Marketing Agencies and Collusive Bidding in Online Ad Auctions

Francesco Decarolis, Maris Goldmanis, Antonio Penta*

April 7, 2017

Abstract

The transition of the advertising market from traditional media to the internet has induced a proliferation of marketing agencies specialized in bidding in the auctions that are used to sell ad space on the web. We analyze how collusive bidding can emerge from bid delegation to a common marketing agency and how this can undermine the revenues and allocative efficiency of both the Generalized Second Price auction (GSP, used by Google and Microsoft-Bing and Yahoo!) and the of VCG mechanism (used by Facebook). We find that, despite its well-known susceptibility to collusion, the VCG mechanism outperforms the GSP auction both in terms of revenues and efficiency.

JEL: C72, D44, L81.

Keywords: Collusion, Digital Marketing Agencies, Facebook, Google, GSP, Internet Auctions, Online Advertising, VCG.

*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, Jean-Pierre Benoit, Nicholas Bloom, Yeon-Koo Che, Kerry Gabrielson, 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, London Business School, 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

Online advertising is the main source of revenues for important firms such as Google, Facebook, Twitter, etc., and it represents one of the largest and fastest growing industries in the US: in 2013, for instance, the value of advertising on search engines alone amounted to 50 billion dollars in the U.S., with an annual growth of 10% (PwC (2015)), and 96% of Google's global revenues in 2011 were attributed to advertisement (Blake, Nosk and Tadelis (2015)). Almost all online ads are sold through auctions, in which bidders compete for the adjudication of one of a given number of 'slots' available in various online venues, such as search engine result pages, social networks feeds, and so on. With the significant exception of Facebook, which recently adopted the Vickerey-Clarke-Groves (VCG) mechanism, for a long time this market has been dominated by the Generalized Second Price (GSP) auction (used, for instance, by Google, Microsot-Bing, Yahoo! and Taobao).

The VCG is a classic and widely studied mechanism: it involves fairly complex payments that price externalities, but it has the advantage of being strategy-proof and efficient. The GSP auction in contrast has very simple rules (the k-highest bidder obtains the k-highest slot at a price-per-click equal to the $(k+1)$-highest bid), but it gives rise to complex strategic interactions. Varian (2007) and Edelman, Ostrovsky and Schwarz (2007, EOS) pioneered the study of the GSP auction. Their results provided a rationale for the GSP's striking success and, until recently, its almost universal diffusion. But these models do not account for a recent trend in this market, which is bound to to alter the functioning of these auctions and has thus the potential to shake up the entire industry.

We allude to the fact that, at least since 2011, an increasing number of advertisers are delegating their bidding campaigns to specialized digital marketing agencies (DMAs), many of which belong to a handful of networks (seven in the US) that conduct all bidding activities through centralized agency trading desks (ATDs).[1] As a result, with increasing frequency, the same entity (be it DMA or ATD) bids in the same auction on behalf of different advertisers. But this clearly changes the strategic interaction, as these agencies have the opportunity to lower their payments by coordinating the bids of their clients.

This paper proposes a theoretical analysis of the impact of agency bidding on the two main auction formats: the VCG and the GSP. We find that, in the presence of marketing agencies, the VCG outperforms the GSP both in terms of revenues and efficiency. This is a strong result because the VCG is well-known to be highly susceptible to collusion, but it is especially noteworthy if one considers the sheer size of transactions currently occurring under the GSP. It also suggests a rationale for Facebook's recent adoption of the VCG mechanism, which – despite the early surprise it provoked (e.g., *Wired* (2015)) –

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

has proven remarkably successful. The striking fragility of the widespread GSP auction we uncover suggests that further changes are likely to occur in this industry, raising important questions from a market-design perspective. But since agencies' behavior in our model is analogous to that of buying consortia, which have been sanctioned in the past, our results are also relevant from an antitrust perspective.[2] (To the best of our knowledge, this is the first study to point at the central role of marketing agencies in this market.) The specificities of the market, however, suggest a more nuanced view of the harm to consumers. We discuss this and other policy implications in the conclusions.

The study of agency bidding in the GSP auction presents numerous difficulties. First, it is important to develop a model in which collusive and competitive behavior coexist, because agencies in these auctions operate side by side with independent advertisers. But the problem of 'partial cartels' is acknowledged as a major difficulty in the literature (e.g., Hendricks, Porter and Tan (2008)).[3] Second, strategic behavior in the GSP auction is complex and brings forth a plethora of equilibria (Borgers et al. (2013)). Introducing a tractable refinement has been a key contribution of EOS and Varian (2007), to cut through this complexity and bring out the economics of these auctions.[4] But their refinement is not defined in the agency model. Thus, a second challenge we face is to develop a model of agency bidding that is both tractable and ensures clear economic insights.

To achieve these goals, we modify EOS and Varian's baseline model by introducing a marketing agency, which we model as a player choosing bids for its clients in order to maximize the total profits. Bidders that do not belong to the agency are referred to as 'independents', and have the usual objectives. To overcome the curse of multiplicity in the GSP auction, and ensure a meaningful comparison with the competitive benchmark, we introduce a refinement of bidders' best responses that distills the individual-level underpinnings of EOS' equilibrium, and assume that independents place their bids accordingly. This stratagem enables us to maintain the logic of EOS' refinement for the independents, even if their equilibrium is not defined in the game with collusion. The marketing agency 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 of others, and taking into account the possible unraveling of the rest of the coalition, no client has an incentive to abandon the agency and bid as an independent. Hence, the outside options of the coalition's members are equilibrium objects themselves, and implicitly

[2]See, 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).*

[3]The 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 has assumed 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), or in the dynamic auctions of Chassang and Ortner (2016)). The main focus of that literature is on the coalition members' incentives to share their private information so as to implement collusion (see also Che and Kim (2006, 2009) and Che et al. (2016)), a moot point under complete information, as EOS, Varian's (2007) and our settings. Other mechanisms for collusion have been considered, for instance, by Harrington and Skrzypacz (2007, 2011).

[4]On a similar note, by Levin and Skrzypacz (2016) strike a fine balance between tractability and realism of the assumptions, to deliver clear economic insights on an otherwise very complex auction.

incorporate the restrictions entailed by the underlying coalition formation game. The logic of our model is therefore closely related to the idea of 'equilibrium binding agreements' (Ray (2008)), in that it involves both equilibrium and recursive stability restrictions.

We consider different models of collusive bidding within this general framework. First, we assume that the agency is constrained to placing bids that cannot be detected as collusive by an external observer, such as an antitrust authority or the auction platform. We show that, under this constraint, the GSP auction is efficient and its revenues are identical to those obtained if the same coalition structure (viz., agency) bid in a VCG auction. We then relax this 'undetectability constraint', and show that in this case the revenues in the GSP auction are never higher, and are in fact typically lower, than those obtained in the VCG mechanism with the same agency configuration. Furthermore, once the 'undetectability constraint' is lifted, efficiency is no longer guaranteed by the GSP. Since the VCG is well-known to be highly susceptible to collusion, finding that it outperforms the GSP both in terms of revenues and efficiency is remarkably negative for the GSP auction.

The source of the GSP's fragility, and the complexity of agency bidding in this context, can be understood thinking about an agency that controls the first, second, and fourth highest bidders in an auction. The agency in this case can lower the highest bidder's payment by shading the bid of the second, without necessarily affecting either his position or his payment. Given the rules of the GSP auction, the agency can benefit from this simple strategy only if two of her members occupy adjacent positions. But due to the GSP's complex equilibrium effects, the agency can do more than that. For instance, suppose that the agency shades the bid of her lowest member, with no direct impact on her other clients' payments. Intuitively, if this bid is kept persistently lower, then the logic of EOS' refinement suggests that the third highest bidder, who is an independent, would eventually lower his bid. But not only would this lower the second bidder's payment, it would also give the agency extra leeway to lower the second highest bid, to the greater benefit of the highest bidder. Revenues in this case diminish for both the *direct effect* (lowering the 2-nd highest bid lowers the highest bidder's payment) and for the *indirect effect* (lowering the 4-th highest bid induces a lower bid for the independent, which in turn lowers the second bidder's payment). Hence, even a small coalition may have a large impact on total revenues. Our general results show that this impact is larger if the agency includes members which occupy low or adjacent positions in the ranking of valuations, but it also depends on the rate at which click-through-rates vary from one position to another, and on how independents' valuations compare to those of the coalition members.

