Marketing Agencies and Collusive Bidding* in Online Ad Auctions†

Francesco Decarolis, Maris Goldmanis, Antonio Penta‡

March 27, 2018

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.

*The use of the word *collusion* in this essay is unrelated to any assessment of the legal implications of agencies or advertisers behavior under the competition laws of the US or other countries.

†We are grateful for the comments received from Susan Athey, Jean-Pierre Benoit, Yeon-Koo Che, Kerry Gabrielson, Ken Hendricks, Jon Levin, Massimo Motta, Marco Ottaviani, Rob Porter, Marc Rysman, Andy Skrzypacz and Steve Tadelis and from the participants at the seminars at Berkeley-Haas Business School, Boston University, Columbia University, CREST-Paris, European Commission DG Competition, Facebook Research Lab, Harvard-MIT joint IO seminar, HEC-Lausanne, London Business School, Microsoft Research Lab, Princeton University, Stanford University, University of British Columbia, University of California Davis, University of Chicago, University of Mannheim, University of Toronto, University of Wisconsin Madison and at the conferences of the EIEF-University of Bologna-Bocconi joint IO workshop, Society for Economic Dynamics and Toulouse Network for Information Technology Meeting where earlier versions of this paper were presented. Antonio Penta acknowledges the financial support of the ERC Starting Grant #759424.

‡Decarolis, Bocconi University, *francesco.decarolis@unibocconi.it*. Goldmanis, Department of Economics, Royal Holloway, University of London, *Maris.Goldmanis@rhul.ac.uk*. Penta, UW-Madison, UPF and Barcelona GSE, *antonio.penta@upf.edu*.

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, Nosko and Tadelis (2015)). Almost all online ads are sold through auctions, in which bidders compete for the adjudication of one of a given number of 'slots' available in various online venues, such as search engine result pages, social networks feeds, and so on. With the significant exception of Facebook, which recently adopted the 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.[1] But these models do not account for a recent trend in this market, which is bound 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).[2] 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, thus representing a blatant instance of the possible anti-competitive effects posed by algorithm-based pricing which have recently become a central antitrust concern (OECD, 2017).

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 the agency's equilibrium bids are akin to implementing a certain form of collusion (even if none of its clients explicitly

[1] Gomes and Sweeney (2014) provide a more critical assessment of the GSP auction in an environment with incomplete information. In constrast, EOS, Varian (2007) and the present model maintain a complete information setting. This modeling choice as well as other variations are discussed in Section 3.

[2] 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.

attempt it), and that in this situation 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 (e.g., Ausubel and Milgrom, 2006), but it is especially noteworthy if one considers the sheer size of transactions currently occurring under the GSP. It also suggests a rationale for why Facebook's recent adoption of the VCG mechanism was so successful, despite the early surprise it provoked (e.g., *Wired* (2015)), and for why the last few years have recorded a steady decline in ad prices.[3] 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.[4] To the best of our knowledge, this is the first study to point at the central role of marketing agencies in these auctions, and to their potential policy implications (see below).

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)).[5] Second, strategic behavior in the GSP auction is complex and brings forth a plethora of equilibria. 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.[6] 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.

[3]Google, for instance, reports passing from a positive growth rate in its average cost-per-click of about 4 percent per year in the four years before 2012, to a negative growth rate in each year since then, with an average yearly decline of 9 percent. Source: 10-k filings of Alphabet inc.

[4]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).*

[5]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), or Chassang and Ortner (2017) in a different setting). The main focus of that literature is on the coalition members' incentives to share their private information so as to implement collusion, 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).

[6]On a similar note, 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.

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 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. (2018)), 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.

Finally, it is important to stress that the type of collusion that we discuss does not require explicit communication and can be implemented tacitly through the mere functioning of bid optimization algorithms. As discussed in the next section, the success of agencies is closely connected to the shift to algorithmic bidding. OECD (2017) reviews the key elements of the ongoing debate on algorithms and collusion identifying two main features. First, thanks to high price transparency and high-frequency trading, algorithms increase the ability to react fast and aggressively, thus making collusive strategies more stable in virtually any market structure. Second, by using automated mechanisms to implement common policies or optimize joint profits with deep learning techniques, algorithms can lead to the same outcomes of traditional *hard core* cartels through tacit collusion. How exactly are online ad auctions affected by this latter force is what the following analysis uncovers. We return to its nuanced policy implications in the conclusions.

