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

Francesco Decarolis
Maris Goldmanis
Antonio Penta
Ksenia Shakhgildyan

May 2021

Barcelona GSE Working Paper Series

Working Paper nº 1258

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

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

May 21, 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. Using proprietary data from auctions held on a major search engine, this study develops a methodology to detect bid coordination. It also presents a strategy to estimate a bound on the search engine revenue losses imposed by coordination relative to a counterfactual benchmark of competitive bidding. In the data, coordination is detected in 55 percent of the cases of delegated bidding observed and the associated upper bound revenue loss for the search engine 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 the participants at the seminars 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. Francesco Decarolis acknowledges financial support from the PRIN 2017. Antonio Penta acknowledges financial support from the ERC Starting Grant #759424. This research also benefited from the support of the Spanish Ministry of Science and Innovation, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S)

†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 GSE, TSE and TSE Digital Center, *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. 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 features such as the location of the coalition advertisers in the ad ranking, the value that different

[1] As 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.

[2] Competitive 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.

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. This 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 develop in this paper consists of two steps. First, by exploiting repeated observations (i.e., auctions) for the same keyword, our methodology determines which behavioral model between collusion and competition fits the data. Then, the second step of the procedure uses bids, and DGP's theoretical results on the coordination strategy 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 and to separately quantify the *direct* and *indirect effects* of coordination mentioned earlier.

The novelty of our contribution is not the application of the method in the second step, which closely tracks ideas in Varian (2007), but the development of a new method to detect bid coordination in the first step. Indeed, while DGP offers a theoretical characterization of bid coordination in the GSP and VCG auctions, the present paper develops a way to operationalize those ideas into an empirical approach that can detect coordination and, hence, enable us to unfold the quantitative implications of DGP's theoretical analysis. All this methodology is novel relative to DGP, although it is obviously built starting from their theoretical framework.

An important assumption that we maintain from DGP is that of complete information regarding bidder valuations. We follow the same modelling assumption, but also provide some new results to support it. In particular, we also derive an estimator for bidders' valuations under the incomplete information setting proposed by Gomes and Sweeney (2014), and show through Monte Carlo simulation that the estimates based on the complete information model approximate well those obtained from the incomplete information model.

Finally, we present 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). The application of the two-step method reveals that: (i) coordinated bidding is detected in 55 percent 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 undistinguishable 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 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.[3] 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 selection from a representative sample of the search auctions involving some of the most frequently searched keywords.[4] The selection contains all the auctions for 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.[5] 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.[6] Hence, to help

[3]The company name cannot be disclosed, due to a confidentiality agreement.

[4]The search engine gave us access to these data for research purposes. The data are considered to be sufficiently old not to represent a threat to its business. However, extensive discussion with the managers at this search engine convinced us that, despite old, the data is representative of bidding behavior in the last decade. Where we do foresee the possibility of more substantial qualitative changes, is in the data from the last couple of years, during which artificial intelligence algorithms have started to become more prominent, somewhat changing the patters in bidding behavior, but not before that. Finally, we shall remark that the novel methodology we put forward can be applied broadly to data from any period in which the rules of the auction are set according to the GSP auction rules.

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

[6]In 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.

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.[7] 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 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 full sample of 1,402 keywords (last four table columns), from those for the subsample of 1,102 keywords with at least one ad placed by an intermediary (first four table columns). 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, especially for the subsample of delegated bidding, thus underscoring the relevance of intermediaries for keywords with high potential revenues for the search engine. 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:

[7] As stated earlier, the theoretical counterpart of this paper is DGP. However, in the closely connected realm of display ad auctions Mansour, Muthukrishnan and Nisan (2012), Balseiro and Candogan (2017) and Allouah and Besbes (2017) all offer insightful theoretical models on the role of intermediaries in these auctions. At a more general level, the analysis of mediators in games has been introduced by Monderer and Tennenholtz (2009) for complete information settings and by Ashlagi, Monderer and Tennenholtz (2009) for incomplete information ones, with important extension offered in Kalai (2010) and Roth and Shorrer (2018).

[8] We combine the dataset in Decarolis and Rovigatti (2021) with the list of the top 10,000 keywords on Google US (in terms of number of searches) in 2017. We restrict the attention to these popular keywords to enhance the comparability with the proprietary data. Further details on the data are presented in Decarolis and Rovigatti (2021). An important feature discussed there regards the ownership structure of the intermediaries and, in particular, the fact that it is not the case that they are owned by Google, see their appendix K on this.

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

	Keywords with at Least 1 Network				Full Sample			
	Mean	Median	SD	Obs	Mean	Median	SD	Obs
Cost-per-click (CPC)	1.53	0.81	2.66	1,102	1.76	0.85	3.51	1,402
Search Volume	6.51	0.25	2.54	1,102	1.13	0.25	12.21	1,402
# of Words	1.85	2.00	0.82	1,102	1.87	2.00	0.81	1,402
# of Characters	10.86	10.00	4.96	1,102	11.31	11.00	5.00	1,402
Long Tail	0.03	0.00	0.18	1,102	0.03	0.00	0.18	1,402
Coalition	0.41	0.00	0.49	1,102				
Coalition Size	2.78	2.00	1.17	449				

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.