We also explore whether these concerns on the GSP auction may be mitigated by competition between agencies. Although multiple agencies each with multiple bidders in the same auction are rare (Decarolis et al. (2016)), the question has theoretical relevance because the phenomenon may become more common in the future. If an increase in agency competition restored the good properties of these auctions, then the diffusion of marketing agencies need not lead to major structural changes in this industry. Our results, however, suggest otherwise: for certain coalition structures, agency competition as expected mitigates the revenue losses in both mechanisms (while preserving their relative performance); but for other coalition structures, it has a particularly perverse impact on both mechanisms. That is because, from the viewpoint of an agency bidding for multiple clients, these auction mechanisms have a flavor of a first-price auction: even holding positions constant, the price paid depends on the agency's own bids. With multiple agencies, this feature of agency bidding may lead to non-existence of pure equilibria, very much like the case of competitive (non-agency) bidding in a Generalized First Price (GFP) auction. But as seen in the early days of this industry, when the GFP was adopted (see Section 2), lack of pure equilibria may generate bidding cycles which eventually lead to a different form of collusion. In fact, these bidding cycles are often cited as the primary cause for the transition, in the early '00s, from the GFP to the GSP auction (Edelman and Schwarz (2007)). Hence, not only does agency competition not solve the problems with these auctions, but it appears likely to exacerbate them, giving further reasons to expect fundamental changes in this industry.

The rest of the paper is organized as follows: Section 2 provides a brief history of the market and illustrates the basic stylized facts that motivate our model. Section 3 reviews the competitive benchmarks. Section 4 introduces the model of collusion, and Section 5 presents the main results. Section 6 develops a method for detecting collusion in search auctions data and to quantify the revenue losses. Section 7 discusses the main policy implications of our results and directions for future research.

2 A brief history of the online ad market

In 1998, the search engine GoTo.com revolutionized the world of online advertising by introducing auctions to sell ad space on its search results pages. This company, later renamed Overture and acquired by Yahoo! in 2001, had devised the so called Generalized First Price (GFP) auction, in which advertisement space was assigned to advertisers by the ranking of their bids, with each advertiser paying his own bid for each click he received. But as Yahoo!'s auctions grew in volume, and advertisers became acquainted with their operation, this initially very successful model became problematic (cf. Ottaviani (2003)). The reason is that, after an initial period in which advertisers cycled through phases of aggressive and conservative bidding, their bids eventually settled at very low levels, with the GFP indirectly favoring the diffusion of collusive bidding strategies. This phenomenon, later attributed to the lack of pure equilibria in the GFP auction (Edelman and Schwarz (2007)), led to the creation of a new auction format, which would soon dominate this market: the Generalized Second Price (GSP) auction.

In February 2002, Google introduced the GSP as part of its *AdWords Select* bidding platform. Key to Google's success was the ability to incorporate advertisement in the clean layout of its pages, without diluting the informative content for the consumers (cf.

Wu (2016)).[5] But the strategic structure of the GSP, as well as the simplicity of its rules, turned out to be fundamental to ensure stable bidding behavior, and hence a solid revenue base, which boosted Google's business in an unprecedented way: on August 19th, 2004, Google went public with a valuation of $27 billion. In 2011, the company registered $37.9 billion in global revenues, of which $36.5 billion (96%) were attributed to advertising (Google Inc., Blake et al. (2015)). Google is now worth close to $300 billion. Google's success turned the GSP into the mechanism of choice of all other major search engines, including earlier incumbent Yahoo!, its subsequent partner Microsoft-Bing, and Taobao in China. The GSP's supremacy among online ad auctions went essentially undisputed, until recently, when another major player in the industry attempted an alternative route.

In 2015, Facebook introduced the VCG for its own display ad auctions. These display ad auctions are different from those of search engines, in that they are not generated by keywords and raise specific challenges to integrate ads within Facebook's organic content. But these technicalities aside, they boil down to the same kind of economic problem: a multi-unit auction. Before John Hegeman, a Stanford economics MA graduate, took the role of Facebook's chief economist, the (multi-unit) VCG had had a limited impact outside of academia. Perhaps for this reason, or for the somewhat byzantine VCG payment rule, the industry's initial reaction was one of surprise (cf. *Wired* (2015)). But Facebook and its VCG auction are now essential parts of this industry: in the second quarter of 2015, Facebook pulled in $4.04 billion and, together with Twitter, it has become one of the largest players in display ad auctions. Together, sponsored search and display ad auctions represent nearly the entirety of how online ads are sold.

Alongside the evolution of auction platforms, this market witnessed profound changes on the advertisers' side as well. In the early days of online ad auctions, advertisers bid through their own individual accounts, often managed separately across platforms. But already back in 2011 (see footnote 1), a large share of advertisers in the US delegated their bidding activities to specialised digital marketing agencies (DMAs), whose diffusion quickly led to the issue of common agency discussed in the introduction. The case of Merkle, one of the major agencies in the U.S., provides a clear example of this phenomenon. Crucially for our purposes, many of Merkle's clients operate in the same industries, and are therefore likely to bid on the same keywords.[6] For instance, data from Redbook (the leading public database to link advertisers to their agencies) confirm that Merkle managed the campaigns of two leading charities in 2016, *Habitat for Humanity* and *Salvation Army*, both of which were bidding in the same auctions for hundreds of keywords.[7] Table 1 reports

[5] In the seminal paper which marked the birth of Google, its founders Sergey Brin and Larry Page complain that earlier advertising-funded search engines were "inherently biased towards the advertisers and away from the needs of consumers" (Brin and Page (1998)), which they deemed a major pitfall. The concern for building and maintaining a long-lasting consumer base is a central concern in Google's history, which is also reflected in the introduction of 'quality scores' in the payment rule of its ad auctions (see Section 6). Wu (2016) provides a thorough account of the history of the advertising industry.

[6] See: https://www.merkleinc.com/who-we-are-performance-marketing-agency/our-clients.

[7] Similar examples can be identified for nearly every industry: for electronics, *HP* and *IBM/Lenovo* use iProspect as their agency; for clothing, *Urban Outfitters* and *Eddie Bauer* use Rimm-Kaufman; for

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

Table 1: CPC is the average cost-per-click in $US. Volume is the number of monthly searches, in thousands. Position refers to rank among paid search links on Google's results page for the relevant keyword. Source: 2016 US Google sponsored search data from SEMrush.

the top six of these keywords, in terms of their average cost-per-click (CPC).

The common agency problem is made even more relevant by yet another recent phenomenon, the formation of 'agency trading desks' (ATDs). While several hundred DMAs are active in the US, most of them belong to one of the seven main agency networks (Aegis-Dentsu, Publicis Groupe, IPG, Omnicom Group, WPP/Group M, Havas, MDC), which operate through their corresponding ATDs (respectively: Amnet, Vivaki, Cadreon, Accuen, Xaxis, Affiperf and Varick Media). ATDs' importance is growing alongside another trend in this industry, in which DMAs also play a central role. That is, the ongoing shift towards the so called 'programmatic' or 'algorithmic' real time bidding: the algorithmic purchase of ad space in real time over all biddable platforms through specialized software. ATDs are the units that centralize all bidding activities within a network for 'biddable' media like Google, Bing, Twitter, iAd, and Facebook. Hence, while DMAs were originally not much more sophisticated than individual advertisers, over time they evolved into more and more sophisticated players, and their diffusion and integration through ATDs has made the issue of common agency increasingly frequent.

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

In July 2016, Aegis-Dentsu acquired Merkle, which was not previously affiliated to any network. At that time, many of Merkle's clients were bidding on the same keywords as some of Aegis-Dentsu's advertisers.[8] This acquisition therefore further increased the po-

pharmaceuticals, *Pfizer* and *Sanofi* use Digitas; etc. (Source: Redbook.)

