non-parametric analysis of optimizing behavior with measurement error.” *Journal of Econometrics*, 30(1-2): 445–458.