(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 the following, 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 $i \in I = \{1, \ldots, n\}$ to slots $s = 1, \ldots, S$, $n \geq S$ 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 to 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 $x^1 > x^2 > \cdots > x^S > 0$, 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 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 such 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, $\bar{b}_i = b_i \cdot e_i$: 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 setting is characterized by the following recursion: $b_i^{EOS} = v_i$ for all $i > S$, and for all $i = 2, \ldots, S$,

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

In turn, this characterization implies that the resulting allocation is efficient, in the sense that positions are assigned so that $v_1 e_1 \geq v_2 e_2 \geq \cdots \geq v_n e_n$. This characterization will represent our competitive benchmark. It is particularly convenient because, as shown by Varian (2007), 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 $j = 2, ..., S$:

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

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 modelled as a subset of bidders $C \subseteq I$, which places bids jointly for their members

[9] As 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.

in order to maximize their joint surplus, subject to participation and stability constraints. 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 which 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 instead place bids which are consistent with the same recursion, except that they replace the true valuations with *feigned valuation*, $(v_i^f)_{i\in I}$, optimally chosen in order to maximize the agency's surplus, subject 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 (equation (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) $j \in C \backslash \min\{i : i \in C\}$ who is placed immediately above an independent (i.e., such that $j+1 \notin C$), the condition in equation (2) must hold with equality.

In the second model, the UC-restrictions is lifted but the agency is restrained to preserving the allocative efficiency (so called Eff-RAE). DGP show that bids can be further lowered compared to the UC-RAE, which in turn directly implies that the condition in equation (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]

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

[10]DGP'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).

[11]See DGP for the recursive characterization analogous to EOS' recursion in (2). (DGP's main result refer 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 $\bar{b}_i = e_i b_i$.)

[12]In 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.

4 Empirical Analysis

The system of inequalities (3) can be verified using sponsored search data on $(b_i, e_i, x^s, \mathcal{C})_{i \in I, s \in S}$. Bringing the model to the data, however, requires making a few additional assumptions. First of all, the previous model is static, and it characterizes the equilibrium for a single auction, considered in isolation. We maintain this structure in the empirical work by treating the multiple auctions in the data as repetitions of the same auction game, each with its own idiosyncratic bidder valuations and quality scores. Thus, we do not model potential linkages between auctions, either over time or across keywords. Second, inequalities (3) characterize equilibrium strategies in an environment with complete information. Applying this framework to real world data, in which the search engine updates quality scores in real time, as in our case, requires considering that bidders do not necessarily know the exact quality scores at the time when they submit their bids. But, in principle, bidders' uncertainty might be even more extensive, and perhaps also involve features like others' valuations and bids. The question of what information structure best describes GSP auctions is extensively debated in the literature, and most studies agree that complete information offers an adequate approximation of the richer incomplete information models. In section 5 below, we review this literature and also present some novel results comparing the estimates of bidder valuations obtained from the complete information model, with those obtained from the incomplete information model of Gomes and Sweeney (2014), suitably extended to account for quality scores. Since these results confirm the prevailing view in the literature, on the adequacy of the complete information approximation, in the analysis below we maintain the complete information framework.

Nevertheless, the DGP model does not include any *noise* or error term, and, hence, it is not directly 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.

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 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.

We follow this same approach and introduce a stochastic element in our context by perturbing ad quality, $d_i e_i$. Based on this empirical model, we develop an empirical method to determine whether intermediaries are adopting competitive or coordinated bidding strategies. Furthermore, when coordination is present, we illustrate how to estimate bounds on the search engine revenue losses, relative to a competitive benchmark. The next two sections discuss these two elements.

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 undistinguishable as they both satisfy inequalities (2). By construction, under UC-RAE for all coalition bidders, except the one with the highest valuation, the conditions in (2) must hold with equality. This, however, suggests that, when multiple auctions are observed for the same keyword, the UC-RAE and EOS equilibria might become distinguishable. In particular, since quality scores are changed nearly in real time by the search engine, while valuations and position effects are likely more persistent, suppose 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 position effects, but differ for the quality scores. Equilibrium bidding requires bids to vary across auctions to ensure the satisfaction of the (observable) restrictions imposed on the data by the systems either in (2) or (3), both of which are indeed functions of quality scores.

The specific way in which bids change across auctions differs depending on which equilibrium is played. For all b_j, with $j \in \mathcal{C} \backslash \min\{i : i \in \mathcal{C}\}$, it must be that (2) holds with equality under the UC-RAE, holds with weak inequalities under competitive EOS bidding and is violated under the Eff-RAE. This is the key idea behind our method to detect bid coordination, which we illustrate through a simple example.

Illustrative Example:

A. Simulated Data To illustrate how the detection criterion works, we first simulate a dataset based on the leading example in DGP. We thus consider an auction with four 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 and 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.[13] 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 then compute the equilibrium bid vectors separately under the three equilibrium models of EOS, UC-RAE and Eff-RAE. We thus end up with a dataset $(b_{ti}^{EOS}, b_{ti}^{UC-RAE}, b_{ti}^{Eff-RAE}, e_{ti}, x^s)$, with $i = 1, ..., 5$, $s = 1, ..., 5$ and $t = 1, ..., 1000$, and where the position effects are not indexed by t as they are assumed to stay fixed across auctions.

