

Marketing Agencies and Collusive Bidding in Sponsored Search Auctions

Francesco Decarolis, Maris Goldmanis, Antonio Penta*

December 29, 2016

Abstract

The transition of the advertisement market from traditional media to the internet has induced a proliferation of marketing agencies specialized in bidding in the auctions used to sell advertisement space on the web. We analyze how collusive bidding can emerge from bid delegation to a common marketing agency and how this undermines both revenues and efficiency of both the generalized second price auction (GSP, used by Google and Microsoft-Yahoo!) and the VCG mechanism (used by Facebook). Surprisingly, we find that both in terms of revenues and efficiency the GSP auction is outperformed by the VCG, despite the latter mechanism is known to perform poorly under collusive bidding. We also propose a criterion to detect collusion in the data.

JEL: C72, D44, L81.

Keywords: Online Advertising, Internet Auctions, Collusion.

*Decarolis, Department of Economics and Hariri Institute, Boston University, *fdc@bu.edu*. Goldmanis, Department of Economics, Royal Holloway, University of London, *Maris.Goldmanis@rhul.ac.uk*. Penta, Department of Economics, University of Wisconsin-Madison, *apenta@ssc.wisc.edu*. We are grateful for the comments received from Susan Athey, Nicholas Bloom, Yeon-Koo Che, Ken Hendricks, Jon Levin, Massimo Motta, Marco Ottaviani, 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, HEC-Lausanne, Microsoft Research Lab, Stanford University, University of British Columbia, University of California Davis, University of Mannheim, University of Toronto, University of Wisconsin Madison and at the conferences of the EIEF-University of Bologna-Bocconi joint IO workshop, Society for Economic Dynamics and Toulouse Network for Information Technology Meeting where earlier versions of this paper were presented.

1 Introduction

In two influential papers, Edelman, Ostrovsky and Schwarz (2007, EOS hereafter) and Varian (2007) pioneered the study of the Generalized Second Price (GSP) auction, the mechanism used to allocate advertisement space on the results web page of search engines like Google, Microsoft Bing and Yahoo!.¹ These auctions represent one of the fastest growing and most economically relevant forms of online advertisement, with an annual growth of 10% and a total value of 50 billion dollars in 2013 (PwC (2015), see also Blake, Nosko and Tadelis (2015)). A recent trend in this industry, however, has the potential to alter the functioning of these auctions – and, hence, the profits in this industry – in ways that are not accounted for by the existing models. In particular, an increasing number of bidders are delegating their bidding campaigns to specialized Search Engine Marketing Agencies (SEMAs).² As a result, SEMAs often bid in the same auction on behalf of different advertisers. But this clearly changes the strategic interaction, as SEMAs have the opportunity to lower their payments by coordinating the bids of their clients.

In this paper we explore the impact of bidding delegation to a common SEMA on the performance of the GSP auction. Our theoretical analysis uncovers a striking fragility of this mechanism to the possibility of bid coordination. This is underscored by our finding that the GSP auction is outperformed, both in terms of revenues and efficiency, by the Vickrey-Clarke-Groves (VCG) mechanism, recently adopted by Facebook. The superiority of the VCG relative to the GSP is both surprising, given that the VCG is well known in the literature to perform poorly under coordinated bidding, and noteworthy, given the sheer size of transactions occurring under the GSP.

Studying bidding coordination in the GSP auction presents a number of difficulties, which

¹Varian (2007) and EOS study a complete information environment. More recent work by Athey and Nekipelov (2012) maintains common knowledge of valuations but introduces uncertainty on the set of bidders and quality scores. Gomes and Sweeney (2014) instead study an independent private values model. In related work on the ad exchanges auctions, Balseiro, Besbes and Weintraub (2015) study the dynamic interaction among advertisers under limited budgets constraints.

²A survey by the Association of National Advertisers (ANA) among 74 large U.S. advertisers indicates that about 77% of the respondents fully outsource their search engine marketing activities (and 16% partially outsource them) to specialized agencies, see ANA (2011). Analogously, a different survey of 325 mid-size advertisers by Econsultancy (EC) reveals that the fraction of companies not performing their paid-search marketing in house increased from 53% to 62% between 2010 and 2011, see EC (2011).

are only partly due to the inherent complexity of the mechanism. An insightful analysis of the problem thus requires a careful balance between tractability and realism of the assumptions.

For instance, since SEMAs in this market operate side by side with independent advertisers, it is important to have a model in which coordinated and competitive bidding coexist. But the problem of ‘partial cartels’ is acknowledged as a major difficulty in the literature (e.g., Hendricks, Porter and Tan (2008)).³ To address this problem, our model combines elements of cooperative and non-cooperative game theory. We thus introduce a general notion of ‘Recursively-Stable Agency Equilibrium’ (RAE) which involves both equilibrium and stability restrictions, and which provides a unified framework to study the impact of SEMAs under different mechanisms. Second, it is well-known that strategic behavior in the GSP auction is complex and gives rise to a plethora of equilibria (Borgers, Cox, Pesendorfer and Petricek (2013)). Introducing a tractable and insightful refinement has been a key contribution of EOS and Varian (2007). But their refinements are not defined in a context in which some advertisers coordinate their bids. Thus, a second challenge we face is to develop a refinement for the model with SEMA, which ensures both tractability as well as clear economic insights.

To achieve these goals, we modify EOS and Varian’s baseline model by introducing a SEMA, which we model as a player that chooses the bids of its clients in order to maximize the total surplus. Bidders that do not belong to the SEMA are referred to as ‘independents’, and place bids in order to maximize their own profits. To ensure a meaningful comparison with EOS competitive benchmark, and to avoid the severe multiplicity of equilibria in the GSP auction, we introduce a refinement of bidders’ best responses that distills the individual-level underpinnings of EOS ‘lowest-revenue envy-free’ equilibrium, and assume that independents place their bids accordingly. This device enables us to maintain the logic of EOS refinement for the independent bidders, even if EOS equilibrium is not defined in the game with SEMA. The SEMA in turn makes a proposal of a certain profile of bids to its clients. The proposal is implemented if it is ‘recursively stable’ in the sense that, anticipating the bidding strategies

³The literature on ‘bidding rings’, for instance, has either considered mechanisms in which non-cooperative behavior is straightforward (e.g., second price auctions with private values, as in Mailath and Zemski (1991)), or assumes that the coalition includes all bidders in the auction (as in the first price auctions of McAfee and McMillan (1992) and Hendricks et al. (2008)). Other mechanisms for collusion have been considered, for instance, by Harrington and Skrzypacz (2007, 2011).

of others, and taking into account the possible unraveling of the rest of the coalition, no client has an incentive to abandon the SEMA and bid as an independent. Hence, in our model, the outside options of the members of a coalition are equilibrium objects themselves, and implicitly incorporate the restrictions entailed by the underlying coalition formation game.

Within this general framework, we consider different models of coordinated bidding, in which the SEMA operates under different constraints. We first assume that the agency is constrained to placing bids which cannot be detected as ‘coordinated’ by an external observer. This is a useful working hypothesis, which also has obvious intrinsic interest. Under this assumption, we show that the resulting allocation in the GSP with SEMA is efficient and the revenues are the same as those that would be generated if the same coalition structure was bidding in a VCG auction without any constraint. We then relax this ‘undetectability constraint’, and show that in this case the search engine revenues in the GSP auction are never higher, and are in fact typically lower, than those obtained in the VCG mechanism with coordination. Furthermore, once the ‘undetectability constraint’ is lifted, efficiency in the allocation of bidders to slots is no longer guaranteed by the GSP mechanism. Since the VCG is famously regarded as a poor mechanism under coordinated bidding, finding that it outperforms the GSP both in terms of revenues and allocation is remarkably negative for the GSP. We also show that these insights persist even in the presence of multiple agencies, and hence competition between agencies is not enough to avoid these problems.

This fragility of the GSP auction is due to the complex equilibrium effects it induces. In particular, by manipulating the bids of its members, in equilibrium the SEMA also affects the bids of the independents. The SEMA therefore has both a *direct* and an *indirect effect* in the GSP auction, and hence even a small coalition may have a large impact on total revenues. Depending on the structure of the SEMA, the indirect effect may be first order. As we will explain, our analysis uncovers that this is especially the case if the SEMA includes members which occupy low or adjacent positions in the ranking of valuations.

In the final part of the paper, we supplement our theoretical analysis showing how it can be applied to search auctions data. We use simulated data to show how our theoretical results can be used to detect potentially coordinated behavior, and to choose over alternative models of coordinated bidding. We also argue how this detection tool can be used to quantify

the revenue losses that collusion imposes on the search engine.

The main implication of this study is that, with SEMAs, the revenue sharing and allocative properties of both GSP and VCG will change. This is obviously interesting from a market design perspective. Although the design of ad auctions has received considerable attention in the literature (e.g., Edelman and Schwarz (2010) and Celis et al. (2015)), this is the first study to point to the role of agencies and, importantly, to develop a methodology capable of accommodating the joint presence of competitive and collusive bidding.

From an applied perspective, the growing relevance of the risk of collusion through SEMAs is underscored by the increased concentration in the ad agency market. Although several hundreds of agencies operate in the US, the majority of them currently belongs to one of the seven large ad agency networks which, in turn, all have a single “agency trading desk” to conduct all bidding activities.⁴ These desks are the centralized units within a network, that collect the data and design the algorithms required to optimize the purchase of online advertisement space for “programmatic” (i.e., algorithmic or “biddable”) media like Google, Bing, Twitter, iAd, and Facebook. In a complementary study, Decarolis, Goldmanis and Penta (2016), we empirically analyze ad auctions confirming both that bidding delegation to agencies is pervasive and that a common agency placing bids in a single auction on behalf of different advertisers occurs in an economically relevant amount of auctions. While that study seeks to look at multiple aspects of marketing agencies that go beyond collusive bidding, such as their role in attracting new customers and improving ad quality, the model of collusive bidding developed here is the crucial building block to understand the incentives that a SEMA faces whenever two or more of its clients are interested in the same keyword.

Finally, our results are also relevant from an antitrust perspective. Our characterization of the agency behavior is analogous to that of buying consortia, which have been sanctioned in the past.⁵ Nevertheless, the specificities of the market suggest a more nuanced view of the harm to consumers. We discuss this point and other policy implications in the conclusions.

⁴The networks are: Aegis-Dentsu, Publicis Groupe, IPG, Omnicom Group, WPP/Group M, Havas, MDC. Their corresponding trading desks are: Amnet, Vivaki, Cadreon, Accuen, Xaxis, Affiperf and Varick Media.

⁵See, for instance, the case of the tobacco manufacturers consortium buying in the tobacco leaves auctions, *United States v. American Tobacco Company*, 221 U.S. 106 (1911).