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)). 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 has since boosted Google's business: 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.

Facebook, in particular, introduced the VCG for its own display ad auctions.[7] 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 Facebook, 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 to Facebook's innovation 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

[7] After an initial period in which also Facebook adopted the GSP auction, they started experimenting with the VCG around 2007, and then progressively expanded its application. By 2015, most of Facebook's auctions were operated through a VCG auction format. In 2012, also Google introduced a version of the VCG auction for its contextual ads, but not for the search ads (cf. Varian and Harris (2014)).

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.

with Twitter, it has become one of the largest players in display ad auctions. Together, 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 2), a large share of advertisers in the US delegated their bidding activities to specialized 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.[8] 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.[9] Table 1 reports 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 orig-

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

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

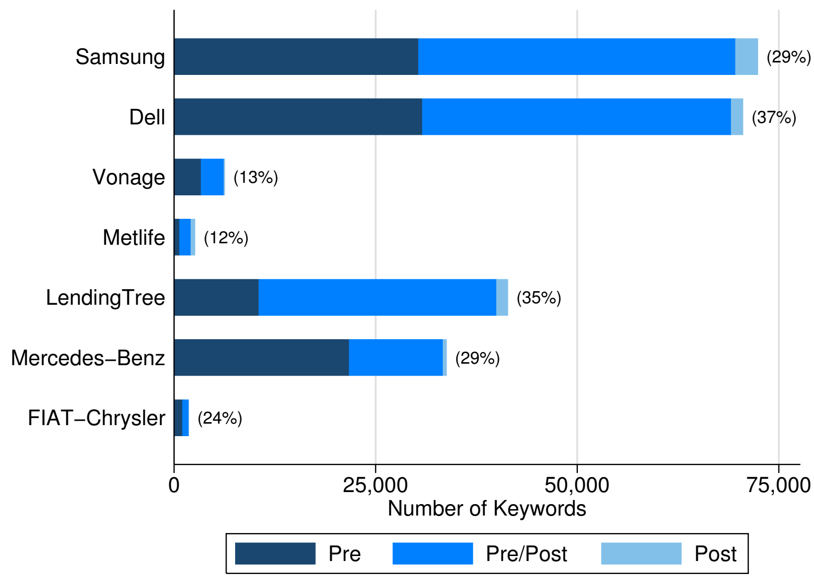

Figure 1: For each of Merkle's advertisers in footnote 10, the figure represents the number of keywords on which it bid alongside at least one member of the Aegis-Dentsu network (and as a share of the total number of keywords on which it bid, in parenthesis) between June 2015 and January 2017 (Merkle's 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.

inally 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.[10] This acquisition therefore further increased the po-

[10]For instance, in the electronics sector, *Dell* and *Samsung* were in Merkle's portfolio, placing bids on keywords also targeted by Aegis-Dentsu's clients *Apple*, *HP*, *IBM/Lenovo* and *Intel*. Other examples include: in the financial sector, Merkle's *Lending Tree* and *Metlife* were bidding in auctions alongside Aegis-Dentsu's *Capitalone, Discover, Fidelity, Equifax, JP Morgan-Chase*; for car manufacturers, Merkle's *FIAT-Chrysler* and *Mercedes-Benz USA* bid alongside Aegis-Dentsu's *Toyota, Volkswagen, Subaru*; in

tential for coordinated bidding. Figure 1 reports, for each of Merkle's advertisers listed in footnote 10, 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 relevant.

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 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.

While uncommon in broader auction theory, this complete information environment is the main benchmark for the literature on the GSP auction.[11] It thus represents the natural starting point to focus on novel substantive problems, and to be able at the same time to make clear comparisons with the existing benchmarks. But there are deeper economic reasons which make the complete information assumption desirable in this context. First and foremost, bidders can easily access detailed information on their rivals' bids using both search engines' tools (like Google's *Keyword Planner*) and direct experimentation.[12] Indeed, this is why from the very beginning (Ottaviani, 2003) the literature has consistently preferred the complete information structure.