[13]In 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.

(a) No Errors (b) Small Errors (c) Large Errors

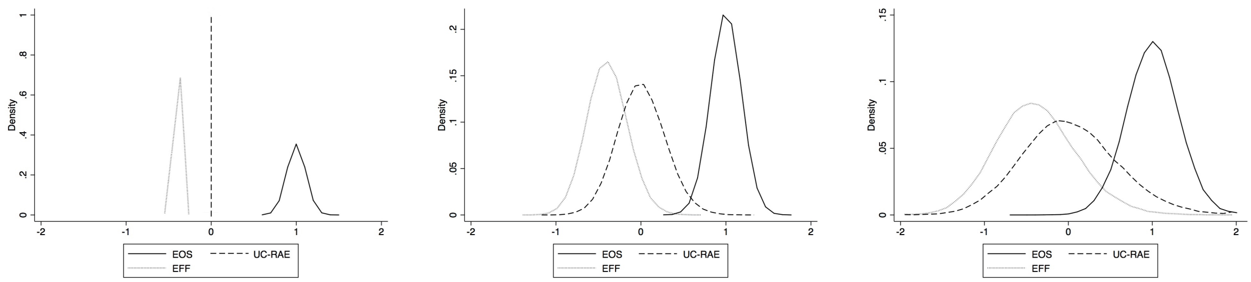

B. Detection Method In this dataset, in each auction t there is always the same 2-bidder coalition formed by the first and third highest value bidders. This implies that the only observable difference between the three equilibria must involve b_3, the bid of the lowest-value coalition bidder.[14] In particular, calculate for each auction $t = 1, ..., T$ the following quantity J_t:

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

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 is as low as to make $J_t = 0$ and under the Eff-RAE it is even lower so that $J_t < 0$. Clearly enough, observing the distribution of J_t in the data is fully revealing of 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 $d_{it} \sim \mathcal{N}(1, 0.05^2)$, and a 'larger error' case, with $d_{it} \sim \mathcal{N}(1, 0.1^2)$. 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

[14] Notice that the lowest ranked member of the 2-bidder coalition is the coalition bidder who has the lowest adjusted bid (defined as $b_i e_i$). As a result, this bidder occupies the lowest position among coalition bidders. In our illustrative example, where quality scores are just a small fraction of valuations, the ranked coalition bidder is always the lowest-value coalition bidder.

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 case of Eff-RAE and it is approximately centred 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 in 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, $k = 1, ..., 71$, 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

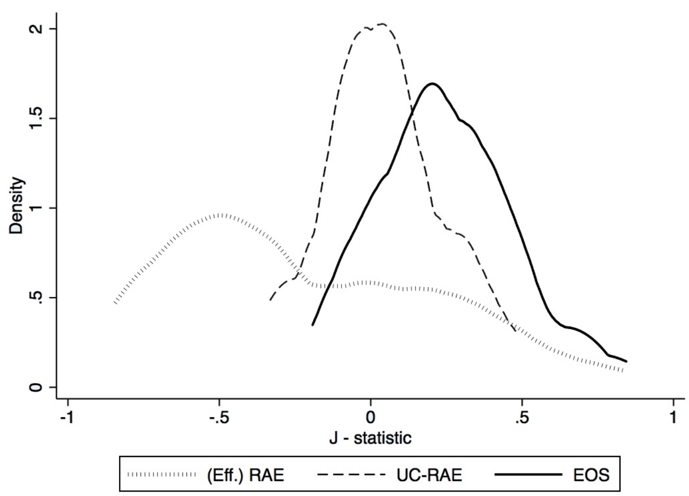

Distribution of J_t for three keywords exemplifying the different equilibria.

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

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

needed to perform such an in depth analysis, we propose a simple classification that partitions at keyword level the instances of coordination and of competition.

Therefore, for each of the keywords, we calculate a 95 percent confidence interval for the median of J_{kt}: 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 instead classified as competitive. 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}. This is the reason why we do not provide an analysis of the power and statistical significance for our results.[16] Nevertheless, it would be feasible to incorporate this analysis within the proposed approach provided that a reference distribution is available to test the observed distribution of J_{kt} against a null of no-coordination. For instance, if the dataset included 'comparable keywords' where no coalition is present, then one can consider testing how the distribution of J_{kt} compares between keywords with and without coalitions. This is not feasible in our data, however, since the search engine selected keywords for which a coalition is always present and, moreover, would require a careful evaluation of what keywords can be considered as comparable. We postpone the discussion of these features and of other extensions to the conclusions.

4.2 Step 2: Revenue Loss Quantification

The second step consists in 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.[17]

A key departure from the literature, however, is that not all valuations can be point identified under a bid coordination equilibrium. 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

[16] Such an analysis is straightforward for the Monte Carlo simulations in Figure 1. Indeed, the frequency with which the value of J_{kt} is equal or less than zero when calculated for a non coordinating bidder would give us the size of the test, while the same frequency when calculated for a coordinating bidder (other than the one ranked the highest among the coordinators) would give us the power of the test.