2 The GSP auction

In this section we introduce the GSP auction and the necessary notation. The only original result in this section is Lemma 1, which shows that EOS *lowest-revenue envy-free* (LREF) equilibrium – originally defined as a refinement of the Nash equilibrium correspondence – can be equivalently defined as the fixed point of a refinement of individuals’ best responses. This result will play an important role in the next sections. That is because, by distilling the individual underpinnings of EOS refinement, it enables us to extend EOS approach to the analysis of the GSP auction with SEMA, in which LREF equilibrium is not defined.

Consider the problem of assigning agents $i \in I = \{1, \dots, n\}$ to slots $s = 1, \dots, S$, where $n \geq S$. In our case, agents are advertisers, and slots are positions for ads in the page for a given keyword. Slot $s = 1$ corresponds to the highest position, and so on until $s = S$, which is the slot in the lowest position. For each s , we let x^s denote the ‘click-through-rate’ (CTR) of slot s , that is the number of clicks that an ad in position s is expected to receive, and assume that $x^1 > x^2 > \dots > x^S > 0$. We also let $x^t = 0$ for all $t > S$. Finally, we let v_i denote the per-click-valuation of advertiser i , and we label advertisers so that $v_1 > v_2 > \dots > v_n$.

In the GSP auction, each advertiser submits a bid $b_i \in \mathbb{R}_+$. The advertiser who submits the highest bid obtains the first slot and pays a price equal to the second highest bid every time his ad is clicked; the advertiser with the second highest bid obtains the second slot and pays a price-per-click equal to the third highest bid, and so on. We denote bid profiles by $b = (b_i)_{i=1, \dots, n}$ and $b_{-i} = (b_j)_{j \neq i}$. For any profile b , we let $\rho(i; b)$ denote the rank of i ’s bid in b (ties are broken according to bidders’ labels).⁶ When b is clear from the context, we omit it and write simply $\rho(i)$. For any $t = 1, \dots, n$ and b or b_{-i} , we let b^t and b_{-i}^t denote the t -highest component of the vectors b and b_{-i} , respectively.

The rules of the auction are formalized as follows. For any b , if $\rho(i) \leq S$ bidder i obtains position $\rho(i)$ at price-per-click $p^{\rho(i)} = b^{\rho(i)+1}$. If $\rho(i) > S$, bidder i obtains no position.⁷ We maintain throughout that advertisers’ preferences, the CTRs and the rules of the auction are

⁶Formally, $\rho(i; b) := |\{j : b_j > b_i\} \cup \{j : b_j = b_i \text{ and } j < i\}| + 1$. This tie-breaking rule is convenient for the analysis of coordinated bidding. It can be relaxed at the cost of added technicalities (see footnote 15).

⁷In reality, bidders allocation to slots is determined adjusting advertisers’ bids by some ‘quality scores’. To avoid unnecessary complications, we only introduce quality scores in section 5 (see also Varian (2007)).

common knowledge. We thus model the GSP auction as a game $\mathcal{G}(v) = (A_i, u_i^{\mathcal{G}})_{i=1, \dots, n}$ where $A_i = \mathbb{R}_+$ denotes the set of actions of player i (his bids), and payoff functions are such that, for every i and every $b \in \mathbb{R}_+^n$, $u_i^{\mathcal{G}}(b) = (v_i - b^{\rho(i)+1}) x^{\rho(i)}$. Although it may seem unrealistic, EOS complete information assumption has been shown to be an effective modeling proxy (e.g., Athey and Nekipelov (2012), Che, Choi and Kim (2013) and Varian (2007)).

Note that any generic profile $b_{-i} = (b_j)_{j \neq i}$ partitions the space of i 's bids into $S + 1$ intervals. The only payoff relevant component of i 's choice is in which of these intervals he should place his own bid: any two bids placed in the same interval would grant bidder i the same position at the same per-click price (equal to the highest bid placed below b_i). So, for each $b_{-i} \in \mathbb{R}_+^{n-1}$, let $\pi_i(b_{-i})$ denote i 's favorite position, given b_{-i} .⁸ Then, i 's best-response to b_{-i} is the interval $BR_i(b_{-i}) = (b_{-i}^{\pi_i(b_{-i})}, b_{-i}^{\pi_i(b_{-i})-1})$. This defines the best-response correspondence $BR_i : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$, whose fixed points are the set of (pure) Nash equilibria: $\mathcal{EG}^0(v) := \{b \in \mathbb{R}_+^n : b_i \in BR_i(b_{-i}) \forall i \in I\}$.

It is well-known that the GSP auction admits a multiplicity of equilibria (Borgers et al. (2013)). For this reason, EOS introduced a refinement of the set $\mathcal{EG}^0(v)$, the *LREF equilibrium*. As anticipated above, we consider instead a refinement of individuals' best response correspondence: for any $b_{-i} \in \mathbb{R}_+^{n-1}$, let

$$BR_i^*(b_{-i}) = \left\{ b_i^* \in BR_i(b_{-i}) : (v_i - b_{-i}^{\pi_i(b_{-i})}) x^{\pi_i(b_{-i})} = (v_i - b_i^*) x^{\pi_i(b_{-i})-1} \right\}. \quad (1)$$

In words, of the many $b_i \in BR_i(b_{-i})$ that would grant player i his favorite position $\pi_i(b_{-i})$, he chooses the bid b_i^* that makes him indifferent between occupying the current position and climbing up one position paying a price equal to b_i^* . The set of fixed points of the BR_i^* correspondence are denoted as $\mathcal{EG}(v) = \{b \in \mathbb{R}_+^n : b_i \in BR_i^*(b_{-i}) \text{ for all } i \in I\}$.

Lemma 1. *For any $b \in \mathcal{EG}(v)$, $b_1 > b_2$, $b_i = v_i$ for all $i > S$, and for all $i = 2, \dots, S$,*

$$b_i = v_i - \frac{x^i}{x^{i-1}} (v_i - b_{i+1}). \quad (2)$$

Hence, the fixed points of the BR^ correspondence coincide with EOS' LREF equilibria.*

⁸ Allowing ties in individuals' bids or non-generic indifferences complicates the notation, without affecting the results and the main insights. We thus ignore these issues here and leave the details to Appendix ??.

In section 4 we will assume that independents in the GSP auction bid according to BR_i^* , both with and without the agency. Since, by Lemma 1, this is precisely the same assumption on individuals' behavior that underlies EOS' analysis, our approach ensures a meaningful comparison with the competitive benchmark. Lemma 1 obviously implies that our formulation inherits the many theoretical arguments in support of EOS refinement (e.g. EOS, Edelman and Schwarz (2010), Milgrom and Mollner (2014)). But it is also important to stress that, independent of equilibrium restrictions, this individual-level refinement is particularly compelling because it conforms to the tutorials on how to bid in these auctions provided by the search engines.⁹ This formulation is thus preferable from both a conceptual and a practical viewpoint.

For later reference, it is useful to rearrange (2) to obtain the following characterization of the LREF equilibrium testable implications (cf. EOS and Varian (2007)):

Corollary 1. *For any $b \in EG(v)$, for all $i = 2, \dots, S$:*

$$\underbrace{\frac{b_i x^{i-1} - b_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i} > \underbrace{\frac{b_{i+1} x^i - b_{i+2} x^{i+1}}{x^i - x^{i+1}}}_{=v_{i+1}} \quad (3)$$

The main alternative to the GSP auction is represented by the VCG: it is used by Facebook, as well as by Google for a subset of its auctions, and it also represents the standard benchmark for the GSP auction in the theoretical literature (e.g., EOS, Varian (2007) and Athey and Nekipelov (2012)). Given valuations v , the VCG mechanism is formally defined as a game $\mathcal{V}(v) = (A_i, u_i^\mathcal{V})_{i=1, \dots, n}$, where $A_i = \mathbb{R}_+$ and $u_i^\mathcal{V}(b) = v_i x^{\rho(i)} - \sum_{t=\rho(i)+1}^{S+1} b^t (x^{t-1} - x^t)$ for each i and $b \in \mathbb{R}_+^n$.

It is well-known that bidding $b_i = v_i$ is a dominant strategy in this game. In the resulting equilibrium, advertisers are efficiently assigned to positions. Furthermore, EOS (Theorem 1) showed that the position and payment of each advertiser in the dominant strategy equilibrium of the VCG are the same as in the LREF equilibrium of the GSP auction.

The next example will be used repeatedly throughout the paper to illustrate the relative performance of the GSP and VCG mechanisms:

⁹See, for instance, the Google Ad Word tutorial in which Hal Varian teaches how to maximize profits by following this bidding strategy: <http://www.youtube.com/watch?v=jRx7AMb6rZ0>.

Example 1. Consider an auction with four slots and five bidders, with the following valuations: $v = (5, 4, 3, 2, 1)$. The CTRs for the five positions are the following: $x = (20, 10, 5, 2, 0)$. In the VCG mechanism, bids are $b_i = v_i$ for every i , which induces total expected revenues of 96. Bids in the LREF of the GSP auction instead are as follows: $b_5 = 1$, $b_4 = 1.6$, $b_3 = 2.3$ and $b_2 = 3.15$. The highest bid b_1 is not uniquely determined, but it does not affect the revenues, which in this equilibrium are exactly the same as in the VCG mechanism: 96. Clearly, also the allocation is the same in the two mechanisms, and efficient.

3 Search Engine Marketing Agencies

Our analysis of the Search Engine Marketing Agencies (SEMAs) focuses on their opportunity to coordinate the bids of different advertisers. We thus borrow the language of cooperative game theory and refer to the clients of the agency as ‘members of a coalition’ and to the remaining bidders as ‘independents’. In this Section we focus on environments with a single SEMA. We will extend the analysis to multiple SEMAs in Section 4.4.

Modeling coordinated bidding, it may seem natural to consider standard solution concepts such as Strong Nash or Coalition Proof Equilibrium. Unfortunately, these concepts have no bite in the GSP auction, as it can be shown that the LREF equilibrium satisfies both refinements. As already shown by EOS in the competitive setting, resorting to non-standard solution concepts is a more promising route to generate meaningful insights on the elusive GSP auction. We thus model the SEMA as a player that makes proposals of binding agreements to its members, subject to certain stability constraints. The independents then play the game which ensues from taking the bids of the agency as given. We assume that the agency seeks to maximize the coalition surplus, but is constrained to choosing proposals that are *stable* in two senses: first, they are consistent with the independents’ equilibrium behavior; second, no individual member of the coalition has an incentive to abandon it and play as an independent. We also assume that, when considering such deviations, coalition members are *farsighted* in the sense that they anticipate the impact of their deviation on both the independents and the remaining members of the coalition (cf. Ray and Vohra (1997, 2014)). The constraint for a coalition of size C thus depends on the solutions to the

problems of all the subcoalitions of size $C - 1$. Therefore, the solution concept for the game with the agency will be defined recursively.