[8]For instance, in the electronics sector, *Dell* and *Samsung* were in Merkle's portfolio, placing bids on keywords also targeted by Aegis-Dentsu's clients *Apple*, *HP*, *IBM/Lenovo* and *Intel*. Other examples

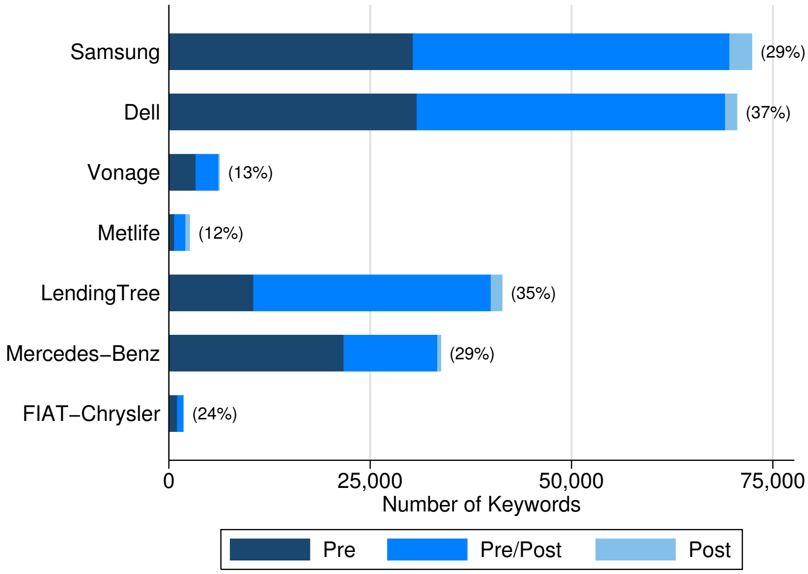

Figure 1: For each of Merkle's advertisers in footnote 8, the figure represents the number of keywords on which it bid alongside at least one member of the Aegis-Dentsu network (and as a share of the total number of keywords on which it bid, in parenthesis) between June 2015 and January 2017 (Merkle's acquisition by Aegis-Dentsu was in July 2016). The graph shows whether bids on these 'shared' keywords occurred only pre-acquisition (dark blue: all keywords appearing only before July 2016), only post-acquisition (turquoise: all keywords appearing only after July 2016), or both pre- and post-acquisition (blue: all keywords appearing both before and after July 2016.) Source: keyword-level data provided by SEMrush.

tential for coordinated bidding. Figure 1 reports, for each of Merkle's advertisers listed in footnote 8, the fraction of the total keywords on which they were bidding at the same time as some of Aegis-Dentsu's clients, and whether joint targeting of such keywords happened only pre-acquisition, only post-acquisition, or both pre- and post-acquisition. Although there is some variation among these advertisers, we clearly see that shared keywords are a quantitately large phenomenon also post-acquisition (interestingly, a small fraction of keywords are shared *only* post-acquisition). Hence, coordinated bidding through a common agency in the same auction – the focus of our model – is clearly a relevant phenomenon.

3 Competitive Bidding in Online Ad Auctions

Stripped down to their essence, online ad auctions are mechanisms to solve the problem of assigning agents $i \in I = \{1, \ldots, n\}$ to slots $s = 1, \ldots, S$, where $n \geq S$. In our case, agents are advertisers, and slots are positions for ads in a given webpage (e.g., on a social

include: in the financial sector, Merkle's *Lending Tree* and *Metlife* were bidding in auctions alongside Aegis-Dentsu's *Capitalone, Discover, Fidelity, Equifax, JP Morgan-Chase*; for car manufacturers, Merkle's *FIAT-Chrysler* and *Mercedes-Benz USA* bid alongside Aegis-Dentsu's *Toyota, Volkswagen, Subaru*; in phone services, Merkle's *Vonage* bid alongside Aegis-Dentsu's *T-Mobile.* (Source: Redbook.)

media's newsfeed for a certain set of cookies, on a search-engine result page for a given keyword, etc.). Slot $s = 1$ corresponds to the highest/best position, and so on until $s = S$, which is the slot in the lowest/worst 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 $x^1 > x^2 > \cdots > x^S > 0$. 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 $v_1 > v_2 > \cdots > v_n$. As in Varian (2007) and EOS, we maintain that valuations and CTRs are common knowledge. Although it may seem unrealistic, this complete information assumption has been shown to be an effective modeling proxy for these settings (e.g., Athey and Nekipelov (2012), Che et al. (2013) and Varian (2007)).[9]

3.1 Rules of the auctions

Both in the VCG and in the GSP auction, advertisers submit bids $b_i \in \mathbb{R}_+$, and slots are assigned according to their ranking: first slot to the highest bidder, second slot to the second-highest bidder, 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).[10] 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^t_{-i} denote the t-highest component of the vectors b and b_{-i}, respectively. Hence, with this notation, for any profile b, in either mechanism bidder i obtains position $\rho(i)$ if $\rho(i) \leq S$, and no position otherwise.[11] The resulting utility, ignoring payments, is thus $v_i x^{\rho(i)}$.

The GSP and VCG mechanisms only differ in their payment rule. In the GSP mechanism, the k-highest bidder gets position k and pays a price-per click equal to the $(k+1)$-highest bid. Using our notation, given a profile of bids b, agent i obtains position $\rho(i)$ and pays a price-per-blick equal to $b^{\rho(i)+1}$. Bidder i's payoff in the GSP auction, given a bids profile $b \in \mathbb{R}^n_+$, can thus be written as $u^{\mathcal{G}}_i(b) = \left(v_i - b^{\rho(i)+1}\right) x^{\rho(i)}$.

In the VCG auction, an agent pays the total allocation externality he imposes on others. In this setting, if the advertiser in position k were removed from the auction, all bidders below him would climb up one position. Hence, if other bidders are bidding truthfully (i.e., $b_j = v_j$, as will be the case in equilibrium), the total externality of the k-highest bidder is equal to $\sum_{t=k+1}^{S+1} b^t(x^{t-1} - x^t)$. We can thus write i's payoff in the VCG mechanism, given a bids profile $b \in \mathbb{R}^n_+$, as $u^{\mathcal{V}}_i(b) = v_i x^{\rho(i)} - \sum_{t=\rho(i)+1}^{S+1} b^t(x^{t-1} - x^t)$.

In the rest of this section we review known results on the competitive benchmarks for these two mechanisms. The only original result will be Lemma 1, which provides an alternative characterization of EOS' *lowest envy-free equilibrium* of the GSP auction.

[9] For an independent private values model, see Gomes and Sweeney (2014). Borgers et al. (2013) maintain the complete information assumption, but consider a more general model of CTRs and valuations.

[10] Formally, $\rho(i; b) := |\{j : b_j > b_i\} \cup \{j : b_j = b_i \text{ and } j < i\}| + 1$. 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).

[11] In 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 6 (cf. Varian (2007)).

3.2 Equilibria

As we mentioned in the introduction, despite the relative complexity of its payment rule, bidding behavior in the VCG is very simple, as truthful bidding (i.e., $b_i = v_i$) is a dominant strategy in this game. In the resulting equilibrium, advertisers are efficiently assigned to positions. The VCG mechanism therefore is efficient and strategy-proof.

Equilibrium behavior in the GSP auction is much more complex. To see this, first note that a generic profile of bids for i's opponentes, $b_{-i} = (b_j)_{j \neq i}$, 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 price-per-click (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}.[12] 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 of the GSP auction.

The GSP auction has many equilibria. For this reason, EOS introduced a refinement of the equilibrium correspondence, the *lowest envy-free equilibrium*, which was crucial to cut through the complexity of the GSP auction. We consider instead a refinement of individuals' best response correspondence: for any $b_{-i} \in \mathbb{R}_+^{n-1}$, let

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

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, given valuations v, are denoted as $E^*(v)$.

Lemma 1 *For any profile of valuations $v = (v_i)_{i=1,\dots,n}$, and for any $b \in E^*(v)$, $b_1 > b_2$, $b_i = v_i$ for all $i > S$, and for all $i = 2, \dots, S$,*

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

Hence, the fixed points of the BR^* correspondence coincide with EOS' lowest envy-free equilibrium*, and it induces the same allocation and payments as in the dominant strategy equilibrium of the VCG mechanism.*

This lemma shows that EOS' 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. Hence, BR_i^* distills the individual level underpinnings of EOS' equilibrium. In Section 5 we will assume that independents in the GSP

[12] Allowing ties in individuals' bids or non-generic indifferences complicates the notation, without affecting the results and the main insights. See Appendix **??** for details on this.

auction bid according to BR_i^*, and play their dominant strategy in the VCG, 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 also 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)). Finally, 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.[13]

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

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 *lowest envy-free equilibrium* 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 $b_1 > b_2$ 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. $\square$

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

Corollary 1 *For any $b \in E^*(v)$, for all $i = 2, \ldots, S$:*

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

4 A Model of Agency Bidding

Our analysis of marketing agencies 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 agency, and postpone the analysis of the multiple agency case to Section 5.3.