[17] Since intermediaries are many and heterogenous, assuming that all of them play a (specific type of) coordination equilibrium would have clearly been more restrictive relative to proceeding in steps 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.

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, suppose that we are in the context of 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.

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. Therefore, define $\tilde{b}_i = d_i e_i b_i$. Valuations are recovered by first estimating the perturbations d_i, and then by including these estimates in the bids' inversion procedure. Suppose that the data are assumed to be generated under a UC-RAE, then:

$$min_d \sum_{i>1}(d_i - 1)^2 \qquad \text{subject to:}$$

$$\begin{cases} \frac{\tilde{b}_i x^{i-1} - \tilde{b}_{i+1} x^i}{x^{i-1} - x^i} \geq \frac{\tilde{b}_{i+1} x^i - \tilde{b}_{i+2} x^{i+1}}{x^i - x^{i+1}}, & \text{if } i \notin \mathcal{C} \text{ or } i \in \{min(\mathcal{C})\}; \\ \tilde{b}_i x^{i-1} = \frac{x^{i-1} - x^i}{x^{i+1} - x^{i+2}}[\tilde{b}_{i+2} x^{i+1} - \tilde{b}_{i+3} x^{i+2}] + \gamma d_i e_i [x^{i-1} - x^i] + \tilde{b}_{i+1} x^i, & \text{if } i \in \mathcal{C} \backslash \{min(\mathcal{C})\}; \end{cases}$$

where γ is the minimum bid increment (5 cents in the data). In words, the solution to the quadratic program above finds the smallest d such that the UC-RAE restrictions are satisfied.[18] Provided with the estimated $\hat{d}_i$, 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 get 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]

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

[18] A similar formula applies in the case of the Eff-RAE. It can be obtained by replacing the constraint for $i \in \mathcal{C} \backslash \{min(\mathcal{C})\}$ in the formula above with the corresponding ones in DGP.

[19] The upper bound is the most relevant benchmark since the lower bound 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.

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. Columns two through five report respectively the lower bound of the counterfactual revenues under competitive bidding, the difference between the observed revenues and the lower bound, the upper bound of the counterfactual revenues and the difference between the upper bound and the observed revenues. In squared brackets we report 95% confidence interval of a keyword matched-pairs t-test: for each keyword, the mean difference in revenue between the observed and the counterfactual auctions is computed.

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

	Observed	Counterfactual Lower Bound	Difference $\Delta = Obs.\text{-}LowerB.$	Counterfactual Upper Bound	Difference $\Delta = UpperB.\text{-}Obs.$
Normalized Total Revenues	100	96.30	3.7	107.90	7.9
			[2.47; 4.92]		[5.32; 10.44]
Payments from Agency Advertisers	33.20	32.14	1.06	35.28	2.08
			[0.69; 1.42]		[1.49; 2.68]
Payments from Independent Advertisers	66.80	64.16	2.64	72.62	5.82
			[1.70; 3.58]		[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 five columns report the observed (normalized) revenues, the lower bound of the counterfactual revenues, the difference between the observed revenues and the lower bound, the upper bound of the counterfactual revenues and the difference between the upper bound and the observed revenues. 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 a range between 5.2 percent and 10.4 percent. The result on the lower bound, instead, suggests that coordinating bidders might be failing to extract full rents through lowering their bids. Interestingly, in both counterfactual scenarios 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.

5 Complete vs Incomplete Information GSP Auction

The information setup adopted to model the GSP auction influences how bidding coordination can be detected from the data. 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.[20]

A first case in point is 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 (Varian (2007)). 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.

In addition to these findings in the literature, we present next some novel results on the estimation of an incomplete information model with uncertainty over both the quality scores and valuations of other bidders. Once again, it will turn out that the complete information model produces estimates of valuations that are close to the ones of the incomplete information model, thereby lending further support to the approach undertaken in the previous sections.

5.1 Incomplete Information Model

We consider next an incomplete information setting based on Gomes and Sweeney (2014), with the only difference that we introduce quality scores, which are absent from the analysis in Gomes and Sweeney (2014). Specifically, we assume that per-click valuations v_i and quality scores are private information of the agents, with v_i and e_i are drawn independently of each other, and independently across bidders.[21] In addition to the notation introduced in Section 3, we let $\bar{v}_i = v_i \cdot e_i$ denote the quality.adjusted valuations, and let $F(\cdot)$ denote their (common known) distribution. Similarly, given a bid b_i and quality score e_i, recall that we defined the corresponding adjusted bids as $\bar{b}_i := e_i \cdot b_i$. As before, 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 it, which is $\frac{e_{k_{s+1}} b_{k_{s+1}}}{e_{k_s}} = \frac{\bar{b}_{k_{s+1}}}{e_{k_s}}$, where k_s denotes the bidder who occupies the s-th position. We consider the strictly monotonic Bayesian-Nash equilibrium of this game.