Besides overcoming the afore-mentioned limitations of standard solution concepts, our approach also addresses some important questions in the theoretical and applied literature, such as: (i) provide a tractable model of *partial cartels*, a well-known difficulty in the literature on bidding rings;¹⁰ (ii) deliver sharp results on the impact of coordinated bidding on the GSP auction, vis-à-vis the lack of bite of standard solution concepts; (iii) provide a model of coordinated bidding that can be applied to different mechanisms; (iv) bridge the theoretical results to the data, by generating easy-to-apply testable predictions to detect coordination and select between different models of bidding coordination.

Another obvious alternative to our approach would be to model bidders' choice to join the SEMA explicitly. This would also be useful from an empirical viewpoint, as it would generate extra restrictions to further identify bidders' valuations. But once again, the structure of the GSP auction raises non trivial challenges. First, it is easy to see that without an exogenous cost of joining the agency, the only outcome of a standard coalition formation game would result in a single agency consisting of the grand-coalition of players. Thus, the 'obvious' extension of the model would not be capable of explaining the lack of grand coalitions in the data. Hence, at a minimum, some cost of joining the coalition should be introduced. Clearly, there are many possible ways in which participation costs could be modeled (e.g., costs associated to information leakage, management practices, agency contracts, etc.). But given the still incomplete understanding of SEMAs in online advertising, it is not obvious which should be preferable. Regardless of the specific modeling choices, however, the cost of joining the SEMA would ultimately have to be traded-off against the benefit, which in turn presumes solving for the equilibrium of the auction *given* the resulting coalition structure. Our work can thus be seen as a necessary first step in developing a full-blown model of

¹⁰That literature typically focuses on the incentives that members of the coalition have to share their private information, which is not an issue in our setting. But more importantly, allowing for the co-presence of coordinated and independent bidding is a well-known difficulty in that literature, which either considers mechanisms in which non-cooperative behavior is straightforward (e.g., second price auction with private values, as in Graham and Marshall (1987) and Mailath and Zemski (1991)), or assumes that the coalition includes all bidders in the auction (as in the first price auctions of McAfee and McMillan (1992) and Hendricks et al. (2008)). Our equilibrium notion enables us to combine cooperative and non-cooperative interaction in general mechanisms, even if non-cooperative behavior is complex, as in the GSP auction.

endogenous agency participation, which also clarifies that participation costs must be taken into account, to make sense of the empirical evidence.

In the next section we introduce a general definition of the ‘Recursively Stable Agency Equilibrium’ (RAE), which allows for arbitrary underlying mechanisms. This is useful in that it provides a general framework to analyze the impact of SEMAs under different mechanisms. We then specialize the analysis to the GSP and VCG mechanisms in Section 4.

3.1 The Recursively Stable Agency Equilibrium

Let $G(v) = (A_i, u_i^G)_{i=1, \dots, n}$ denote the baseline game (without a coalition) generated by the underlying mechanism (e.g., the GSP ($G = \mathcal{G}$) or the VCG ($G = \mathcal{V}$) mechanism), given the profile of valuations $v = (v_i)_{i \in I}$. For any $\mathcal{C} \subseteq I$ with $|\mathcal{C}| \geq 2$, we let \mathcal{C} denote the agency, and we refer to advertisers $i \in \mathcal{C}$ as ‘members of the coalition’ and to $i \in I \setminus \mathcal{C}$ as ‘independents’. The coalition chooses a vector of bids $b_{\mathcal{C}} = (b_j)_{j \in \mathcal{C}} \in \times_{j \in \mathcal{C}} A_j$. Given $b_{\mathcal{C}}$, the independents $i \in I \setminus \mathcal{C}$ simultaneously choose bids $b_i \in A_i$. We let $b_{-\mathcal{C}} := (b_j)_{j \in I \setminus \mathcal{C}}$ and $A_{-\mathcal{C}} := \times_{j \in I \setminus \mathcal{C}} A_j$. Finally, given profiles b or $b_{-\mathcal{C}}$, we let $b_{-i, -\mathcal{C}}$ denote the subprofile of bids of all independents other than i (that is, $b_{-i, -\mathcal{C}} := (b_j)_{j \in I \setminus \mathcal{C}: j \neq i}$).

We assume that the agency maximizes the sum of the payoffs of its members,¹¹ denoted by $u_{\mathcal{C}}(b) := \sum_{i \in \mathcal{C}} u_i(b)$, under three constraints. Two of these constraints are stability restrictions: one for the independents, and one for the members of the coalition. The third constraint, which we formalize as a set $R_{\mathcal{C}} \subseteq A_{\mathcal{C}}$, allows us to accommodate the possibility that the agency may exogenously discard certain bids (this restriction is vacuous if $R_{\mathcal{C}} = A_{\mathcal{C}}$). For instance, in section 4.2.1 we will consider the case of an agency whose primary concern is not being identified as inducing coordinated bids. In that case, $R_{\mathcal{C}}$ would be comprised of only those bids that are ‘undetectable’ to an external observer as coordinated. We denote the collection of exogenous restrictions for all possible coalitions as $\mathcal{R} = (R_{\mathcal{C}})_{\mathcal{C} \subseteq I: |\mathcal{C}| \geq 2}$.

Stability-1: The first stability restriction on the agency’s proposals is that they are stable

¹¹This is a simplifying assumption, which can be justified in a number of ways. From a theoretical viewpoint, our environment satisfies the informational assumptions of Bernheim and Whinston (1985, 1986). Hence, as long as the SEMA is risk-neutral, this particular objective function may be the result of an underlying common agency problem. More relevant from an empirical viewpoint, the agency contracts most commonly used in this industry specify a lump-sum fee per advertiser and per campaign. Thus, the SEMA’s ability to generate surplus for its clients is an important determinant of its long run profitability.

with respect to the independents. For any $i \in I \setminus \mathcal{C}$, let $BR_i^G : A_{-i} \rightrightarrows A_i$ denote some refinement of i 's best response correspondence in $G(v)$ (e.g., BR_i^* in $\mathcal{G}(v)$ or weak dominance in $\mathcal{V}(v)$). Define the *independents' equilibrium correspondence* $BR_{-\mathcal{C}}^G : A_{\mathcal{C}} \rightrightarrows A_{-\mathcal{C}}$ as

$$BR_{-\mathcal{C}}^G(b_{\mathcal{C}}) = \{b_{-\mathcal{C}} \in A_{-\mathcal{C}} : \forall j \in I \setminus \mathcal{C}, b_j \in BR_j^G(b_{\mathcal{C}}, b_{-j, -\mathcal{C}})\}. \quad (4)$$

If the agency proposes a profile $b_{\mathcal{C}}$ that is not consistent with the equilibrium behavior of the independents (as specified by $BR_{-\mathcal{C}}^G$), then that proposal does not induce a *stable agreement*. We thus incorporate this stability constraint into the decision problem of the agency, and assume that the agency can only choose bid profiles from the set

$$S_{\mathcal{C}} = \{b_{\mathcal{C}} \in A_{\mathcal{C}} : \exists b_{-\mathcal{C}} \text{ s.t. } b_{-\mathcal{C}} \in BR_{-\mathcal{C}}^G(b_{\mathcal{C}})\}. \quad (5)$$

Clearly, the strength of this constraint in general depends on the underlying game $G(v)$ and on the particular correspondence $BR_{-\mathcal{C}}^G$ that is chosen to model the independents' behavior. This restriction is conceptually important, and needed to develop a general framework for the analysis of coordinated bidding in arbitrary mechanisms. Nonetheless, the restriction plays no role in our results for the GSP and VCG mechanisms, because (5) will be either vacuous (Theorem 1) or a redundant constraint (Theorems 2 and 3).

Stability-2: When choosing bids $b_{\mathcal{C}}$, the agency forms conjectures about how its bids would affect the bids of the independents. We let $\beta : S_{\mathcal{C}} \rightarrow A_{-\mathcal{C}}$ represent such conjectures of the agency. For any profile $b_{\mathcal{C}} \in S_{\mathcal{C}}$, $\beta(b_{\mathcal{C}})$ denotes the agency's belief about the independents' behavior, if she chooses profile $b_{\mathcal{C}}$. It will be useful to define the set of conjectures β that are consistent with the independents playing an equilibrium:

$$B^* = \{\beta \in A_{-\mathcal{C}}^{S_{\mathcal{C}}} : \beta(b_{\mathcal{C}}) \in BR_{-\mathcal{C}}^G \text{ for all } b_{\mathcal{C}} \in S_{\mathcal{C}}\}. \quad (6)$$

The second condition for stability requires that, given the conjectures β , the members of the coalition have no incentives to leave the agency and start acting as independents. Hence, the outside option for coalition member $i \in \mathcal{C}$ is determined by the equilibrium outcomes of the game with coalition $\mathcal{C} \setminus \{i\}$. This constraint thus requires a recursive definition.

Let $E^{C'}(BR^G, \mathcal{R})$ denote the set of Recursively Stable Agency Equilibrium (RAE) outcomes for the game with coalition C' (given restrictions \mathcal{R} and refinement BR^G). For coalitions of size $C = 1$ (that is, the coalition-less game $G(v)$), define the set of equilibria as

$$E^1(BR^G) = \{b \in \mathbb{R}_+^n : b_i \in BR_i^G(b_{-i}) \text{ for all } i \in I\}. \quad (7)$$

Now, suppose that $E^{C'}(BR^G, \mathcal{R})$ has been defined for all subcoalitions $C' \subset \mathcal{C}$. For each $i \in \mathcal{C}$, and for each $C' \subseteq \mathcal{C} \setminus \{i\}$, define

$$\bar{u}_i^{C'} = \begin{cases} \min_{b \in E^1(BR^G)} u_i(b) & \text{if } |C'| = 1 \\ \min_{b \in E^{C'}(BR^G, \mathcal{R})} u_i(b) & \text{if } |C'| \geq 2 \end{cases}.$$

The set of RAE of the game with coalition \mathcal{C} is defined as follows:

Definition 1. A Recursively Stable Agency Equilibrium (RAE) of the game G with coalition \mathcal{C} , given restrictions \mathcal{R} and independents' equilibrium refinement BR^G , is a profile of bids and conjectures $(b^*, \beta^*) \in A_{\mathcal{C}} \times B^*$ such that:¹²

1. The independents play a mutual best response: for all $i \in I \setminus \mathcal{C}$, $b_i^* \in BR_i^G(b_{-i}^*)$.
2. The conjectures of the agency are correct: $\beta^*(b_{\mathcal{C}}^*) = b_{-\mathcal{C}}^*$.
3. The agency best responds to the conjectures β^* , given the exogenous restrictions (R) and the stability restrictions about the independents and the coalition members (S.1 and S.2, respectively):

$$\begin{aligned} b_{\mathcal{C}}^* &\in \arg \max_{b_{\mathcal{C}}} u_{\mathcal{C}}(b_{\mathcal{C}}, \beta^*(b_{\mathcal{C}})) \\ \text{subject to} & : (R) \ b_{\mathcal{C}} \in R_{\mathcal{C}} \\ & : (S.1) \ b_{\mathcal{C}} \in S_{\mathcal{C}} \\ & : (S.2) \ \text{for all } i \in \mathcal{C}, \ u_i(b_{\mathcal{C}}, \beta^*(b_{\mathcal{C}})) \geq \bar{u}_i^{\mathcal{C} \setminus \{i\}} \end{aligned}$$

¹²Note that, by requiring $\beta^* \in B^*$, this equilibrium rules out the possibility that the coalition's bids are sustained by 'incredible' threats of the independents.