Modeling coordinated bidding, it may seem natural to consider standard solution concepts such as strong Nash (Aumann (1959)) or coalition proof equilibrium (Bernheim et al. (1987)). Unfortunately, these concepts have no bite in the GSP auction, as it can be

[13] See, for instance, the Google AdWord tutorial in which Hal Varian teaches how to maximize profits by following this bidding strategy: http://www.youtube.com/watch?v=jRx7AMb6rZ0. Borgers et al. (2013) provide a more critical view of Varian and EOS' refinement. Nonetheless, those refinements are the established benchmark in the literature, and hence our modeling choice enables us to focus on the issue of agency bidding while allowing a meaningful comparison with the competitive benchmark.

shown that EOS' equilibrium satisfies both refinements.[14] As EOS showed, resorting to non-standard concepts is a more promising route to get some insights into the elusive GSP auction. We thus model the marketing agency 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 her proposals can be implemented only if they are *stable* in two senses: **(S.1)** first, if they are consistent with the independents' equilibrium behavior, which in turn is defined as the fixed-point of the same refinements of the individual-best responses used in the competitive benchmarks (i.e., truthful bidding in the VCG, and BR_i^* in the GSP); **(S.2)** second, if no individual member of the coalition has an incentive to abandon it and bid 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 (2008)). Hence, given a coalition C, the outside option for each member $i \in C$ is his equilibrium payoff in the game with coalition $C \backslash \{i\}$, in which i bids as an independent. The constraint for coalition C thus depends on the solutions to the problems of all the subcoalitions $C' \subseteq C$, and hence the solution concept for the game with the agency will be defined recursively. We thus call it the *'Recursively-Stable Agency Equilibrium' (RAE).*

Before getting into the intricacies of agency bidding in the GSP auction, and in the formal definition of RAE for general mechanisms, we illustrate its basic logic in the context of the simpler VCG mechanism.

4.1 RAE in the VCG: Informal Explanation

We begin by considering an example of RAE in the VCG mechanism. In the example, as well as in some results in Section 5, equilibrium bids will sometime be such that $b_i = b_{i+1}$ for some i. Since ties are broken according to bidders' labels (cf. footnote 10), in that case bidder i obtains the position above $i+1$. To emphasize this, we will write $b_i = b_{i+1}^+$.[15]

Example 2 Consider an environment with five bidders who compete for the allocation of four slots sold through the VCG mechanism. Bidders' valuations are $v = (40, 25, 20, 10, 9)$, and the CTRs are $x = \{20, 10, 9, 1, 0\}$. As discussed in Section 3, in this mechanism

[14] These standard solution concepts therefore fail to capture any difference between competitive and collusive bidding in the GSP auction. On the other hand, we envision bid delegation to a common agency as more than just a channel for non-binding communication, which is the focus of those concepts.

[15] Without the tie-breaking rule embedded in ρ (footnote 10), the agency's best replies may be empty valued. In that case, our analysis would go through assuming that bids are placed from an arbitrarily fine discrete grid (i.e., $A_i = (\mathbb{R}_+ \cap \varepsilon\mathbb{Z})$ where ε is the minimum bid increment). In that setting, $b_i = b_{i+1}^+$ 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 $\varepsilon \to 0$ (the notation b_{i+1}^+ is thus reminiscent of this alternative interpretation, as the right-hand limit $b_{i+1}^+ := \lim_{\varepsilon^+ \to 0} (b_{i+1} + \varepsilon)$).

advertisers bid truthfully in the competitive benchmark, and hence equilibrium payoffs for the five bidders are $u^{Comp} = (441, 141, 91, 1, 0)$.

Now consider a setting in which bidders 1 and 5 belong to the same agency, $C' = \{1, 5\}$, and everyone else is an independent. Bidding truthfully remains a dominant strategy for the independents, but clearly this is not the case for the agency: since 1's payment is strictly decreasing in b_5, it is clear that bidding $(b_1, b_5) = (40, 0)$ is a profitable deviation from truthful bidding for the agency. In fact, it is not difficult to see that this bid profile is optimal for the agency: given the bids of the independents, there would be no benefit in lowering b_1 to the point of losing the highest position, nor in increasing b_5 so as to obtain a higher slot. So, holding constant the allocation, the optimal solution for the agency is to lower b_5 as much as possible, while maintaining $b_1 > b_2 = 25$. Hence, any profile $b' = (b'_1, 25, 20, 10, 0)$ such that $b'_1 > 25$ is an equilibrium, and the resulting payoffs are $u' = (450, 150, 100, 10, 0)$, with a total 450 for the coalition. Comparing u' with u^{Comp}, it is also clear that no member of the coalition would rather bid as an independent.

Next, suppose that the coalition also includes bidder 2: $C'' = \{1, 2, 5\}$. We next show that in this case the RAE-bids are $b'' = (b''_1, 20^+, 20, 10, 0)$, where $b''_1 > 20$, which induce payoffs $u'' = (500, 150, 100, 10, 0)$ and a total of 650 for the coalition. To see that this is a RAE, recall that truthful bidding is still dominant for the independent bidders. The argument for keeping $b''_5 = 0$ and $b''_1 > 20$ are the same as above. So, let's focus on the agency-optimal positioning of b_2. First note that, if the agency set $b_2 = 10^+$, pushing bidder 2 down to the third slot, then the coalition payoff would be 655, which is higher than 650, as in our candidate RAE. But in that profile, 2's payoff would be 145, lower than $u'_2 = 150$, which he could obtain if he left the coalition and bid as an independent in the game with $C' = \{1, 5\}$. Hence, lowering b_2 to the point of obtaining a lower position, would increase the overall coalition payoff (by decreasing bidder 1's payment), but would violate the stability constraint (S.2) for bidder 2, who in that case would rather abandon the coalition and bid as an independent. The optimal b''_2 therefore is the lowest bid which ensures bidder 2 maintains the second position.[16] □

Note that the recursive definition of the outside option matters in this example. If outside options were defined with respect to the competitive case, bidder 2 would remain in the coalition even when forced to take the lower position, since his payoff in the competitive benchmark are $u_2^{Comp} = 141 < 145$. But we find it unreasonable to model 2's outside option this way: why would an agency client assume that, were he to abandon the agency, the entire coalition would be disrupted and full competition restored? Hence, while it will necessarily require a more involved definition, the recursivity of the stability constraint

[16] This argument also shows that the RAE-profile $b'' = (40, 20^+, 20, 10, 0)$ is not a Nash equilibrium of the game in which C'' is treated as a single player, nor a 'plausible' refinement of the original game, as bidders 2 and 5 play weakly dominated strategies. The example's result also relies on the fact that direct transfers are ruled out in our model. If transfers were allowed, the impact of collusion would be even stronger. Our results can thus be seen as a conservative assessment of the impact of collusion. Che et al. (2016) discuss other arguments for the no-transfers assumption in general settings.

for the coalition members captures an important aspect of the environments we attempt to model, and poses economically meaningful restrictions on the agency's freedom to manipulate the bids of its clients.

Our approach also addresses several 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 (cf. footnote 3); (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 (see Section 6).

We conclude this discussion by noting that an obvious alternative to our approach would be to model bidders' choice to join the agency 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. 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 digital marketing agencies, it is not obvious which should be preferable.[17] Independent of these modeling choices, however, the cost of joining the agency would ultimately have to be traded-off against the benefit, which in turn presumes solving for the equilibrium for a *given* coalition structure. Our work can thus be seen as a necessary first step in developing a full-blown model of agency formation.

The next subsection contains the formal definition of the 'Recursively Stable Agency Equilibrium', which allows for arbitrary underlying mechanisms. This is useful in that it provides a unified framework to analyze the impact of marketing agencies under different mechanisms. Section 5 contains the analysis for the GSP and VCG mechanisms.

4.2 The Recursively Stable Agency Equilibrium: General Definition

Let $G = (A_i, u_i)_{i=1,\ldots,n}$ denote the baseline game (without a coalition) generated by the underlying mechanism (e.g., GSP or VCG). We let $\mathcal{C}$ denote the collection of all sets $C \subseteq I$ such that $|C| \geq 2$. For any $C \in \mathcal{C}$, we let C denote the agency, and we refer to advertisers $i \in C$ as 'members of the coalition' and to $i \in I \backslash C$ as 'independents'. The coalition chooses a vector of bids $b_C = (b_j)_{j \in C} \in \times_{j \in C} A_j$. Given b_C, the independents

[17] Moreover, costs need not be symmetric, and hence it may be that an advertisers is willing to join the coalition, but current members are better-off without him. Whereas the decision to abandon an agency is unilateral, the decision to join it is not, raising further modeling questions.