For the version of this model without quality scores, Gomes and Sweeney (2014) characterize the efficient Bayes-Nash equilibrium of the game, provide conditions for its existence, and show

[20] The complete information model of competitive bidding that we adopt also has its own theoretical appeal, as shown by EOS, Varian (2007) and (Milgrom and Mollner, 2018). Moreover, the model of individial 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 on the theoretical merits of the approach, see Decarolis, Goldmanis and Penta (2020).

[21] The assumption that advertisers know their own quality score is realistic since platforms usually reveal their own quality scores to the advertisers. For example, Google allows advertisers to check their own quality scores `https://support.google.com/google-ads/answer/2454010?hl=en`.

that there are no inefficient equilibria in symmetric strategies.[22] As it is easy to see from equation (4) below, our variation with quality scores is equivalent to the model of Gomes and Sweeney (2014), up to rescaling both the valuations and the resulting bids by the quality scores. Gomes and Sweeney (2014) existence results, therefore, extend to our setting. Theoretical analysis aside, however, this model has never been brought to estimation. The non-parametric identification and estimation results that we provide next are thus novel contributions on their own. In pursuing such results, we follow the first-order approach introduced by Guerre, Perrigne and Vuong (2000), which allows the point-wise identification of the (unobservable) bidders' valuations, with no need to solve for the equilibrium of the game.

For each quality-adjusted valuation $\bar{v}_i = v_i \cdot e_i$, we let $\sigma(\bar{v}_i) = \bar{b}_i$ denote the corresponding adjusted bid that maximizes i's expected payoff in the symmetric equilibrium. Since $\sigma(\bar{v}_i)$ is strictly monotonic in $\bar{v}_i$, it is invertible, and $\sigma^{-1}(\bar{b}_i) = \bar{v}_i$. Hence, under the assumption that σ is differentiable, for any possible adjusted bid $\bar{b}'_i$, we have $G(\bar{b}'_i) = Pr(\bar{b}_i \leq \bar{b}'_i) = Pr(\bar{v}_i \leq \sigma^{-1}(\bar{b}'_i)) = F(\sigma^{-1}(\bar{b}'_i)) = F(\bar{v}'_i)$, $g(\bar{b}'_i) = \frac{f(\bar{v}'_i)}{\sigma'(\bar{v}'_i)}$, where G denotes the distribution of adjusted bids, $g(\cdot)$ denotes their density, and $\bar{v}'_i = \sigma^{-1}(\bar{b}'_i)$. Given this, the expected equilibrium probability of winning the s-th slot, for a bidder with adjusted valuation $\bar{v}_i$, is:

$$z_s(\bar{v}_i) = \binom{n-1}{s-1}(1 - G(\sigma(\bar{v}_i)))^{s-1}G^{n-s}(\sigma(\bar{v}_i)) = \binom{n-1}{s-1}(1 - F(\bar{v}_i))^{s-1}F^{n-s}(\bar{v}_i).$$

The expected equilibrium payoff to bidder i, given his valuation v_i and quality weight e_i, when he submits a bid that induces an associated adjusted bid such that $\bar{b}'_i = \sigma(v'_i \cdot e_i) := \sigma(\bar{v}'_i)$ equals:

$$\begin{aligned} E[U_i|v_i, e_i, v'_i; F] &= \sum_{s=1}^{S} e_i x^s z_s(\bar{v}'_i)\left[v_i - \frac{1}{e_i}\int_0^{\bar{v}'_i} \sigma(y)\frac{(n-s)F^{n-s-1}(y)f(y)}{F^{n-s}(\bar{v}'_i)}dy\right] = \\ &= \sum_{s=1}^{S} x^s z_s(\bar{v}'_i)\left[\bar{v}_i - \int_0^{\bar{v}'_i} \sigma(y)\frac{(n-s)F^{n-s-1}(y)f(y)}{F^{n-s}(\bar{v}'_i)}dy\right]. \end{aligned} \tag{4}$$

Since in equilibrium such a bidder must find it optimal to submit a bid that induces $\bar{b}_i = \sigma(v_i \cdot e_i) = \sigma(\bar{v}_i)$, rather than one that induces some other $\bar{b}'_i = \sigma(v'_i \cdot e_i)$ for some $v'_i \neq v_i$, a necessary equilibrium condition is that $\frac{\partial E[U_i|v_i, e_i, v'_i; F]}{\partial v'_i}$, evaluated at the true valuation (i.e., for $v'_i = v_i$), is equal to zero. Rearranging such a first-order condition yields the following:[23]

$$v_i = \frac{\sum_{s=1}^{S} x^s \binom{n-1}{s-1}(n-s)(1-F(\bar{v}_i))^{s-2}\left[\sigma(\bar{v}_i)(1-F(\bar{v}_i))F^{n-s-1}(\bar{v}_i) - (s-1)\int_0^{\bar{v}_i}\sigma(y)F^{n-s-1}(y)f(y)dy\right]}{e_i \sum_{s=1}^{S} x^s \binom{n-1}{s-1}(1-F(\bar{v}_i))^{s-2}F^{n-s-1}(\bar{v}_i)\left[n - s - (n-1)F(\bar{v}_i)\right]}$$

[22] See Appendix A for the details.
[23] The derivation of this result is in the Appendix.