The set of RAE outcomes for the game with coalition \mathcal{C} (given BR^G and $R_{\mathcal{C}}$) is:

$$E^{\mathcal{C}}(BR^G, \mathcal{R}) = \{b^* \in A : \exists \beta^* \text{ s.t. } (b^*, \beta^*) \text{ is a RAE}\}. \quad (8)$$

In section 4 we will apply the RAE to study the impact of a SEMA on the GSP auction, and compare it to the benchmark VCG mechanism. The RAE in the GSP and the VCG mechanism are obtained from Definition 1, once the game G , the correspondence BR^G and the exogenous restrictions \mathcal{R} are specified accordingly. After presenting the main results below, we return in section ?? to discuss how the RAE relates to the existing literature.

In the meantime, it is useful to note that RAE outcomes in general are not Nash Equilibria of the baseline game, nor of the game in which the coalition is replaced by a single player: similar to Ray and Vohra (1997, 2014, RV) ‘equilibrium binding agreements’, the stability restrictions do affect the set of equilibrium outcomes, not merely as a refinement.¹³

4 Analysis

In the previous section we developed a machinery (the general notion of RAE) to study bid coordination in arbitrary mechanisms. In the following we apply it to the GSP auction and to the VCG mechanism, the traditional benchmark for the GSP auction in the literature.

Definition 2. *The RAE-outcomes in the GSP and VCG mechanism are defined as follows.*

Given constraints \mathcal{R} :

1. *The RAE of the GSP auction is obtained setting $G(v)$ and BR_i^G in Definition 1 equal to $\mathcal{G}(v)$ and to BR_i^* , respectively.*
2. *The RAE of the VCG mechanism are obtained from Definition 1 letting $G(v) = \mathcal{V}(v)$ and letting BR_i^G be such that independents play the dominant (truthful) strategy.*

¹³There are two main differences between RV and our model. First, our stability restriction (S.2) only allows the agency’s proposal to be blocked by individual members. In contrast, RV’s solution concept presumes that the coalition’s proposals may be blocked by the joint deviation of any set of members of the coalition. That advertisers can make binding agreements outside of the SEMA, and jointly block its proposals, seems unrealistic in our context. Second, the underlying model of non-cooperative interaction is Nash equilibrium in RV (this can be accommodated in our setting letting BR_i^G in Definition 1 coincide with the standard best-response correspondence), whereas we allow for refinements, which in the next Section will be used for the analysis of the VCG and GSP.

For both mechanisms, we will refer to the case where \mathcal{R} is such that $R_{C'} = A_{C'}$ for all $C' \subseteq I$ as the ‘unconstrained’ case.¹⁴

We first present the analysis of RAE in the VCG mechanism (section 4.1), and then proceed to the analysis of the GSP auction (section 4.2). Our main conclusion is that, when a SEMA is present, the VCG mechanism outperforms the GSP auction both in terms of revenues and allocative efficiency. These results therefore uncover a striking fragility of the GSP mechanism with respect to the possibility of coordinated bidding.

In some of the results below we have that the equilibrium bid for some i is such that $b_i = b_{i+1}$. Since ties are broken according to bidders’ labels (cf. footnote 6), in this case bidder i obtains the position above $i + 1$. To emphasize this, we will write $b_i = b_{i+1}^+$.¹⁵

4.1 Coordinated Bidding in the VCG mechanism

As anticipated in Definition 2, we apply RAE to the game $\mathcal{V}(v)$, assuming that the independents play the dominant strategy (truthful bidding), and letting the exogenous restrictions \mathcal{R} be vacuous (i.e., such that $R_C = A_C$ for every $C \subseteq I$). It is easy to check that, under this specification of the independents’ equilibrium correspondence, the set $S_C = A_C$. Hence, constraint (S.1) in Definition 1 plays no role in the results of this subsection.

Theorem 1. *For any \mathcal{C} , the unconstrained RAE of the VCG is unique up to the bid of the highest coalition member. In this equilibrium, advertisers are assigned to positions efficiently ($\rho(i) = i$), independents’ bids are equal to their valuations and all the coalition members (except possibly the highest) bid the lowest possible value that ensures their efficient position.*

¹⁴Under this definition, the RAE-outcomes for the coalitions of size one (eq.7) coincide precisely with the non-agency equilibria introduced in section 2. Namely, the LREF (EOS) equilibria in the GSP and the dominant strategy equilibrium in the VCG.

¹⁵Without the tie-breaking rule embedded in ρ (footnote 6), the SEMA’s best replies may be empty valued. In that case, our analysis would go through assuming that SEMA’s bids are placed from an arbitrarily fine discrete bid (i.e., $A_C = (\mathbb{R}_+ \cap \varepsilon\mathbb{Z})^{|C|}$ where ε is the minimum bid increment). In such alternative model, the case $b_i = b_{i+1}^+$ can be thought of as i bidding the lowest feasible bid higher than b_{i+1} , i.e. $b_i = b_{i+1} + \varepsilon$. All our results would hold in such a discrete model, once the equilibrium bids in the theorems are interpreted as the limit of the equilibria in the discrete model, letting $\varepsilon \rightarrow 0$ (the notation b_{i+1}^+ is thus reminiscent of this alternative interpretation, as the right-hand limit $b_{i+1}^+ := \lim_{\varepsilon \rightarrow 0} (b_{i+1} + \varepsilon)$). Embedding the tie-breaking rule in ρ allows us to avoid these technicalities, and focus on the key economic insights.

Formally, in any RAE of the VCG mechanism, the bid profile \hat{b} is such that

$$\hat{b}_i \begin{cases} = v_i & \text{if } i \in I \setminus \mathcal{C}; \\ = \hat{b}_{i+1}^+ & \text{if } i \in \mathcal{C} \setminus \{\min(\mathcal{C})\} \text{ and } i \leq S; \\ \in (\hat{b}_{i+1}^+, v_{i-1}) & \text{if } i = \min(\mathcal{C}) \text{ and } i \leq S. \end{cases} \quad (9)$$

where we denote $v_0 := \infty$ and $\hat{b}_{n+1} := 0$.

The RAE of the VCG mechanism therefore are efficient, with generally lower revenues than in the VCG without a SEMA. Moreover, the presence of a SEMA has no impact on the bids of the independents (this follows from the strategy-proofness of the mechanism).

The efficiency result in Theorem 1 is due to the stability restrictions in RAE, which limits the SEMA's freedom to place bids. Restriction (S.2), in particular, requires that the agency's proposal gives no member of the coalition an incentive to abandon it and bid as an independent. A recursive argument further shows that the payoff that any coalition member can attain from abandoning the coalition is bounded below by the equilibrium payoffs in the baseline (coalition-less) game, in which assignments are efficient. The 'Pigouvian' logic of the VCG payments in turn implies that such (recursive) participation constraints can only be satisfied by the efficient assignment of positions. It is interesting to notice that the recursive stability restriction is key to this result. As shown by the next example, without the recursive stability restriction (S.2), inefficiency is possible in the VCG with bid coordination:

Example 2. Let $n = 5$, $v = (40, 25, 20, 10, 9)$. The CTRs for the five positions are $x = \{20, 10, 9, 1, 0\}$, and suppose that $\mathcal{C} = \{1, 2, 5\}$. If 2 remains in the coalition and keeps his efficient position, RAE-bids are $\hat{b} = (\hat{b}_1, 20^+, 20, 10, 0)$, the coalition's payoff is 650, and 2 obtains 150 (again, \hat{b}_1 is not pinned down, but revenues and payoffs are). If 2 were to stay in, but drop one position down, the bids would be $\hat{b} = (\hat{b}_1, 10^+, 20, 10, 0)$, and payoffs for the coalition and 2 would be equal to 655 and 145, respectively. If 2 abandons the coalition, the bids in the game in which $\mathcal{C}' = \{1, 5\}$ are $\hat{b} = (\hat{b}_1, 25, 20, 10, 0)$, and 2 obtains a payoff of 150. Thus, the coalition would benefit from lowering 2's position, but the recursive stability condition does not allow such a move. Note also that the recursive definition matters: If the

outside option were defined by the case with no coalition at all, 2 would not drop out even when forced to take the lower position, since 2's payoff in that case would be $141 < 145$.

Whereas the presence of a SEMA does not alter the allocation of the VCG mechanism, it does affect its revenues: in any RAE of the VCG mechanism, the SEMA lowers the bids of its members (except possibly the one with the highest valuation) as much as possible, within the constraints posed by the efficient ranking of bids. Since, in the VCG mechanism, lowering the i -th bid affects the price paid for all slots $s = 1, \dots, \min\{S + 1, i - 1\}$, even a small coalition can have a significant impact on the total revenues.

Example 3. Consider the environment in Example 1, and suppose that $\mathcal{C} = \{1, 3\}$. Then, the RAE of the VCG mechanism is $\hat{b} = (\hat{b}_1, 4, 2^+, 2, 1)$. The resulting revenues are 86, as opposed to 96 of the non-agency case.

4.2 Coordinated Bidding in the GSP auction

We turn next to the GSP auction. According to the refinement of the best responses introduced above, we set BR_i^G in equation (4) equal to the marginally envy-free best response correspondence (eq. 1). The resulting correspondence $BR_{-\mathcal{C}}^G$ therefore assigns, to each profile $b_{\mathcal{C}}$ in the set of exogenous restrictions $R_{\mathcal{C}}$, the set of independent profiles that are fixed points of the BR_i^* correspondence for all $i \notin \mathcal{C}$.

We begin our analysis characterizing the RAE when the agency is constrained to placing bid profiles that could not be detected as ‘coordinated’ by an external observer (the ‘undetectable coordination’ restriction). Theorem 2 shows that the equilibrium outcomes of the GSP with this restriction are exactly the same as the unrestricted RAE of the VCG mechanism. We consider this problem mainly for analytical convenience, but the result of Theorem 2 also has independent interest, in that it characterizes the equilibria in a market in which ‘not being detectable as a bid coordinator’ is a primary concern of the SEMA.

We lift the ‘undetectable coordination’ restriction in section 4.2.2. We show that, unlike the VCG mechanism, the unrestricted RAE of the GSP auction in general may be inefficient and induce strictly lower revenues than their counterparts in the VCG mechanism. In light of the efficiency of the VCG mechanism we established earlier (Theorem 1), it may be tempting

to impute the lower revenues of the GSP auction to the inefficiencies that it may generate. To address this question, in section 4.2.2 we also consider the RAE of the GSP auction when the agency is constrained to inducing efficient allocations. With this restriction, we show that the equilibrium revenues in the GSP auction are always lower than in the VCG mechanism (Theorem 3). The revenue ranking therefore is not a consequence of the allocative effects.