A case in point is provided by the bidding cycles this market witnessed under the Generalized *First* Price auction (cf. Section 2). This phenomenon could be hardly explained by the standard incomplete information model, but has a very simple explanation if looked at through the lens of complete information. The key fact is that these auctions

phone services, Merkle's *Vonage* bid alongside Aegis-Dentsu's *T-Mobile*. (Source: Redbook.)

[11] A notable exception is Gomes and Sweeney (2014), which provides a thorough analysis of competitive bidding in the GSP auction with independent private values, with a much more pessimistic outlook on both the allocative and revenue properties of the GSP auction. Borgers et al. (2013) maintain the complete information assumption, but consider a more general model of CTRs and valuations.

[12] Learning rivals' bids by trying out different bids is typically fast - as Google, for instance, allows to change bid every two hours - and not particularly expensive for most keywords.

are repeated over time, and bidders receive information at high frequency on earlier bids. If valuations are fairly stable, the problem at 'steady state' is better approximated by a complete information model.[13] Finally, there is now a diverse set of empirical analysis suggesting that uncertainty over valuations plays a small role in these environments, and that the complete information assumption generally provides an excellent approximation, from a quantitative viewpoint, even in settings in which advertisers may be uncertain over various features of the environment like the quality scores or the set of rivals (e.g., Athey and Nekipelov (2012), Che et al. (2017) and Varian (2007)).

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).[14] 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.[15] 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-click 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

[13] Intuitively, if valuations are fairly stable over time, relative to the frequency with which bidders may adjust their bids, then bidders would learn others' bids. They would thus have the opportunity to lower their own bid (and hence their payment) with essentially no risk of losing their position. As bids are lowered over time, at some point some bidder would find it profitable to increase his bid and jump to a higher position. But this would induce others to increase their bids, so as to restore the original ranking. But then again bids could not remain stable: if bidders know others' bids, they have incentives to lower their own, and the cycles starts again. This logic is neatly distilled by the lack of pure strategy equilibria in the GFP auction with complete information (Edelman and Schwarz, 2007). Note, however, that none of the above would happen if players faced sufficient uncertainty about others' bids (for instance, if valuations were sufficient variable): in that case, bidders could not lower their bids without risking to lose their position. This is in fact the reason why, pure equilibria exist in the GFP auction with incomplete information, and why the observed phenomenon of bidding cycles is better understood with a complete information model.

[14] 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 20).

[15] 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)).

mechanism, given a bids profile $b \in \mathbb{R}^n_+$, as $u_i^{\mathcal{V}}(b) = v_i x^{\rho(i)} - \sum_{t=\rho(i)+1}^{S+1} b^t(x^{t-1} - x^t)$.

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.

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}.[16] 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-revenue locally envy-free equilibrium*, which was crucial to cut through the complexity of the GSP auction.[17] EOS' refinement has proven very successful both as a theoretical construction and for its performance as basis of empirical analysis of GSP auctions (see, e.g., Varian (2007), EOS (2007) and Edelman and Schwarz (2010), Milgrom and Mollner (2018), etc.). It is also particularly important because it conforms with the tutorials that the search engines provide on how to bid in these auctions.[18]

The problem with EOS refinement for our purposes is that it is not defined in our context, in which both agency and independent bidders are present. We thus consider instead a refinement of individuals' best response correspondence, which distills the individual-level underpinnings behind the basic idea EOS refinement. Intuitively, what we propose is a selection from individual best-responses in which advertisers place bids which make them indifferent between getting their position at a price equal to the next highest bid,

[16] Allowing ties in individuals' bids or non-generic indifferences complicates the notation, without affecting the results and the main insights. See Appendix A.1 for details on this.

[17] Formally, a Nash equilibrium $(b_i)_{i \in I}$ is locally envy-free if it satisfies $x^{\rho(i)}(v_i - b^{\rho(i)+1}) \geq x^{\rho(i)-1}(v_i - b^{\rho(i)})$ for every i. EOS refinement is the lowest-revenue Nash equilibrium which satisfies this condition.