Rewriting the right-hand side in terms of observables, using $F(\bar{v}_i) = G(\bar{b}_i)$ and $g(\bar{b}_i) = \frac{f(\bar{v}_i)}{\sigma'(\bar{v}_i)}$, where we let $\bar{v}_i = \sigma^{-1}(\bar{b}_i)$, we identify the pseudo-valuations based on the bids, the quality scores, the distribution of adjusted bids, and the click-through rates:

$$v_i = \frac{\sum_{s=1}^{S} x^s \binom{n-1}{s-1}(n-s)(1-G(\bar{b}_i))^{s-2}\left[\bar{b}_i(1-G(\bar{b}_i))G^{n-s-1}(\bar{b}_i) - (s-1)\int_0^{\bar{b}_i} yG^{n-s-1}(y)g(y)dy\right]}{e_i \sum_{s=1}^{S} x^s \binom{n-1}{s-1}(1-G(\bar{b}_i))^{s-2}G^{n-s-1}(\bar{b}_i)\left[n-s-(n-1)G(\bar{b}_i)\right]} \tag{5}$$

To estimate the valuations we use a plug-in estimator based on equation (5). In particular, we estimate the density of the bids using kernel density estimator, and we estimate the distribution of bids using the empirical CDF.

Monte Carlo Simulation: Consider the GSP with four positions, five bidders, and click-through rates $(x^s)_{s=1,\ldots,5} = (20, 10, 5, 2, 0)$ (the $x^5 = 0$ is a useful notational device to indicate that there is no fifth slot). The quality-adjusted value per click of each bidder (multiplication of valuation and quality) is an independent draw from the uniform distribution: $\bar{v}_i \sim U[0, 6]$. We simulate 1,000 auctions: for each of these auctions, we draw five quality-adjusted valuations from the uniform distribution and calculate the corresponding adjusted bids based on the incomplete information model described above. [24] For the realizations of the quality-adjusted valuations $\bar{v} = (5, 4, 3, 2, 1)$, the equilibrium adjusted bids in incomplete information model are: $b_5 = 0.85$, $b_4 = 1.60$, $b_3 = 2.35$, $b_2 = 3.12$ and $b_1 = 3.94$. These are quite close to the EOS bids discussed earlier.

Given the equilibrium adjusted bids, we estimate the adjusted valuations both as if the bidders were playing the EOS equilibrium, as well as the pseudo-valuations based on equation (5).[25] We find that the average difference between the estimates obtained from the two models is very small:[26] the mean deviation, as a fraction of the estimated adjusted valuations, is 2% for the estimates based on the incomplete information model and 4% for those based on the complete information one. The mean adjusted valuation is 2.50 for the incomplete information model and 2.48 for the EOS model. Figure 3 plots the estimated adjusted valuations based on the two models (incomplete information on the x-axis, complete information on the y-axis). As can be seen from this picture, on average the adjusted valuations are very similar. The fact that the estimated adjusted valuations are not the same for a particular auction is to be expected: the reason is that the optimal bids in the EOS setting depend on the exact realizations of rivals' valuations, whereas in the incomplete information

[24] We use the numerical approximation of the equilibrium strategy. Details are in the Appendix.

[25] We consider adjusted valuations instead of valuations without loss of generality since quality scores are equal under both models. With a dataset where quality scores are observed, valuations are identified separately from the quality scores.

[26] We remind that the highest bidder's valuation cannot be identified within the EOS model. Hence, this comparison as well as all the ones that follow concern all valuations but the highest.

model they only depend on their distribution.

Figure 3: Comparison of the estimated adjusted valuations of the incomplete information model and of the estimated adjusted valuations as if the bidders were playing EOS

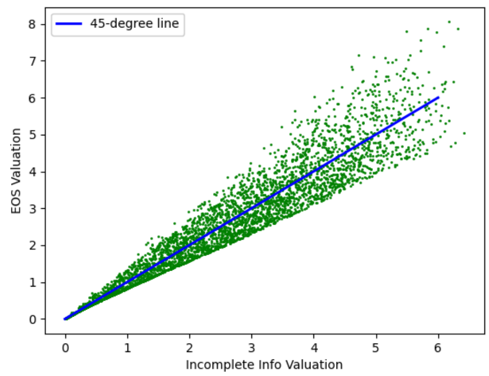

6 Conclusions

The analysis above illustrates the potential of using search auction data to detect bid coordination and to quantify its effects on revenues. We conclude with a discussion of 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. 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. However, 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.[27] 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 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.,

[27] Differences in serial correlation across keywords would also be one of the factors complicating the selection of a set of comparable keywords where no agency is present. As stated earlier, this might be a useful route if one would like to use such a group of keywords to devise a test comparing the distribution of the J_t statistic across keywords with and without potential coordinators.

criteria based on the mode, or on the concentration of the mass of the distribution, etc.). Our results are robust to several of these changes.[28] An interesting extension for future work would be to base the decision on whether the coalition's bids are too low to be competitive 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 calculation of the counterfactual revenues in the second step. This is the main reason why in this study we did not pursue this strategy.