4.2.1 The ‘Undetectable Coordination’ Restriction: A VCG-Equivalence Result

Consider the following set of exogenous restrictions: for any $\mathcal{C} \subseteq I$ s.t. $|\mathcal{C}| > 1$,

$$R_{\mathcal{C}}^{UC} := \left\{ b_{\mathcal{C}} \in A_{\mathcal{C}} : \exists v' \in \mathbb{R}_+^{|\mathcal{C}|}, b_{-C} \in \mathbb{R}_+^{n-|\mathcal{C}|} \text{ s.t. } (b_{\mathcal{C}}, b_{-C}) \in \mathcal{EG}(v'_{\mathcal{C}}, v_{-C}) \right\}.$$

These exogenous restrictions have a clear interpretation: the set $R_{\mathcal{C}}^{UC}$ is comprised of all bid profiles of the agency that could be observed as part of a LREF equilibrium in the GSP auction without the agency, given the valuations of the independents $v_{-C} = (v_j)_{j \in I \setminus \mathcal{C}}$. For instance, consider a SEMA whose primary interest is not being detectable as inducing bid coordination by an external observer. The external observer (e.g., the search engine or the anti-trust authority) can only observe the bid profile, but not the valuations $(v_i)_{i \in \mathcal{C}}$. Then, $R_{\mathcal{C}}^{UC}$ characterizes the bid profiles that ensure the agency would not be detected as inducing coordination, even if the independents had revealed their own valuations. The next result characterizes the RAE of the GSP auction under these restrictions, and shows its revenue and allocative equivalence to the unrestricted RAE of the VCG mechanism:

Theorem 2. *For any \mathcal{C} , in any RAE of the GSP auction under the ‘undetectable coordination’ (UC) restriction, the bids profile \hat{b} is unique up to the highest bid of the coalition and up to the highest overall bid. In particular, let $v_{n+1}^f = 0$, and for each $i = n, \dots, 1$, recursively define $v_i^f := v_{i+1}^f$ if $i \in \mathcal{C}$ and $v_i^f = v_i$ if $i \notin \mathcal{C}$. Then, for every i ,*

$$\hat{b}_i \begin{cases} = v_i^f - \frac{x^i}{x^{i-1}} (v_i^f - \hat{b}_{i+1}), & \text{if } i \neq 1 \text{ and } i \neq \min(\mathcal{C}); \\ \in \left[v_i^f - \frac{x^i}{x^{i-1}} (v_i^f - \hat{b}_{i+1}), \hat{b}_{i-1} \right) & \text{otherwise} \end{cases}, \quad (10)$$

where $\hat{b}_0 := \infty$ and $x^i/x^{i-1} := 0$ whenever $i > S$. Moreover, in each of these equilibria,

advertisers are assigned to positions efficiently ($\rho(i) = i$), and advertisers' payments are the same as in the corresponding 'unrestricted RAE' of the VCG mechanism (Theorem 1).

Note that, in the equilibrium of (10), every bidder i other than the highest coalition member and the highest overall bidder bids as an independent with valuation v_i^f would bid in the baseline competitive model (first line of eq. 10). For the independent bidders ($i \notin \mathcal{C}$), such v_i^f coincides with the actual valuation v_i . For coalition members instead, $v_i^f \neq v_i$ is a 'feigned valuation'. Though notationally involved, the idea is simple and provides a clear insight on the SEMA's equilibrium behavior: intuitively, in order to satisfy the UC-restriction, the SEMA's bids for each of its members should mimic the behavior of an independent in the competitive benchmark, for some valuation. The SEMA's problem therefore boils down to 'choosing' a feigned valuation, and bid accordingly. The optimal choice of the feigned valuation is the one which, given others' bids, and the bidding strategy of an independent, induces the lowest bid consistent with i obtaining the i -th position in the competitive equilibrium of the model with feigned valuations, which is achieved by $v_i^f = v_{i+1}^f$. Note that the fact that bidder i cannot be forced to a lower position is not implicit in the UC-restriction, but the result of the equilibrium restrictions.¹⁶ The last line of (10) corresponds to the bid of the highest coalition member and the highest overall bidder, required to be placed in their efficient position. The logic of the equilibrium implies that it results in an efficient allocation. Moreover, these equilibria induce the same individual payments (hence total revenues) as the unrestricted RAE of the VCG mechanism.

To understand the implications of this equilibrium, notice that in the GSP auction, the i -th bid only affects the payment of the $(i - 1)$ -th bidder. Hence, the 'direct effect' of bids manipulation is weaker in the GSP than in the VCG mechanism, where the payments for all positions above i are affected. Unlike the VCG mechanism, however, manipulating the bid of coalition member i also has an 'indirect effect' on the bids of all the independents placed above i , who lower their bids according to the recursion in (10).

Example 4. Consider the environment of Example 3, with $\mathcal{C} = \{1, 3\}$. Then, the UC-RAE is $\hat{b} = (\hat{b}_1, 2.9, 1.8, 1.6, 1)$, which results in revenues 86. These are the same as in

¹⁶The reason is similar to that discussed for Theorem 1, only here is more complicated due to the fact that, in the GSP auction, the bids of the SEMA alter the bids placed by the independents.

the VCG mechanism, and 10 less than in the non-agency case (Example 1). Note that the bid $\hat{b}_3 = 1.8$ obtains setting $v_3^f = v_4 = 2$, and then applying the same recursion as for the independents. Also note that the ‘direct effect’, due to the reduction in \hat{b}_3 , is only equal to $(b_3^{EOS} - \hat{b}_3) \cdot x_2 = 5$ (where b_3^{EOS} denotes 3’s bid in the non-agency benchmark). Thus, 50% of the revenue loss in this example is due to the SEMA’s ‘indirect effect’ on the independents.

Thus, despite the simplicity of the payment rule in the GSP auction, the equilibrium effects in (10) essentially replicate the complexity of the VCG payments: once the *direct* and *indirect effects* are combined, the resulting revenue loss is the same in the two mechanisms. This result also enables us to simplify the analysis of the impact of the SEMA on the GSP, by studying the comparative statics of the unconstrained RAE in the VCG mechanism. We can thus obtain simple qualitative insights for this complicated problem:

Remark 1. *Both in the VCG mechanism and in the UC-RAE of the GSP auction:*

1. *Holding everything else constant, the revenue losses due to the SEMA increase with the differences $(x_{i-1} - x_i)$ associated to the SEMA’s clients $i \in \mathcal{C}$.*
2. *Holding $(x_s - x_{s+1})$ constant (i.e., if it is constant in s), the revenue losses due to the SEMA are larger if: (i) The coalition includes members that occupy a lower position in the ranking of valuations; (ii) The coalition includes members that occupy adjacent positions in the ranking of valuations; (iii) The difference in valuations between members of the coalition and the independents immediately below them (in the ranking of valuations) is larger.*

Part 1 is immediate from Theorem 1 and the transfers of the VCG payment (section 2). Part 2 is also straightforward: point (i) follows from the fact that, for any given decrease of the bid of a single bidder, the total reduction in revenues of the VCG mechanism increases with the number of agents placed above him. Point (ii) is due to the fact that, for any $i \neq \min \{\mathcal{C}\}$, if $i+1$ also belongs to the coalition, then the SEMA can lower i ’s bid below v_{i+1} , still preserving an efficient allocation. Point (iii) follows because, the lower the valuations of the independents ranked below a member of the coalition, the more the SEMA has freedom to lower the bid of that member without violating the efficient ranking of bids (Theorem 1).

We conclude the analysis of the UC-RAE of the GSP auction providing a characterization of its testable implications. In particular, the equilibrium characterization in Theorem 2

involves bidders' valuations. But since valuations are typically not observable to an external analysts, the conditions in (10) may appear to be of little help for empirical analysis. Those terms, however, can be rearranged to obtain a characterization that only depends on the CTRs and the individual bids.

Corollary 2. *For any \mathcal{C} , in any RAE of the GSP auction under the ‘undetectable coordination’ (UC) restriction, the bids profile \hat{b} satisfies the following conditions:*

- if $i \notin \mathcal{C}$:

$$\underbrace{\frac{\hat{b}_i x^{i-1} - b_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i} > \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \quad (11)$$

- if $i \in \mathcal{C}$ and $i \neq \min(\mathcal{C})$:

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i^f (\leq v_i)} = \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \quad (12)$$

These conditions are easily comparable to the analogous characterization obtained for the competitive benchmark (equation 3), and will provide the basic building block for the applications to ad auctions data in Section 5.

4.2.2 Lifting the UC-Restriction: Revenue Losses and Inefficiency

As discussed in section 4.1, the revenues generated by the VCG may be largely affected by the presence of a SEMA, even if it comprises a small number of members. Theorem 2 therefore already entails a fairly negative outlook on the sellers' revenues in the GSP auction when a SEMA is active, even when it cannot be detected as explicitly inducing any kind of ‘coordinated bidding’. In this section we show that, when the undetectability constraint is lifted, a SEMA may induce larger revenue losses as well as inefficient allocations in the GSP auction. Before doing that, however, we first consider a weaker set of exogenous restrictions, which force the SEMA to induce efficient allocations. This is useful to isolate the price-reducing effect of bidding coordination separately from its potential allocative effect.

Theorem 3 shows that, even with this restriction, the auctioneer’s revenues are no higher than in the unrestricted equilibria of the VCG mechanism.

Formally, let $\mathcal{R}^{EFF} = \{R_C^{EFF}\}_{C \subseteq I: |C| \geq 2}$ be such that, for each non trivial coalition $C \subseteq I$,

$$R_C^{EFF} := \{b_C \in A_C : \exists b_{-C} \in BR_{-C}^*(b_C) \text{ s.t. } \rho(i; (b_C, b_{-C})) = i \ \forall i \in I\}.$$

Definition 3. *An efficiency-constrained RAE of the GSP auction is a RAE of the GSP auction where the exogenous restrictions are given by $\mathcal{R} = \mathcal{R}^{EFF}$ and the agency’s conjectures β^* satisfy $\rho_i(b_C, \beta^*(b_C)) = i$ for all $b_C \in R_C^{EFF}$ and all $i \in I$.*

Theorem 3. *Efficiency-constrained RAE of the GSP auction exist; in any such RAE: (i) the agency’s payoff is at least as high as in any RAE of the VCG mechanism, and (ii) the auctioneer’s revenue is no higher than in the corresponding equilibrium of the VCG auction. Furthermore, there exist parameter values under which both orderings are strict.*

By imposing efficiency as an exogenous constraint, Theorem 3 shows that the structure of the payments in the GSP auctions determines a fragility of its performance in terms of revenues, independent of the allocative distortions it may generate. The intuition behind Theorem 3 is simple, in hindsight: In the VCG mechanism, truthful bidding is dominant for the independents, irrespective of the presence of the SEMA. Hence, by manipulating the bids of its members, the agency cannot affect the bids of the independents (though it may affect their payments, as seen in Example 3). The SEMA’s manipulation of the bids of its members therefore only has a *direct effect* on the total revenues. In the GSP auction, in contrast, the SEMA has both a *direct* and an *indirect effect* on the total revenues. The latter is due to its ability to affect the equilibrium bids of the independents.