$i \in I \backslash C$ simultaneously choose bids $b_i \in A_i$. We let $b_{-C} := (b_j)_{j \in I \backslash C}$ and $A_{-C} := \times_{j \in I \backslash C} A_j$. Finally, given profiles b or b_{-C}, we let $b_{-i,-C}$ denote the subprofile of bids of all independents other than i (that is, $b_{-i,-C} := (b_j)_{j \in I \backslash C : j \neq i}$).

We assume that the agency maximizes the sum of its members' payoffs,[18] denoted by $u_C(b) := \sum_{i \in C} u_i(b)$, 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 $R_C \subseteq A_C$, allows us to accommodate the possibility that the agency may exogenously discard certain bids. For instance, in Section 5.2.1 we will consider the case of an agency whose primary concern is not being identified as inducing collusion. In that case, R_C would be comprised of only those bids that are 'undetectable' to an external observer as collusive. We denote the collection of exogenous restrictions for all possible coalitions as $\mathcal{R} = \{R_C\}_{C \in \mathcal{C}}$.

Stability-1: The first stability restriction on the agency's proposals requires that they are stable with respect to the independents. For any $i \in I \backslash C$, let $BR_i^* : A_{-i} \rightrightarrows A_i$ denote some refinement of i's best response correspondence in the basline game G (e.g., truthful bidding in the VCG, or (1) in the GSP). Define the *independents' equilibrium correspondence* $BR_{-C}^* : A_C \rightrightarrows A_{-C}$ as

$$BR_{-C}^*(b_C) = \left\{ b_{-C} \in A_{-C} : \forall j \in I \backslash C,\ b_j \in BR_j^*(b_C, b_{-j,-C}) \right\}. \tag{4}$$

If the agency proposes a profile b_C that is not consistent with the equilibrium behavior of the independents (as specified by BR_{-C}^*), 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

$$S_C = \left\{ b_C \in A_C : \exists b_{-C} \text{ s.t. } b_{-C} \in BR_{-C}^*(b_C) \right\}. \tag{5}$$

Clearly, the strength of this constraint in general depends on the underlying game G and on the particular correspondence BR_{-C}^* that is chosen to model the independents' behavior. This restriction is conceptually important, and needed to develop a general framework for 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_C, the agency forms conjectures about how its bids would affect the bids of the independents. We let $\beta : S_C \to A_{-C}$ represent such conjectures of the agency. For any profile $b_C \in S_C$, $\beta(b_C)$ denotes the agency's belief

[18] This 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 agency 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 agency's ability to generate surplus for its clients is an important determinant of its long run profitability.

about the independents' behavior, if she chooses profile b_C. It will be useful to define the set of conjectures β that are consistent with the independents playing an equilibrium:

$$B^* = \left\{\beta \in A_{-C}^{S_C} : \beta(b_{\mathcal{C}}) \in BR_{-C}^* \text{ for all } b_C \in S_C\right\}. \tag{6}$$

The second condition for stability requires that, given conjectures β, no client of the agency has an incentive to leave and bid as an independent. Hence, the outside option for coalition member $i \in C$ is determined by the equilibrium outcomes of the game with coalition $C\backslash\{i\}$. This constraint thus requires a recursive definition.

First, we let $E^* = \left\{b \in \mathbb{R}_+^n : b_i \in BR_i^*(b_{-i}) \text{ for all } i \in I\right\}$ denote the set of equilibria in the game without coalition, given refinement BR_i^*. Letting $E^{\mathcal{R}}(C')$ denote the set of *Recursively Stable Agency Equilibrium (RAE)* outcomes of the game with coalition C', given restrictions $\mathcal{R}$ (and refinement BR_i^*), we initialize the recursion setting $E^{\mathcal{R}}(C') = E^*$ if $|C'| = 1$ (that is, if an agency controls only one bidder, then the RAE are the same as the competitive equilibria). Suppose next that $E^{\mathcal{R}}(C')$ has been defined for all subcoalitions $C' \subset C$. For each $i \in C$, and $C' \subseteq C\backslash\{i\}$, let $\bar{u}_i^{C'} = \min_{b \in E^{\mathcal{R}}(C')} u_i(b)$. Then, recursively:

Definition 1 *A* Recursively Stable Agency Equilibrium (RAE) *of the game G with coalition C, given restrictions $\mathcal{R} = \{R_C\}_{C \in \mathcal{C}}$ and refinement BR_i^*, is a profile of bids and conjectures $(b^*, \beta^*) \in A_C \times B^*$ such that:*[19]

1. The independents play a best response: for all $i \in I\backslash C$, $b_i^* \in BR_i^*(b_{-i}^*)$.

2. The conjectures of the agency are correct: $\beta^*(b_C^*) = b_{-C}^*$.

3. The agency best responds to conjectures β^*, subject to the exogenous restrictions (R) and the stability restrictions (S.1) and (S.2):

$$\begin{aligned} & b_C^* \in \arg\max_{b_C} u_C(b_C, \beta^*(b_C)) \\ \textit{subject to} & : \textit{(R) } b_C \in R_C \\ & : \textit{(S.1) } b_C \in S_C \\ & : \textit{(S.2) for all } i \in C,\ u_i(b_C, \beta^*(b_C)) \geq \bar{u}_i^{C\backslash\{i\}} \end{aligned}$$

The set of ($\mathcal{R}$-constrained) RAE outcomes *for the game with coalition C is:*

$$E^{\mathcal{R}}(C) = \{b^* \in A : \exists \beta^* \textit{ s.t. } (b^*, \beta^*) \textit{ is a RAE}\}. \tag{7}$$

We will refer to the case in which $\mathcal{R}$ is such that $R_{C'} = A_{C'}$ for all $C' \subseteq I$ as the 'unconstrained' case, and denote the set of unconstrained RAE outcomes as $E(C)$.

In the next section we apply this definition to study agency bidding in the GSP and VCG mechanism. Here we provide some general considerations on the solution concept.

[19] Note 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.

First, as we mentioned in Section 4.1, 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's (1997, 2014, RV) equilibrium binding agreements, the stability restrictions do affect the set of equilibrium outcomes, not merely as a refinement.

Relative to RV, our approach differs mainly in that our stability restriction (S.2) only allows agency proposals to be blocked by individual members, whereas RV allow for any joint deviation of coalition members. That advertisers can make binding agreements outside the agency, and jointly block its proposals, seems unrealistic in this context. A direct application of their concept to this setting therefore seems inappropriate. Also, unlike RV (in which the non-cooperative interaction is based on Nash equilibrium), our definition also allows for refinements. As already explained, this is crucial here, especially for the analysis of GSP auction.

5 Agency Bidding in VCG and GSP: Results

In this Section we specialize the general notion of RAE to the GSP and VCG mechanisms:

Definition 2 (RAE in the GSP and VCG) *Given a set of exogenous restrictions $\mathcal{R}$, the $\mathcal{R}$-constrained RAE of the GSP and VCG mechanisms are obtained from Definition 1 letting G denote the corresponding game, and BR_i^* be defined, respectively, as in (1) for the GSP and as the dominant (truthful) strategy in the VCG.*

We first present the analysis of the VCG mechanism (Section 5.1), and then proceed to the GSP auction (Section 5.2). Our main conclusion is that the VCG outperforms the GSP both in terms of revenues and allocative efficiency, thereby uncovering a striking fragility of the GSP with respect to agency bidding.

5.1 Agency Bidding in the VCG mechanism

Our first result characterizes the *unconstrained RAE* of the VCG mechanism:

Theorem 1 (RAE in the VCG) *For any 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, 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. Formally: in the VCG mechanism, $\hat{b} \in E(C)$ if and only if*

$$\hat{b}_i \begin{cases} = v_i & \text{if } i \in I \backslash C; \\ = \hat{b}^+_{i+1} & \text{if } i \in C \backslash \{\min(C)\} \text{ and } i \leq S; \\ \in \left(\hat{b}^+_{i+1}, v_{i-1}\right) & \text{if } i = \min(C) \text{ and } i \leq S. \end{cases} \tag{8}$$

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's competitive benchmark. Moreover, the presence of a marketing agency has no impact on the bids of the independents, which follows from the strategy-proofness of the mechanism, embedded in the independents' refinement BR_i^*. (This property also ensures that $S_C = A_C$, and hence constraint (S.1) in Def. 1 plays no role in the result.) As we discussed in Section 4.1, the recursive stability restriction (S.2) is key to this result.[20] The proof of Theorem 1 is based on a recursive argument, which 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.