Third, we do not consider the possibility that the quality scores might be endogenous. The model in which the quality score is part of the advertiser's decision is very interesting but, due to data limitations, it is outside the scope of this paper. It would be impossible given the data to distinguish the change in quality score due to the improvement of the ad made by the advertiser from the change due to other determinants of the quality scores that are not under the control of an advertiser. The main reason is that generally neither Google nor the other search engines reveal the exact algorithms that they use for determining the scores. Moreover, the search engine itself might use its ability to pivot quality scores in order to reduce its revenue losses due to bid coordination. Abou Nabout and Skiera (2012) show that the quality improvements might even lead to the decrease in advertiser's profits. Understanding the algorithms behind the quality scores assignments would thus be crucial to further consider these important effects.

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 coordinated strategies, 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 it normally is in standard single-item auctions, since with multiple items this strategy comes at the cost of forgoing the surplus that the agencies' clients may still make by obtaining a different slot, and by 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 in the SEMrush data

[28] 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 lower bound of this support is positive, as Eff-RAE if the support upper bound is negative and as UC-RAE if the interval includes zero. On the contrary, the classification changes substantially if we use the mean of J_t. In this case, the presence of outliers produces less interpretable results.

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.

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 limiting 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. That is an important direction of future work but we do not implement this empirically since we do not have data from different auction formats.

References

Abou Nabout, Nadia, and Bernd Skiera. 2012. "Return on Quality Improvements in Search Engine Marketing." *Journal of Interactive Marketing*, 26(3): 141–154.

Aghion, Philippe, Pol Antras, and Elhanan Helpman. 2007. "Negotiating free trade." *Journal of International Economics*, 73(1): 1–30.

Allouah, Amine, and Omar Besbes. 2017. "Auctions in the Online Display Advertising Chain: A Case for Independent Campaign Management." *Columbia Business School Research Paper No. 17-60.*

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.*

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

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.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-price Auctions." *Econometrica*, 68(3): 525–574.

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.

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

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

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.

Appendices

A Incomplete Information GSP Auction

A.1 Equilibrium exhistence

The discussion here is based on Gomes and Sweeney (2014).[29] The authors characterize the efficient Bayes-Nash equilibrium of the GSP without quality scores and provide a necessary and sufficient condition that guarantees existence of such an equilibrium.

Consider the generalized second-price auction (GSP) with n bidders, S positions ($n > S$) and click-through rates $x^1 \geq x^2 \geq ... \geq x^S$. If an efficient Bayes-Nash equilibrium exists for this auction, then its symmetric bidding strategy is

$$\sigma(v) = v - \phi(v) - \sum_{n=1}^{\infty} \int_0^v K_n(v,t)\phi(t)dt,$$

where

$$\phi(v) = \frac{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1}(s-1)(1-F(v))^{s-2} \mathbf{\frac{1}{n-s}} \int_0^v F^{n-s}(x)dx}{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1}(1-F(v))^{s-1} F^{n-s-1}(v)},$$

$$K_1(v,t) = \frac{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1}(s-1)(1-F(v))^{s-2} F^{n-s-1}(t) f(t)}{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1}(1-F(v))^{s-1} F^{n-s-1}(v)},$$

$$K_n(v,t) = \int_{\mathbf{t}}^{v} K_1(v,\epsilon) K_{n-1}(\epsilon,t) d\epsilon \ for \ n \geq 2,$$

and by assumption $\phi(v) \in L^2([0,\bar{v}])$ and $K_1(v,t) \in L^2([0,\bar{v}]^2)$.

If $\sigma(v)$ is strictly increasing, then an efficient Bayes-Nash equilibrium exists for the GSP. Otherwise, this auction does not admit an efficient equilibrium. If this auction does not possess an efficient equilibrium, then there exists no symmetric equilibrium.

A.2 First Order Condition

Since the GSP with quality scores incomplete information GSP with the quality scores is equivalent to incomplete information GSP with no quality scores, in which the both the valuations and resulting bids are scaled up by the quality scores, let's consider first the GSP in which all quality scores are asumed to equal one.

[29] Based on Gomes and Sweeney (2014) with two corrections in bold.

First, we rewrite the expected payoff to bidder i when his true valuation is v_i but he bids as if it was v_i' as:

$$E[U_i|v_i, v_i'; F] = \sum_{s=1}^{S} x^s z_s(v_i') \left[v_i - \int_0^{v_i'} \sigma(y) \frac{(n-s)F^{n-s-1}(y)f(y)}{F^{n-s}(v_i')} dy \right] =$$

$$= \sum_{s=1}^{S} x^s z_s(v_i') \left[v_i - \frac{n-s}{F^{n-s}(v_i')} \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \right] =$$

$$= \sum_{s=1}^{S} x^s \left[z_s(v_i') v_i - \binom{n-1}{s-1} (1 - F(v_i'))^{s-1} F^{n-s}(v_i') \frac{n-s}{F^{n-s}(v_i')} \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \right] =$$

$$= \sum_{s=1}^{S} x^s \left[v_i z_s(v_i') - \binom{n-1}{s-1} (n-s)(1 - F(v_i'))^{s-1} \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \right].$$

To prove the result stated in the proposition, we derive the first-order condition by differentiating the expected payoff with respect to v_i', substituting $v_i' = v_i$ and equating the resulting expression to zero.