Under the UC-restrictions, the two effects combined induce exactly the same revenue-loss as in the VCG mechanism. Since the RAE with the UC-restriction also induce efficient allocations, it may seem that Theorem 3 follows immediately from the efficiency constraint being weaker than the UC-restriction. This intuition is incorrect for two reasons. First, the UC-constraint requires the existence of *feigned valuations* which can rationalize the observed bid profile, but does not require that they preserve the ranking of the true valuations. Second, when the exogenous restrictions $\mathcal{R} = (R_C)_{C \subseteq I: |C| \geq 2}$ are changed, they change for all

Table 1: Summary of Results in Examples

Valuations	VCG	GSP (EOS)	RAE in VCG	UC-RAE in GSP	(Eff.) RAE in GSP
5	5	\mathbf{b}_1	\mathbf{b}_1	\mathbf{b}_1	\mathbf{b}_1
4	4	3.15	4	2.9	2.8
3	3	2.3	2⁺	1.8	1.6⁺
2	2	1.6	2	1.6	1.6
1	1	1	1	1	1
Revenues	96	96	86	86	82

Summary of results in examples 1, 3-5. Coalition members' bids and valuations are in bold. The VCG and GSP columns represent the competitive equilibria in the two mechanisms as described in example 1. The RAE in VCG and the revenue equivalent UC-RAE in the GSP are from examples 3 and 4 respectively. The last column denotes both the Efficient RAE and the unrestricted RAE of the GSP auction, which coincide in example 5.

coalitions: hence, even if $R_{\mathcal{C}}$ is weaker for any given \mathcal{C} , the fact that it is also weaker for the subcoalitions may make the stability constraint (S.2) more stringent. Which of the two effects dominates, in general, is unclear. Hence, because of the ‘farsightedness assumption’ embedded in constraint (S.2), the proof of the theorem is by induction on the size of the coalition.

The next example illustrates the Eff-RAE in the environment of Examples 3 and 4. Table 1 compares the bid profiles and revenues of the equilibria illustrated in our leading examples.

Example 5. *Consider the environment of Examples 3 and 4, with $\mathcal{C} = \{1, 3\}$. The efficiency-constrained RAE is $\hat{b} = (\hat{b}_1, 2.8, 1.6^+, 1.6, 1)$, which results in revenues 82, which are lower than the RAE in VCG mechanism (86). Note that, relative to the UC-RAE in Example 4, the coalition lowers b_3 to the lowest level consistent with the efficient ranking. This in turn induces independent bidder 2 to lower his bids, hence the extra revenue loss is due to further direct and indirect effects. We note that the efficiency restriction is not binding in this example, and hence the Eff-RAE and the unconstrained RAE coincide.*

Since, under the efficiency restriction, the GSP auction induces the same allocation as the VCG mechanism, the two mechanisms are ranked in terms of revenues purely due to the SEMA’s effect on prices. Obviously, if allocative inefficiencies were introduced, they would provide a further, independent source of revenue reduction. As already noted, this is not the case in Example 5, in which the efficiency constraint is not binding. Unlike the VCG mechanism, however, the unrestricted RAE of the GSP auction can be inefficient as well:

Example 6. Consider a case with $n = 8$, CTRs $x = (50, 40, 30.1, 20, 10, 2, 1, 0)$ and valuations $v = (12, 10.5, 10.4, 10.3, 10.2, 10.1, 10, 1)$. Let $\mathcal{C} = \{5, 6\}$. The unrestricted RAE is essentially unique (up to the highest overall bid) and inefficient, with the coalition bidders obtaining slots 4 and 6. Equilibrium bids (rounding off to the second decimal) are $b = (b_1, 9.91, 9.76, 9.12, 9.5, 7.94, 5.5, 1)$. The inefficiency arises as follows. The agency drastically reduces the bid of its lower-valued member to benefit the other member. This, however, creates incentives for the independents $i < 5$ to move down to the position just above bidder 6, in order to appropriate some of the rents generated by the reduction of its bid. In order to prevent these independents from doing so, 5's bid must also be reduced, to make the higher positions more attractive. But in this example, the reduction of 6's bid is large enough that the undercut of 4 is low enough that the coalition actually prefers giving up slot 5 to the independent, and climb up to the higher position. Thus, the coalition does not benefit directly from the reduction of 6's bid, but indirectly, by attracting 4 to the lower position.

Similarly to what done before, we conclude this discussion with some testable implications of the Efficiency Constrained RAE in the GSP auction, which will be useful to the empirical analysis in Section 5.

Corollary 3. For any \mathcal{C} , in any Eff-RAE of the GSP auction under, the bids profile \hat{b} satisfies the following conditions:

- if $i \notin \mathcal{C}$:

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i} > \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \quad (13)$$

- if $i \in \mathcal{C}$ and $i \neq \min(\mathcal{C})$:

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{\text{less than } v_i} < \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \quad (14)$$

4.3 Discussion of the theoretical results

The VCG mechanism is typically regarded as a poor mechanism under collusion. Yet, we have shown that, in the presence of a SEMA, it outperforms the GSP auction both in

terms of revenues and allocative efficiency. On the one hand, this conclusion provides a remarkably negative result for the GSP auction, which highlights an important fragility of this mechanism with respect to the possibility of coordinated bidding. On the other hand, the efficiency result in Theorem 1 shows that the desirable allocative properties of the VCG mechanism are robust to the introduction of a SEMA, suggesting that the VCG mechanism may be more resilient than we might have expected.

Independent of the efficiency result, which we discuss below, a key source of the resilience of the VCG mechanism is its strategy-proofness: if individual bidders have a dominant strategy, then a SEMA may have at most a *direct effect* on the total revenues, as independents have no reason to adjust their bids as a function of the strategy and composition of the agency. In the GSP auction, in contrast, the SEMA can manipulate the bids of its members and induce lower bids also amongst the independents. As shown in Example 4, the resulting *indirect effect* may be significant: due to the complex equilibrium effects of the GSP auction, even a small manipulation of the coalition’s bids may have a large effect on total revenues. In general, indirect effects can be completely avoided only if independents’ best responses are unaffected by the bids of the agency. This suggests that strategy-proofness may be a desirable property in the presence of a SEMA.

The efficiency result of Theorem 1 clearly relies on the assumptions of our model. As illustrated in Example 2, without the recursive stability constraint (S.2), coordinated bidding in the VCG mechanism may induce inefficiencies. But as soon as individuals’ incentives to abandon the coalition are taken into account, efficiency is restored.¹⁷ Thus, the significance of the efficiency result depends on whether or not we believe that a mechanism’s performance should be evaluated taking into account the underlying coalition formation process. For a medium or long-run perspective, we think that this is important, hence we built a ‘farsightedness assumption’ into our model.

¹⁷Bachrach (2010) studies collusion in the VCG mechanism in a classical cooperative setting (i.e. without distinguishing the SEMA clients from the independents, and without Ray and Vohra’s farsightedness assumption), finding that the VCG is vulnerable to this form of collusion.

4.4 Extension: Competition between Agencies

The case of multiple SEMAs placing bids for multiple advertisers, ignored in the analysis of the previous sections, rarely happens in the data.¹⁸ Nevertheless, it is interesting to consider how the presence of multiple SEMAs would affect our results.

In this section we show that competition between agencies may have mixed effects on the market. On the one hand, for most configurations of the coalition structure, our earlier results extend to the case with multiple SEMAs essentially unchanged. Clearly, the revenue losses for the search engines will be less pronounced when the same set of coordinating bidders is divided into two (or more) competing subsets of SEMAs, but they would still be substantial. On the other hand, for many configurations of coalitions, equilibria in pure strategies will either not be unique or stop existing altogether. This means that, for certain configuration of coalitions, bidding cycles are likely to emerge under the current mechanism – a phenomenon that had been observed for the earlier mechanisms used in this market, and which is considered to be the main reason of the transition from such earlier mechanisms to the GSP auction.¹⁹ Hence, while competition between agencies may produce the expected result of mitigating the revenue losses due to bidding coordination, it may also impair the working of the current mechanisms in a more fundamental way, with the potential to disrupt the current market arrangement.

For simplicity, we consider the case with two SEMAs (the extension to more than two agencies is cumbersome but straightforward). We also assume that the highest-placed bidder within any coalition bids as if he were an independent. Since his bid does not affect the coalition’s payoff, we had left this bid unrestricted in the earlier analysis. With more than one coalition, however, this way of breaking the indifference helps to characterize the set of equilibria. The next Theorem shows that, for certain coalition structures, the results of the previous sections extend essentially unchanged to the multiple agencies case. But for other coalition structures, bidding cycles may emerge:

¹⁸Decarolis, Goldmanis and Penta (2016) document that, consistent with the agencies’ specialization, it is almost never the case that more than one agency bids on behalf of multiple bidders in the same auction.

¹⁹See Edelman and Ostrovsky (2007) for a discussion of bidding cycles in the Overture’s first price auctions, and Ottaviani (2003) for an early assessment of the transition from first price to GSP auctions.

Theorem 4. *1. If no members of different coalitions occupy adjacent positions in the ordering of valuations, and all members of one coalition are above all members of the other, then the UC-RAE of the GSP with multiple coalitions is unique. In this equilibrium, the allocation is efficient and the search engine revenues are weakly higher than those of the UC-RAE in which all members of the different coalitions bid under the same SEMA, but no higher than under full competition. Moreover, both the allocation and the associated revenues are identical to those resulting in the equilibrium of a VCG mechanism.*

2. If non-top members of different coalitions occupy adjacent positions in the ordering of valuations, then no unconstrained RAE of the VCG and no UC-RAE of the GSP exist.

The first part of the theorem extends Theorems 1 and 2 to the case of multiple SEMAs. The result therefore shows that competition between agencies may mitigate, but not solve, the loss in revenues due to coordinated bidding. In fact, if coalitions have bidders in adjacent positions (part 2 of the Theorem), further problems arise, such as non-existence of pure-strategy equilibria and bidding cycles. Hence, the overall message from the previous section is essentially unaffected by the possibility of competition between agencies: the diffusion of SEMAs is doomed to alter the functioning of the existing mechanisms in a fundamental way, and especially so the GSP auction, which is particularly fragile to the possibility of coordinated bidding.