Whereas the presence of an agency does not alter the allocation of the VCG mechanism, it does affect its revenues: in any RAE of the VCG mechanism, the agency 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 $s = 1, ..., \min\{S+1, i-1\}$, even a small coalition can have a significant impact on the total revenues. On the other hand, the VCG's strategy-proofness ensures that the agency has no impact on the independents, which continue using their dominant strategy and bid truthfully. Hence, while an agency may have a large 'direct effect' on revenues, it has no 'indirect effect' in this mechanism.

Example 3 Consider the environment in Example 1, and suppose that $C = \{1, 3\}$. Then, applying the formula in (8), the RAE of the VCG mechanism is $\hat{b} = \left(\hat{b}_1, 4, 2^+, 2, 1\right)$. The resulting revenues are 86, as opposed to 96 of the competitive benchmark. □

5.2 Agency Bidding in the GSP auction

We begin our analysis of the GSP auction by characterizing the RAE when the agency is constrained to placing bids 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. This result is particularly interesting because it characterizes the equilibria in a market in which 'not being detectable as collusive' is a primary concern of the agency, which appears relevant in the data (Decarolis et al. (2016)). It also enables a tractable comparative statics on the impact of agency bidding in the GSP.

We lift the 'undetectable coordination' restriction in Section 5.2.2. We show that, unlike the VCG mechanism, the unrestricted RAE of the GSP auction may be inefficient and induce strictly lower revenues than their VCG counterparts. In light of the VCG's efficiency (Theorem 1), it may be tempting to impute the lower revenues of the GSP

[20] Bachrach (2010), for instance, studies collusion in the VCG mechanism in a classical cooperative setting (i.e. without distinguishing the agency clients from the independents, and without the 'farsightedness' assumption), finding that the VCG is vulnerable to this form of collusion.

auction to the inefficiencies that it may generate. To address this question, in Section 5.2.2 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 are always lower than in the VCG (Theorem 3). The revenue ranking therefore is not a direct consequence of the allocative distortion.

5.2.1 'Undetectable Coordination': A VCG-Equivalence Result

Consider the following set of exogenous restrictions: for any $C \in \mathcal{C}$,

$$R_C^{UC} := \left\{ b_C \in A_C : \exists v' \in \mathbb{R}_+^{|C|}, b_{-C} \in \mathbb{R}_+^{n-|C|} \text{ s.t. } (b_C, b_{-C}) \in E^* \left(v'_C, v_{-C} \right) \right\}.$$

In words, R_C^{UC} is comprised of all bid profiles of the agency that could be observed as part of a competitive equilibrium in the GSP auction, given the valuations of the independents $v_{-C} = (v_j)_{j \in I \setminus C}$. For instance, consider an external observer (e.g., the search engine or the antitrust authority) who can only observe the bid profile, but not the valuations $(v_i)_{i \in C}$. Then, R_C^{UC} characterizes the bid profiles that ensure the agency could not be detected as 'collusive', even if the independents had revealed their own valuations to the external observer. The next result characterizes the RAE of the GSP under these restrictions, and shows its revenue and allocative equivalence to the unrestricted RAE of the VCG:

Theorem 2 *For any 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 $i = n, ..., 1$, recursively define $v_i^f := v_{i+1}^f$ if $i \in C$ and $v_i^f = v_i$ if $i \notin C$. Then, for every i,*

$$\hat{b}_i \begin{cases} = v_i^f - \frac{x^i}{x^{i-1}} \left(v_i^f - \hat{b}_{i+1} \right), & \text{if } i \neq 1 \text{ and } i \neq \min(C); \\ \in \left[v_i^f - \frac{x^i}{x^{i-1}} \left(v_i^f - \hat{b}_{i+1} \right), \hat{b}_{i-1} \right) & \text{otherwise} \end{cases}, \tag{9}$$

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, and advertisers' payments are the same as in the corresponding unrestricted RAE of the VCG mechanism (Theorem 1).

Note that, in this equilibrium, 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. 9). For the independent bidders ($i \notin C$), 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 agency's equilibrium behavior: intuitively, in order to satisfy the UC-restriction, the agency's bids for each of its members should mimic the behavior of an independent in the competitive benchmark, for some valuation. The agency'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.[21] The last line of (9) corresponds to the bid of the highest coalition member and the highest overall bidder, required to be placed in their efficient positions. The resulting allocation is efficient, and it yields 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 (9).

Example 4 Consider the environment of Example 3, with $C = \{1, 3\}$. Then, applying the formula in (9), the UC-RAE is $\hat{b} = \left(\hat{b}_1, 2.9, 1.8, 1.6, 1\right)$, which results in revenues 86. These are the same as in the VCG mechanism (Example 3), 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 $\left(b_3^{EOS} - \hat{b}_3\right) \cdot x_2 = 5$ (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 agency's 'indirect effect' on the independents. □

Thus, despite the simplicity of the payment rule in the GSP auction, the equilibrium effects in (9) 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 agency bidding 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 *The following comparative statics results hold for both the unconstrained RAE of the VCG and in the UC-RAE of the GSP auction:*

1. Holding everything else constant, the revenue losses due to agency bidding increase with the differences $(x_{i-1} - x_i)$ *associated to the agency's clients* $i \in C$.

2. Holding $(x_s - x_{s+1})$ *constant (i.e., if this difference is constant in* s*), the revenue losses are larger if (i) the agency includes members that occupy adjacent positions in the ranking of valuations, or if (ii) the difference in valuations between the agency's clients and the independents immediately below them (in the ranking of valuations) is larger.*

[21] The 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 agency alter the bids placed by the independents.

3. Holding $(x_s - x_{s+1})$ and $(v_s - v_{s+1})$ constant, the revenue losses due to agency bidding are larger if the agency includes members that occupy a lower position in the ranking of valuations.

Part 1 is immediate from Theorem 1 and the transfers of the VCG payment (Section 3). Part 2 is also straightforward: Point (ii) is due to the fact that, for any $i \neq \max\{C\}$, if $i+1$ also belongs to the coalition, then the agency can lower i's bid below v_{i+1}, still preserving an efficient allocation. Point (ii) follows because, the lower the valuations of the independents ranked below a member of the coalition, the more the agency has freedom to lower the bid of that member without violating the efficient ranking of bids (Theorem 1). Part 3 follows from the fact that, holding constant the size of the 'direct effect' (as entailed by the distances between contiguous valuations), the total reduction in revenues of the VCG mechanism increases with the number of agents placed above him.

The equilibrium characterization in Theorem 2 involves bidders' valuations. But since valuations are typically not observable to an external analyst, the conditions in (9) 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 C, in any UC-RAE of the GSP auction, the bids profile $\hat{b}$ satisfies the following conditions:*

- *if $i \notin C$:*

$$\underbrace{\frac{\hat{b}_i x^{i-1} - b_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i} > \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \tag{10}$$

- *if $i \in C$ and $i \neq \min(C)$:*

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i^f \ (\leq v_i)} = \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \tag{11}$$

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 application in Section 6.

5.2.2 Lifting the UC-Restriction: Revenue Losses and Inefficiency

As discussed in Section 5.1, even a small coalition of bidders may have a large impact on revenues in the VCG. Theorem 2 therefore already entails a fairly negative outlook on the GSP's revenues when an agency is active, even if it cannot be detected as collusive. The next example shows that, when the undetectability constraint is lifted, an agency may induce larger revenue losses as well as inefficient allocations in the GSP auction.