Given that

$$z_s'(v_i') = \binom{n-1}{s-1}(n-s)(1 - F(v_i'))^{s-1} F^{n-s-1}(v_i') f(v_i') - \binom{n-1}{s-1}(s-1)(1 - F(v_i'))^{s-2} F^{n-s}(v_i') f(v_i') =$$

$$= \binom{n-1}{s-1}(1 - F(v_i'))^{s-2} F^{n-s-1}(v_i') f(v_i')[(n-s)(1 - F(v_i')) - (s-1)F(v_i')] =$$

$$= \binom{n-1}{s-1}(1 - F(v_i'))^{s-2} F^{n-s-1}(v_i') f(v_i')[n - s - (n-1)F(v_i')],$$

we get the following equation for the valuation of player i:

$$\sum_{s=1}^{S} x^s \left[v_i z_s'(v_i') - \left(\binom{n-1}{s-1}(n-s)(1 - F(v_i'))^{s-1} \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \right)'_{v_i'} \right] =$$

$$= \sum_{s=1}^{S} x^s \left[v_i \binom{n-1}{s-1}(1 - F(v_i'))^{s-2} F^{n-s-1}(v_i') f(v_i')[n - s - (n-1)F(v_i')] - \right.$$

$$- \left(\binom{n-1}{s-1}(n-s)(1 - F(v_i'))^{s-1} \sigma'(v_i') s(v_i') G^{n-s-1}(\sigma(v_i')) g(\sigma(v_i')) - \right.$$

$$\left. \left. - \binom{n-1}{s-1}(n-s)(s-1)(1 - F(v_i'))^{s-2} f(v_i') \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \right) \right] =$$

$$\left(\text{since } g(\sigma(v_i'))\sigma'(v_i') = f(v_i') \right)$$

$$= \sum_{s=1}^{S} x^s \left[v_i \binom{n-1}{s-1} (1-F(v_i'))^{s-2} F^{n-s-1}(v_i') f(v_i') [n-s-(n-1)F(v_i')] - \right.$$
$$- \binom{n-1}{s-1} (n-s)(1-F(v_i'))^{s-2} \Bigg((1-F(v_i'))\sigma(v_i') G^{n-s-1}(\sigma(v_i')) f(v_i') -$$
$$\left. -(s-1) f(v_i') \int_0^{\sigma(v_i')} y G^{n-s-1}(y) g(y) dy \Bigg) \right] = 0 \text{ when } v_i' = v_i \Rightarrow$$

$f(v_i)$ cancels out and

$$v_i = \frac{\sum_{s=1}^{S} x^s \binom{n-1}{s-1} (n-s)(1-F(v_i))^{s-2} \left[(1-F(v_i))\sigma(v_i) G^{n-s-1}(\sigma(v_i)) - (s-1) \int_0^{\sigma(v_i)} y G^{n-s-1}(y) g(y) dy \right]}{\sum_{s=1}^{S} x^s \binom{n-1}{s-1} (1-F(v_i))^{s-2} F^{n-s-1}(v_i) \Big[n-s-(n-1)F(v_i) \Big]} =$$
$$= \frac{\sum_{s=1}^{S} x^s \binom{n-1}{s-1} (n-s)(1-F(v_i))^{s-2} \left[\sigma(v_i)(1-F(v_i)) F^{n-s-1}(v_i) - (s-1) \int_0^{v_i} \sigma(y) F^{n-s-1}(y) f(y) dy \right]}{\sum_{s=1}^{S} x^s \binom{n-1}{s-1} (1-F(v_i))^{s-2} F^{n-s-1}(v_i) \Big[n-s-(n-1)F(v_i) \Big]}.$$

Now, by using the $\bar{v}_i$ instead of just v_i, we obtain the first order condition of the model with quality scores.

A.3 Monte Carlo simulations

Since for this case the analytical solution based on Volterra equation cannot be easily obtained, we use the following numerical approximation based on Gomes and Sweeney (2014):

$$b = \sigma(v) = v - \phi(v) - \sum_{n=1}^{m} \int_0^v K_n(v,t)\phi(t) dt,$$

where

$$\phi(v) = \frac{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1} (s-1)(1-F(v))^{s-2} \frac{\mathbf{1}}{\mathbf{n-s}} \int_0^v F^{n-s}(x) dx}{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1} (1-F(v))^{s-1} F^{n-s-1}(v)},$$
$$K_1(v,t) = \frac{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1} (s-1)(1-F(v))^{s-2} F^{n-s-1}(t) f(t)}{\sum_{s=1}^{s=S} x^s \binom{n-2}{s-1} (1-F(v))^{s-1} F^{n-s-1}(v)},$$
$$K_n(v,t) = \int_{\mathbf{t}}^{v} K_1(v,\epsilon) K_{n-1}(\epsilon,t) d\epsilon \ for \ n \geq 2, \ m = 3.$$

Figure (4) shows the estimated CDFs of the valuations from the nonparametric estimation of the incomplete information game and the true CDF.

Figure 4: True and Estimated CDF of the adjusted valuations in Monte Carlo simulation of the GSP of incomplete information with 5 bidders, 4 slots and the uniform distribution of adjusted valuations.