We illustrate both these points with the following example. Consider once again our workhorse example of Table 1. Table 2 reports the EOS equilibrium bids (second column) as well as the bids under different coalition structures. First, we look at the case of a single coalition $\mathcal{C} = \{1, 2, 4, 5\}$. According to our earlier results, in the UC-RAE under this SEMA the bottom two bidders bid zero. This has an indirect effect on the independent bidder (3), who lowers his bid from 2.3 to 1.5, thereby lowering the payments and bids for bidders 1 and 2. If we split this coalition into two separate coalitions, however, things will change depending on the way we partition these four bidders. Suppose the split is that in the fourth column of the table: $\mathcal{C}_1 = \{1, 2\}$ and $\mathcal{C}_2 = \{4, 5\}$. In this case, the two coalitions have no adjacent members, hence part 1 of Theorem 4 applies. In equilibrium, the revenues amount to 88 which is above what found for the case of a single coalition (60), but still well below

the EOS case, in which all bidders are independent (96).²⁰

Table 2: Competition between Agencies

Valuations	GSP (EOS)	Single Coalition: $\mathcal{C} = \{1, 2, 4, 5\}$	Two Coalitions: $\mathcal{C}_1 = \{1, 2\}, \mathcal{C}_2 = \{4, 5\}$	Two Coalitions: $\mathcal{C}_1 = \{1, 4\}, \mathcal{C}_2 = \{2, 5\}$
5	b_1	5	5	b_1
4	3.15	2.75	3.05	b_2
3	2.3	1.5	2.1	b_3
2	1.6	0⁺	1.2	b_4
1	1	0	0	b_5
Revenues	96	60	88	—

Suppose next that the coalition of four instead is split as in the last column of Table 2: $\mathcal{C}_1 = \{1, 4\}$ and $\mathcal{C}_2 = \{2, 5\}$. In this case, \mathcal{C}_2 would ideally like to set $b_5 = 0$. If it does so, however, \mathcal{C}_1 will find optimal to set $b_4 = 0^+$. This, however, is incompatible with an equilibrium because once $b_4 = 0^+$, the player in position 5 will prefer to leave the coalition and climb one position up by placing a bid above b_4 . On the other hand, if b_4 is set sufficiently high that the bidder in position 5 does not want to jump up, then, conditional on b_4 , \mathcal{C}_2 would like to place b_5 equal to zero. But then, a strictly positive b_4 cannot be optimal for \mathcal{C}_1 . Part 2 of Theorem 4 says that this instability is widespread, and that bidding cycles will likely emerge for all those configurations in which two bidders from separate coalitions (that are not the top bidder of their respective coalitions) are contiguous in the ordering of valuations. Interestingly, the economics behind this phenomenon is nearly identical to that explained by Edelman and Ostrovsky (2007) in their characterization of the original Generalized First Price (GFP) auction system with which the market started: the same type of bidding cycles that characterized the GFP – and that lead to introduction of the GSP – emerge under the GSP with multiple SEMAs. This is a troubling result for the GSP, showing that competition between agencies, instead of fixing the problems, could exacerbate the instability of the system.

²⁰Note that, if the highest placed member of the lower coalition (i.e., the bidder with a value of 2 in this example) were to slightly increase/decrease his bid, his coalition’s payoffs would not change, but the revenues of the other coalition would correspondingly decrease/increase. Hence, it is evident that without the assumption that top coalition members behave as an independents, a multiplicity of equilibria might arise. Different selections from the best-response correspondence may be used to model other forms of behavior, such as spiteful bidding (e.g., Levin and Skrzypacz, 2016).

5 Applications to Ad Auctions

In this section, we show how our model can be used with typical data available to a search engine to detect collusion in its ad auctions. We first present the method and then illustrate its application through simulated data.

The data observed by search engines like Google or Microsoft-Yahoo! include information on all the variables entering our model, with the exception of advertisers' valuations. In particular, search engines record the advertisers' identity, their SEMA (if any), bids, positions and clicks received. These data would also allow a search engine to estimate CTRs as click frequencies. A major feature of the typical data is that they also record 'quality scores'. This would indeed be the case for Google or Microsoft-Yahoo! (but not, for instance, for Taobao) where a variant of the GSP auction is used: these quality scores are the advertisers' idiosyncratic score assigned by the search engine to account for various quality dimensions, including the CTRs. Quality scores concur in determining the assignment of advertisers to slots and prices: advertisers are ranked by the product of their bid and quality score, and pay a price equal to the minimum bid consistent with keeping that position.

Formally, letting e_i denote the score of bidder i , advertisers are ranked by $e_i \cdot b_i$, and CTRs are equal to $e_i \cdot x^{\rho(i)}$, the product of a 'quality effect' and a 'position effect'. The price paid by bidder i in position $\rho(i)$ is $p_i = e^{\rho(i+1)} b^{\rho(i+1)} / e_i$.²¹ Relabeling advertisers so that $e_i v_i > e_{i+1} v_{i+1}$, the competitive (EOS) equilibrium bids are such that, for all $i = 2, \dots, S$,

$$e_i v_i = \frac{e_i b_i x^{i-1} - e_{i+1} b_{i+1} x^i}{x^{i-1} - x^i} > \frac{e_{i+1} b_{i+1} x^i - e_{i+2} b_{i+2} x^{i+1}}{x^i - x^{i+1}} = e_{i+1} v_{i+1}. \quad (15)$$

This is the analogue, with quality scores, of what is implied by equation (2) in Lemma 1 and by Corollary 1. As shown below, similar modifications apply to the equilibria with agency coordination and this formulation is the basis for our proposed criterion to detect collusion.

²¹In extending the model to accommodate quality scores, we again follow EOS and Varian (2007). Clearly, the baseline model of the previous sections obtains letting $e_i = 1$ for all i .

5.1 Detecting Coordination: Strategy

We seek to devise a criterion that allows a search engine to say whether the data are more likely to be generated by competitive (EOS) bidding or by one of the models of agency coordination (UC-RAE, Eff-RAE and RAE). The latter models differ from EOS in that the bids of all agency bidders, with the exception of the highest value coalition member, are ‘too low’. This property of the coordination models leads to the following classification criterion that we present for the case of a 2-bidder coalition (the extension to larger coalitions is straightforward): let j denote the lowest value agency bidder, then an EOS-compatible bid profile requires that:

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

The key idea of our criterion is immediately evident from the comparison of the testable implication of the EOS model above with those of our collusive models. As shown by equations (12) and (14), under our coordination models, bidder j will lower its bid relative to the EOS case so that the above inequality is violated: either it holds as an equality, in the UC-RAE case, or it is inverted, in Eff-RAE and RAE cases. Importantly, for all other bidders $i \neq j$, our coordination models do not imply any observable difference relative to EOS, as shown by equations (11) and (14). Thus, a criterion based on the above inequality for bidder j captures the *only* feature which differentiates coordinated from competitive EOS bidding.

Thus, if we have as set of T auctions, $t = 1, 2, \dots, T$ for the same keyword/coalition and for which we observe quality scores, bids, CTRs and positions for all bidders, we can construct a quantity which captures the realization of the above inequality in each auction t as follows:

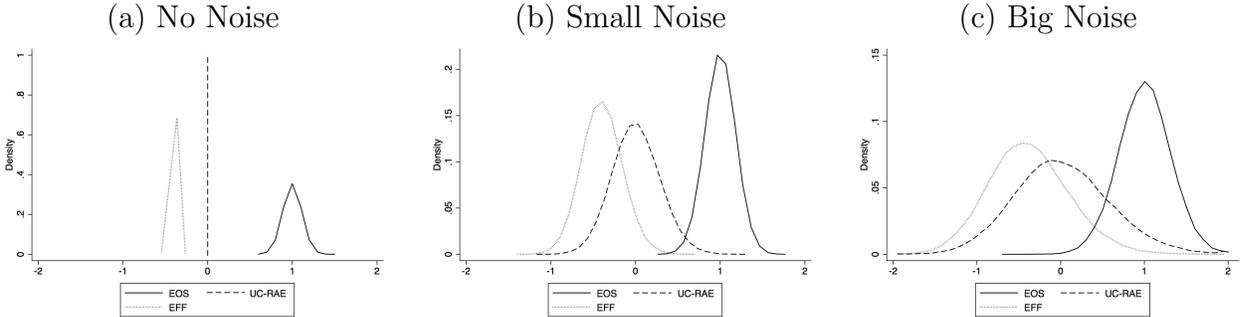
$$J_t = \frac{e_j b_j x^{j-1} - e_{j+1} b_{j+1} x^j}{x^{j-1} - x^j} - \frac{e_{j+1} b_{j+1} x^j - e_{j+2} b_{j+2} x^{j+1}}{x^j - x^{j+1}}.$$

If across many auctions for the same keyword we have evidence in favor of a positive J_t , this will be evidence in favor of competition. Otherwise, the evidence will be in favor of collusion. We next turn to simulated data to illustrate how to operationalize this idea.

5.2 Simulation

Consider once again the example in Table 1. We hold fixed the valuations, CTRs and coalition structure as in Table 1 and construct 100,000 simulated replicas of this auction by randomly drawing quality scores. For each auction and bidder, we take independent draws from a Normal distribution with mean 1 and s.d. 0.03. Since, as reported in Table 1, the lowest value member of the coalition is the bidder with a value of 3, we thus proceed to calculate the value of J_t for this bidder for all simulated auctions under three different equilibrium scenarios and report the resulting distribution of J_t in panel (a) of Figure 1: EOS (solid line), UC-RAE (dashed line) and Eff-RAE (dotted line).

Figure 1: Simulation



The distributions in panel (a) show that, as expected, J_t is never negative when we simulate EOS, it always equals zero when we simulate UC-RAE and it is never positive when we simulate Eff-RAE.²² Under the ideal conditions of the simulation, the observation of the distribution of J_t thus allows us to unambiguously separate the bidding models. Clearly, with real data, this tool should be expected to face some limits.

For instance, search engines update quality scores in real time. Hence, even if bidders can frequently readjust bids, it is not the case that bids are always optimized for the ‘true’ quality scores. Albeit small, the presence of belief errors about quality scores can impact J_t . To illustrate this point, in plot (b) and (c) of Figure 1 we repeat the previous simulation under

²²Detecting bids as coming from UC-RAE, in which coordinated bids were defined as ‘undetectable’, may strike as oxymoronic. The reason is that UC-RAE is undetectable in a single auction, but because it entails that J_t is exactly zero, it becomes detectable once many auctions are considered: $J_t = 0$ in every auction would be possible only if valuations were changing with the quality scores in an ad hoc way, hence the detectability of UC-RAE across auctions.

two scenarios. In both cases, we consider a belief error that enters multiplicatively: for each bidder i and auction t the true quality score is e_{it} , but bidders believe the true score to be \tilde{e}_{it} , where $\tilde{e}_{it} = d_{it} \cdot e_{it}$. Panel (b) considers the case of a small error, with $d_{it} \sim \mathcal{N}(1, 0.05^2)$, while panel (c) considers the case of a larger error, with $d_{it} \sim \mathcal{N}(1, 0.1^2)$. These two cases illustrate that, with any belief error, the distribution of J_t under UC-RAE is no longer degenerate at zero. This implies the need to search for UC-RAE cases by looking at an interval around zero, thus introducing some arbitrariness in the use of the J_t criterion. Moreover, overlaps in the three distributions make it more ambiguous to discriminate between the different models. In panel (b), the relatively small amount of noise still allows us to correctly classify the bidding models by looking at whether most of the mass of the distribution lies to the left of zero, around zero or to the right of zero. In practice, this can be operationalized in many ways by looking, for instance, at the smallest interval including majority of the mass, or, alternatively, by looking at some summary measure like mean, median or mode. As shown by panel (c), however, when the amount of noise is big, none of these methods will give rise to a satisfactory classification. Nevertheless, based on the empirical findings in Varian (2007) and Athey and Nekipelov (2012), it is reasonable to expect that the amount of belief noise is often rather small in the data so that our proposed criterion will typically be a useful tool to detect potential collusion.