Example 5 Consider an environment with 8 bidders and 7 slots, with valuations $v = (12, 10.5, 10.4, 10.3, 10.2, 10.1, 10, 1)$ and CTRs $x = (50, 40, 30.1, 20, 10, 2, 1, 0)$. Let the coalition be $C = \{5, 6\}$. 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 $b = (b_1, 9.91, 9.76, 9.12, 9.5, 7.94, 5.5, 1)$. Note that $b_4 = 9.12 < 9.5 = b_5$, which induces an inefficient allocation. The inefficiency arises as follows: Suppose that the agency drastically lowers b_6 to benefit the other member. If b_6 is very low, it 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 its lower bid. Hence, in order to prevent these independents from doing so, 5's bid must also be reduced, so as to make the higher positions more attractive. But the reduction of 6's bid in this example is large enough that the undercut of 4 is sufficiently low 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. □

Hence, unlike the VCG mechanism, the unrestricted RAE of the GSP auction can be inefficient. In light of this result, it may appear that the unconstrained-RAE in the GSP allows an implausible degree of freedom to the agency, and that this alone is the cause of the low revenues of the GSP auction. To see whether this is the case, we consider next a set of exogenous restrictions that force the agency to induce efficient allocations. This is useful to isolate the price-reducing effect of bidding coordination separately from its potential allocative effect, and provides a less extreme and perhaps more plausible model of agency collusion. Theorem 3 shows that, even with this restriction, the GSP's revenues are no higher than in the unrestricted RAE of the VCG mechanism. Formally, let $\mathcal{R}^{EFF} = \left\{R_C^{EFF}\right\}_{C\in\mathcal{C}}$ be such that, for each non trivial coalition $C \in \mathcal{C}$,

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

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 fragility of the GSP's revenues is 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, and hence the agency's manipulation of its

Table 2: 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**	$\mathbf{2}^+$	**1.8**	$\mathbf{1.6}^+$
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, 4 and 6. 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 6.

members' bids only has a *direct effect* on revenues. In the GSP auction, in contrast, the agency has both a *direct* and an *indirect effect*. Under the UC-restrictions, the two effects combined induce just the same revenue-loss as in the VCG mechanism, but lifting that restriction tilts the balance, to the disadvantage of the GSP.

Since the UC-RAE induce efficient allocations (Theorem 2), 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 $\mathcal{R} = (R_C)_{C \in \mathcal{C}}$ are changed, they change for all coalitions: hence, even if R_C is weaker for any given 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.

Example 6 Consider the environment of Examples 3 and 4, with $C = \{1,3\}$. The efficiency-constrained RAE is $\hat{b} = \left(\hat{b}_1, 2.8, 1.6^+, 1.6, 1\right)$, 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. (Table 2 summarizes and compares the equilibria illustrated in our running examples.)

Summing up, 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 agency'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 6, in which the efficiency constraint is not binding, but it is possible in general (see Example 5).

As done in the earlier sections, we characterize next the testable implications of the Eff-RAE of the GSP auction:

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

- *if $i \notin C$:*

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{=v_i} > \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \tag{12}$$

- *if $i \in C$ and $i \neq \min(C)$:*

$$\underbrace{\frac{\hat{b}_i x^{i-1} - \hat{b}_{i+1} x^i}{x^{i-1} - x^i}}_{\textit{less than } v_i} < \frac{\hat{b}_{i+1} x^i - \hat{b}_{i+2} x^{i+1}}{x^i - x^{i+1}} \tag{13}$$

5.3 Agency Competition

Multiple agencies competing in the same auction appears rarely in the data (Decarolis et al. (2016)), but for the reasons explained in the introduction, it is nevertheless interesting to assess whether competition may soften the impact of agency bidding on online ad auctions. This is a reasonable conjecture, but the results we present in this section suggest a more nuanced view on the impact of agency competition on the VCG and GSP auctions. On the one hand, for certain coalition structures, our earlier results extend to the case with multiple agencies essentially unchanged: the revenue losses will be less pronounced when the same set of coordinating bidders is divided into two (or more) competing coalitions, but they would still be substantial, and preserve the relative performance of the VCG and GSP auctions. On the other hand, for other coalition structures, equilibria in pure strategies will not exist. Hence, bidding cycles are likely to emerge. As discussed in Section 2, a similar phenomenon was observed for the earlier mechanisms used in this market, and is considered to be the main reason for the transition from such earlier mechanisms to the GSP auction.[22] 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.

For simplicity, we consider the case with two agencies (the extension to more than two agencies is cumbersome but straightforward). We also assume that agencies break indifferences over bids in the same way that independents do. This implies that the highest bidder in any coalition bids as if he were an independent. The next result formalizes the discussion above.

[22] See 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 agency, 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 agencies. The result therefore shows that competition between agencies may mitigate, but not solve, the revenue losses due to coordinated bidding. 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. We illustrate both these points in the context of our workhorse example.

Table 3: 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**	$\mathbf{b}_1$
4	3.15	**2.75**	**3.05**	$\mathbf{b}_2$
3	2.3	1.5	2.1	b_3
2	1.6	$\mathbf{0^+}$	**1.2**	$\mathbf{b}_4$
1	1	**0**	**0**	$\mathbf{b}_5$
Revenues	96	60	88	–

Example 7 Consider the environment of the examples in Table 2. Table 3 reports EOS' equilibrium bids (second column) as well as the bids under different coalition structures. We first look at the case of a single coalition $C = \{1,2,4,5\}$. According to our earlier results, in the UC-RAE with this agency configuration 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 do it. If we split C as in the fourth column of the table, $C_1 = \{1,2\}$ and $C_2 = \{4,5\}$, we obtain two coalitions with no adjacent members, as in part 1 of Theorem 4. With this coalition structure, equilibrium revenues amount to 88, which is above the single coalition case (60), but still well below the competitive benchmark (96).[23] If we split C as in the last column

[23]Note 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, without the assumption

of Table 3, $C_1 = \{1, 4\}$ and $C_2 = \{2, 5\}$, pure equilibria would cease to exist. To see this, note that C_2 would ideally like to set $b_5 = 0$. If it does so, however, C_1 will find optimal to set $b_4 = 0^+$. This, however, is incompatible with an equilibrium because once $b_4 = 0^+$, C_2 would find it profitable to increase b_5 so as to obtain a higher position, with a negligible increase in its payments. On the other hand, if b_4 is set sufficiently high that C_2 does not find this deviation profitable, then C_2's optimal response is to set $b_5 = 0$. But then, a strictly positive b_4 cannot be optimal for C_1. Hence, a pure equilibrium does not exist. □

Part 2 of Theorem 4 shows that this phenomenon emerges whenever two coalitions have members (other than their top members) which occupy contiguous positions in the ordering of valuations. It is interesting to note that 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, under which the market started, to explain the bidding cycles observed in the data. As discussed earlier, such bidding cycles are considered to be the main cause for the shift from the GFP to the GSP auction. The fact that a similar phenomenon emerges here with multiple agencies may thus be seen as a troubling result for the existing mechanisms, in that it suggests that agency competition, instead of mitigating the impact of agency bidding, could exacerbate the system's instability.

6 Application: A Method for Detecting Collusion

In this section we show how our model can be used to detect collusion in the typical datasets that are available to search engines. We first present the method and then illustrate its application through simulated data.

A typical search engine's dataset (e.g., Google's or Microsoft-Yahoo!'s) includes information on all variables in our model, except advertisers' valuations. In particular, search engines record advertisers' identity, their agencies (if any), bids, positions and CTRs. But the typical dataset also records information about 'quality scores', which for simplicity we ignored in the previous sections. Quality scores are the advertisers' idiosyncratic score assigned by the search engine to account for various quality dimensions, including the CTRs. In the variant of the GSP auction run by Google or Microsoft-Yahoo! (but not, for instance, by Taobao), 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 'quality 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

that top coalition members behave as independents, a multiplicity of equilibria might arise. Different selections from the best-response correspondence may thus be used to model other forms of behavior, such as spiteful bidding (cf., Levin and Skrzypacz, 2016).

effect'. The price paid by bidder i in position $\rho(i)$ is $p_i = e^{\rho(i+1)}b^{\rho(i+1)}/e_i$.[24] Relabeling advertisers so that $e_i v_i > e_{i+1}v_{i+1}$, the competitive (EOS) equilibrium bids are such that, for all $i = 2, ..., S$,

$$e_i v_i = \frac{e_i b_i x^{i-1} - e_{i+1} b_{i+1} x^i}{x^{i-1} - x^i} > \frac{e_{i+1} b_{i+1} x^i - e_{i+2} b_{i+2} x^{i+1}}{x^i - x^{i+1}} = e_{i+1} v_{i+1}. \tag{14}$$

This is the analogue, with quality scores, of the characterization of EOS' equilibrium in terms of the observable variables we provided in Corollary 1. As shown below, similar modifications apply to various notions of RAE discussed in the earlier sections, and will provide the basis for our proposed criterion to detect collusion.