Finally, for those cases where the above approach reveals the likely presence of collusion, a simple approach can be followed to invert bids and recover estimates of the potential revenue losses, even in the presence of belief errors about quality scores. To see this, suppose that we observe a 2-bidder coalition that bids according to one of our models of coordination. Then, if j is the lowest valued agency member, his value is bounded below by the value of the bidder in position $\rho(j+1)$ and above by the bidder in position $\rho(j-1)$ (or by the bidder in position $\rho(j-2)$, if the two agency bidders are contiguous). Therefore, if we assume that the data are generated by one of our equilibrium models, the one-to-one mapping that these equilibria imply between the independents' bids and their valuations can be used to retrieve their values and, hence, the bounds for the values of the coalition members.²³ Although no

²³To reconcile this approach with the belief errors discussed above, when inverting bids into valuations, an approach similar to Varian (2007) can be followed assuming that the realized belief errors are the smallest errors required to rationalize the data as coming from equilibrium bidding.

bound can be derived when the coalition occupies the top two slots or when its lowest valued member has no bidder below it, in all other cases this approach will be informative and will allow a search engine to compute counterfactual revenues under competitive bidding.

6 Conclusions

This is the first study to focus on the role of agencies on sponsored search auctions, and in particular on their role in coordinating the bids of different advertisers. Our theoretical results uncover a striking fragility of the GSP auction to bid coordination. This is confirmed by the empirical analysis in Decarolis, Goldmanis and Penta (2016), which reveals that even the small 2-bidder coalitions frequently observed in the data can have large effects on revenues. Aside from its theoretical interest, this is a first order finding since most of the online marketing is still passing through GSP auctions. Our finding might also provide a rationale for why Facebook has recently adopted the VCG and Google is said to be considering the transition. Shifts between one mechanism and the other are of tremendous interest given the large stakes involved and the fact that the proper functioning of this market is essential for both advertisers to reach consumers and consumers to learn about products.

From a methodological perspective, we note that the notion of RAE – and particularly the ‘*farsightedness*’ idea – has been key to obtain clear results in this complicated auction, in which competitive and coordinated bidding coexist. This suggests that this broader approach, which combines cooperative and non-cooperative ideas, may be fruitful to address the important problem of partial cartels, an outstanding challenge in the literature.

Clearly, our results are also interesting from a market design perspective. While beyond the scope of this paper, our analysis suggests some possible guidelines for research in this area. For instance, our analysis of the GSP auction with ‘undetectable coordination’ constraints implicitly suggests a way of deriving reservation prices to limit the impact of bids coordination. This kind of intervention would thus reinforce the resilience of the GSP auction, without entailing major changes in the mechanism. More radical modifications of the mechanism may be pursued as well. Theorem 1 shows that, in this setting, the VCG mechanism performs surprisingly well in the presence of bid coordination. As discussed in section

4, this is largely due to the strategy-proofness of this mechanism. While the complexity of the VCG payments is often seen as an impediment to the actual implementability of this mechanism, our analysis suggests that strategy-proofness may be a desirable property for a mechanism to perform well in the presence of bid coordination. Thus, variations of uniform price auctions may also be simpler and more viable options to address bid coordination.

Furthermore, from a broader perspective, we note that our findings are important to understand recent developments in online advertising. In this respect, they complement other the recent work, like Blake, Nosko and Tadelis (2015) and Einav, Farronato and Sundaresan (2014). The former paper explores, through large scale experiments, how eBay could benefit from a more nuanced bidding behavior that distinguishes between brand and non-brand keyword ads. The latter study, instead, focuses on the consumers' side documenting a decline in the importance of consumers' bidding in the eBay auctions with a progressive shift towards purchasing at posted prices. Our results, instead, focus on the advertisers' side analyzing the ongoing switch from advertisers' bidding to delegated bidding via SEMA. Altogether, it emerges the picture that bidding behavior in online marketing platforms is undergoing important transformations that still need careful analysis.

Finally, as pointed out earlier, our findings are also potentially relevant from an antitrust perspective. In particular, the agency behavior in our model is analogous to that of buying consortia, which have been sanctioned in the past (see footnote 5). Nevertheless, the specificities of the sponsored search advertisements market suggest a more nuanced view of the harm to the consumers. First, although multiple search engines exist, the degree of competition between them is likely substantially less than that between most of the advertisers. Since the lower auction prices imply a reduction in the marginal cost advertisers pay to reach consumers, advertiser competition would thus imply that some savings are passed on to consumers. Therefore, harm to consumers would result only if the agency engages in coordinating not only the auction bids, but also the prices charged to consumers. Second, bid coordination can negatively affect the quality of the service received by consumers by exacerbating further the advantage of dominant search engines relative to fringe ones. In Europe, for instance, where 90% of the searches pass via Google, agencies might be rather careful not to harm Google given the risk of being excluded from its results page. Smaller

search engines cannot exert such a threat because agencies are essential to attract new customers. The shift of revenues from small search engines to marketing agencies could thus deprive the former of the essential resources needed for technology investments. Thus, to the extent that competing search engines exert pressure for quality improvements, bid coordination poses a potential threat to consumer welfare.²⁴ All these considerations represent potentially fruitful directions for future research.

References

1. Association of National Advertisers (ANA), (2011). *Trends in Digital Agency Compensation*, 4th Edition, ANA publishing.
2. Athey, S., D. Nekipelov, (2012). “A Structural Model of Sponsored Search Advertising Auctions,” mimeo.
3. Bachrach, Y., (2010). “Honor Among Thieves - Collusion in Multi-Unit Auctions,” *Proc. of 9th int. Conf. on Autonomous Agents and Multiagent Systems* van der Hoek, Kaminka, Luck and Sen (eds.), Toronto Canada.
4. Balseiro, S. r., O. Besbes, G. Weintraub, (2015). “Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design,” *Management Science*, forthcoming.
5. Bernheim, B. D., M. D. Whinston, (1985). “Common Marketing Agency as a Device for Facilitating Collusion,” *The RAND Journal of Economics*, 16(2), 269–281.
6. Bernheim, B. D., M. D. Whinston, (1986). “Common Agency,” *Econometrica*, 54(4), 923–942.
7. Blake, T., C. Nosko, S. Tadelis, (2015). “Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment,” *Econometrica*, 83(1):155-174.
8. Borgers, T., I. Cox, M. Pesendorfer, V. Petricek, (2013). “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” *American Economic Journal: Microeconomics*, 5, 163-187

²⁴Quality of the links is indeed considered relevant for antitrust actions. For instance, one of the claims in the ongoing Google case before the European antitrust authority is the alleged abuse by Google of its dominant position to present links of inferior quality by directing consumers to Google’s own outlets.

9. Che, Y.K., Choi, S. and J. Kim (2013). "An Experimental Study of Sponsored-Search Auctions," *mimeo*.
10. Choi, H. and H. Varian, (2012). "Predicting the Present with Google Trends," *Economic Record*, 88 (s1), 2-9.
11. Celis, L. E., G. Lewis. M. M. Mobius, H. Nazerzadeh, (2015). "Buy-it-now or Take-a-chance: A New Pricing Mechanism for Online Advertising," *Management Science*, *forthcoming*.
12. Decarolis, F., M. Goldmanis, A. Penta, (2016). "Common Agency and Coordinated Bids in Sponsored Search Auctions, *mimeo*.
13. Econconsultancy (EC), (2011). "State of Search Marketing Report 2011," accessed on January 2012 at: <http://econsultancy.com/us/reports/sempo-state-of-search>
14. Edelman, B., M. Ostrovsky, M. Schwarz, (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242-259.
15. Edelman, B., M. Ostrovsky, (2007). "Strategic Bidder Behavior in Sponsored Search Auctions," *Journal of Decision Support Systems*, v.43(1), pp.192-198.
16. Edelman, B., M. Schwarz, (2010). "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100(2), 597-602.
17. Einav, L., C. Farronato, J. Levin, N. Sundaresan, (2014). "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?," *mimeo*.
18. Gomes, R., K. Sweeney, (2014). "Bayes-Nash Equilibria of the Generalized Second-Price Auction," *Games and Economic Behavior*, 86, 421-437
19. Graham, D. and R. Marshall (1987). "Collusive Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, 95, 1217-1239.
20. Harrington, J. E., A. Scrypcacz, (2007). "Collusion under Monitoring of Sales," *RAND Journal of Economics*, 38(2), pp. 314-331.

21. Harrington, J. E., A. Scrkypacz, (2011). "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101(6), 2425-49.
22. Hendricks, K., R. Porter and G. Tan (2008). "Bidding Rings and the Winner's Curse," *RAND Journal of Economics*, 39, 1018-1041.
23. Levin, J., A. Scrkypacz, (2016). "Properties of the Combinatorial Clock Auction" (2016) with Jonathan Levin," *American Economic Review* 106(9), pp. 2528-2551.
24. Mailath, G., P. Zemski, (1991). "Collusion in Second Price Auctions with Heterogenous Bidders," *Games and Economic Behavior*, 3, 467-420.
25. McAfee, P. and J. MacMillan, (1992). "Bidding Rings," *American Economic Review*, 82, 579-599.
26. Milgrom, P., J. Mollner, (2014). "Equilibrium Selection in Multi-Player Games with Auction Applications," mimeo.
27. Ottaviani, M., (2003). "Overture and Google: Internet Pay-Per-Click (PPC) Advertising Auctions," London Business School - Case Study-03-022
28. Pricewaterhouse Coopers, (2015). "Global Entertainment Media Outlook," accessed on April 2015 at: <http://www.pwc.com/gx/en/global-entertainment-media-outlook/segment-insights/internet-advertising.jhtml>
29. Ray, D., R. Vohra, (1997). "Equilibrium Binding Agreements," *Journal of Economic Theory*, 73, 30-78.
30. Ray, D., R. Vohra, (2014). "The Farsighted Stable Set," *Econometrica*, forthcoming.
31. Varian, H., (2007). "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163-1178.