6.1 Detecting Collusion in the GSP: Strategy

We devise next a criterion to say whether a given set of data for the GSP auction is more likely to be generated by competitive (EOS) bidding or by one of the models of agency coordination (UC-RAE, Eff-RAE and RAE). As we showed above, the latter models differ from EOS in that the bids of all agency bidders, with the exception of the highest coalition member, are 'too low'. For 2-bidder coalitions, this property leads to a simple classification criterion (the extension to larger coalitions is straightforward). Let j denote the lowest value agency bidder, and define

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

The key idea of our criterion is to look at the implications that different models of agency bidding have for this quantity J. For instance, it is immediate from eq. (14) that if j's bid is compatible with EOS (competitive) bidding, then it must be the case that $J > 0$. In contrast, as shown by equations (11) and (13), under our models of collusive bidding j's bid will be lower than in the competitive case, so that the above inequality no longer holds: it will either be such that that $J = 0$, as in the UC-RAE case, or it such that $J < 0$, as in the Eff-RAE and unrestricted RAE. Note that, in two-bidder coalitions, this criterion actually captures *all* observable implications (i.e., ignoring valuations) that differentiate collusive from EOS bidding, and UC-RAE from Eff-RAE and Eff-RAE.

Thus, if we have as set of T auctions, $t = 1, 2, ..., T$ for the same keyword/coalition and for which we observe quality scores, bids, CTRs and positions for all bidders, and let J_t denote the value taken by quantity J in auction t, then we can study the distribution of J_t across these auctions to assess whether bidding in these auctions is competitive of collusive. For instance, if we find evidence that J_t is positive, then we can say that there is evidence in favor of competitive bidding. Otherwise, the evidence will be in favor of collusion. We next turn to simulated data to illustrate how to operationalize this idea.

[24] In 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.

6.2 Simulation

Consider once again the example in Table 2. We hold fixed the valuations, CTRs and coalition structure as in Table 2 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 2, the lowest value member of the coalition is the bidder with a value of 3, we calculate the value of J_t for this bidder for all simulated auctions under three different equilibrium scenarios. We report the resulting distributions of J_t in panel (a) of Figure 2: EOS (solid line), UC-RAE (dashed line) and Eff-RAE (dotted line).

Figure 2: Simulation

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

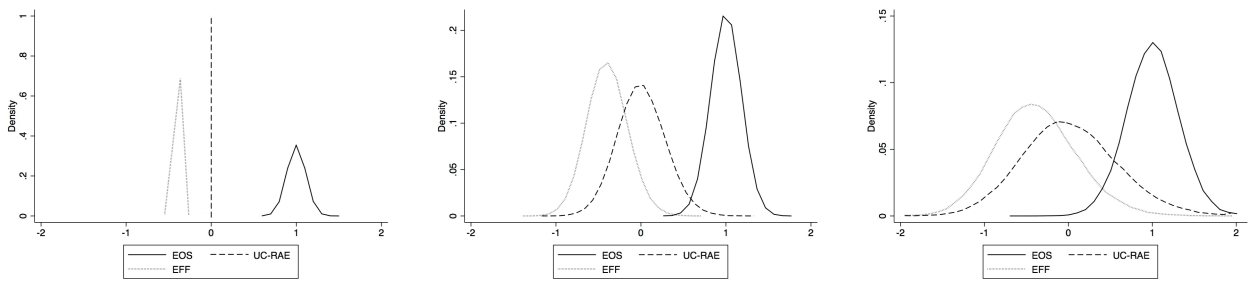

The distributions in panel (a) show that, as expected, J_t is never negative when we simulate EOS, it always equals zero when we simulate UC-RAE, and it is never positive when we simulate Eff-RAE.[25] Under the ideal conditions of the simulation, the observation of the distribution of J_t thus allows us to unambiguously separate the bidding models. Clearly, with real data this tool should be expected to face some limits. For instance, search engines' quality scores are updated in real time, and hence even if bidders can frequently adjust their bids, bids are not always optimized for the 'true' quality scores. That is, there may be 'belief errors' on quality scores, which (albeit small) may impact J_t.

To illustrate this point, in plot (b) and (c) of Figure 2 we repeat the previous simulation under two scenarios. In both cases, we consider a belief error that enters multiplicatively: for each bidder i and auction t, we let e_{it} denote the true quality score, but assume that bidders believe it to be $\tilde{e}_{it}$, where $\tilde{e}_{it} = d_{it} \cdot e_{it}$, where d is drown from a normal distribution centered around 1. Panel (b) considers the case of a small error, with $d_{it} \sim \mathcal{N}(1, 0.05^2)$; 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

[25] Detecting 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 where changing with the quality scores in an ad hoc way, hence the detectability of UC-RAE across auctions.

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 by looking at some summary measure like mean, median or mode. As shown by panel (c), however, when the amount of noise is large, none of these methods will yield an entirely unambiguous 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.

For those cases where the above method reveals the likely presence of collusion, a simple approach can be followed to invert bids and recover estimates of the potential revenue losses. 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 independents' values, and hence the bounds for the values of the coalition members.[26] Although no 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.

7 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

[26] To 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.

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 5, 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 marketing agencies. 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.[27] 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

[27] Quality 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. Che, Y.K., Condorelli, D. and J. Kim (2016). "Weak Cartels and Optimal Auctions," *mimeo.*

11. Che, Y.K., and J. Kim (2006). ""Robustly Collusion-Proof Implementation," *Econometrica*, 74, 1063–1107.

12. Che, Y.K., and J. Kim (2009). "Optimal Collusion-Proof Auctions," *Journal of Economic Theory*, 144, 565–603.

13. Choi, H. and H. Varian, (2012). "Predicting the Present with Google Trends," *Economic Record*, 88 (s1), 2-9.

14. 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.*

15. Decarolis, F., M. Goldmanis, A. Penta, (2016). "Common Agency and Coordinated Bids in Sponsored Search Auctions, mimeo.

16. Econconsultancy (EC), (2011). "State of Search Marketing Report 2011," accessed on January 2012 at: http://econsultancy.com/us/reports/sempo-state-of-search

17. 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.

18. Edelman, B., M. Ostrovsky (2007). "Strategic Bidder Behavior in Sponsored Search Auctions," *Journal of Decision Support Systems*, v.43(1), pp.192-198.

19. Edelman, B., M. Schwarz (2010). "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100(2), 597–602.

20. Einav, L., C. Farronato, J. Levin, N. Sundaresan (2014). "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?," *mimeo.*

21. Gomes, R., K. Sweeney (2014). "Bayes-Nash Equilibria of the Generalized Second-Price Auction," *Games and Economic Behavior*, 86, 421-437

22. Graham, D. and R. Marshall (1987). "Collusive Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, 95, 1217-1239.

23. Harrington, J. E., A. Scrkypacz, (2007)."Collusion under Monitoring of Sales," *RAND Journal of Economics*, 38(2), pp. 314-331.

24. Harrington, J. E., A. Scrkypacz, (2011). “Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices,” *American Economic Review*, 101(6), 2425-49.

25. Hendricks, K., R. Porter and G. Tan (2008). “Bidding Rings and the Winner’s Curse,” *RAND Journal of Economics*, 39, 1018-1041.

26. Levin, J., A. Scrkypacz. (2016). “Properties of the Combinatorial Clock Auction” (2016) with Jonathan Levin,” *American Economic Review* 106(9), pp. 2528-2551.

27. Mailath, G., P. Zemski. (1991). “Collusion in Second Price Auctions with Heterogenous Bidders,” *Games and Economic Behavior*, 3, 467-420.

28. McAffee, P. and J. MacMillan. (1992). “Bidding Rings,” *American Economic Review,* 82, 579-599.

29. Milgrom, P., J. Mollner. (2014). “Equilibrium Selection in Multi-Player Games with Auction Applications,” mimeo.

30. Ottaviani, M. (2003). “Overture and Google: Internet Pay-Per-Click (PPC) Advertising Auctions,” London Business School - Case Study-03-022

31. 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

32. Ray, D., R. Vohra. (1997). “Equilibrium Binding Agreements,” *Journal of Economic Theory*, 73, 30–78.

33. Ray, D., R. Vohra. (2014). “The Farsighted Stable Set,” *Econometrica*, forthcoming.

34. Varian, H., (2007). “Position auctions,” *International Journal of Industrial Organization*, 25(6), 1163–1178.

35. Wired, (2015) https://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/.

36. Wu, T. (2016). *The attention merchants: the epic scramble to get inside our heads*,” Knopf, New York.