[18] 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.

and climbing up one position paying their own bid. In this sense, it formalizes the idea of envy-freeness implicit in EOS' refinement. Formally: for any $b_{-i} \in \mathbb{R}_+^{n-1}$, let

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

In words, of the many $b_i \in BR_i\left(b_{-i}\right)$ that would grant player i his favorite position $\pi_i\left(b_{-i}\right)$, 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 $E^*\left(v\right)$.

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

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

Hence, the fixed points of the BR^* correspondence coincide with EOS' lowest revenue envy-free equilibrium (LREF)*, 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^* provides a model of individual behavior which is consistent with EOS' equilibrium, and which is well-defined in our setting even if EOS' equilibrium is not. In particular, in Section 5 we will assume that independent bidders bid according to BR_i^* in the GSP, 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.

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 shown that EOS' equilibrium satisfies both refinements.[19] 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 $C \setminus \{i\}$, 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

[19] 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. Approaches similar to ours have been previously used, for instance by Aghion, Antras and Helpman (2007) who incorporate insights from Ray and Vohra (1997) to study regionalism versus multilateralism by analyzing whether multilateral or sequential bargaining is more likely to lead to global free trade.

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 14), in that case bidder i obtains the position above $i+1$. To emphasize this, we will write $b_i = b_{i+1}^+$.[20]

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 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,

[20] Without the tie-breaking rule embedded in ρ (footnote 14), 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 $b_{i+1}^+ := \lim_{\varepsilon^+ \to 0} (b_{i+1} + \varepsilon)$).

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.[21] $\square$

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 is $u_2^{Comp} = 141 < 145$. 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 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 5); (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 (which will be discussed in 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.[22] Independent of these modeling choices, however, the cost of joining the

[21] 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.

[22] 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

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, and the extension to the multiple agency case.

4.2 The Recursively Stable Agency Equilibrium: General Definition

Let $G = (A_i, u_i)_{i=1,...,n}$ 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 $i \in I \backslash C$ 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,[23] 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 $R(C) \subseteq A$, allows us to accommodate the possibility that the agency may exogenously discard certain bids. For instance, we will consider the case of an agency whose primary concern is not being identified as inducing collusion (Section 5.2.1) or to induce efficient outcomes (Section 5.2.2). In those cases, $R(C)$ would be comprised respectively of only those profiles that are 'undetectable' to an external observer as collusive, or efficient. We denote the collection of exogenous restrictions for all possible coalitions as $\mathcal{R} = \{R(C)\}_{C \in \mathcal{C}}$, and for any C we also let $R_C \subseteq A_C$ denote the restriction it entails on the coalition bids. That is, $R_C := \{b_C \in A_C : \exists b_{-C} \in A_{-C} \text{ s.t. } (b_C, b_{-C}) \in R(C)\}$.

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 baseline game G (e.g., truthful bidding in the VCG, or (1) in the GSP). Define the *independents' equilibrium*

unilateral, the decision to join it is not, raising further modeling questions.

[23] 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.

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 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:

$$B^* = \left\{ \beta \in A^{S_C}_{-C} : \beta(b_C) \in BR^*_{-C}(b_C) \text{ for all } b_C \in S_C \right\}. \tag{6}$$

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 $\bar{u}^{C'}_i = \min_{b \in E^{\mathcal{R}}(C')} u_i(b)$. 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:*[24]

[24]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.

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

2. The conjectures of the agency are correct and consistent with the exogenous restrictions: $\beta^* \left(b_C^*\right) = b_{-C}^*$ and $\left(b_C, \beta^* \left(b_C\right)\right) \in R\left(C\right)$ for all $b_C \in R_C$.

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

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

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

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

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

The fact that both (S.1) and (S.2) enter as constraints in the agency's optimization problem reflects that, in our model which combines cooperative and non-cooperative ideas, stability considerations have in a certain sense priority over equilibrium logic, which operates within the limits set by the former. In the next section we apply this definition to study agency bidding in the GSP and VCG mechanism. As already noted, for those applications constraint (S.1) will not matter, because (5) will be either vacuous (Theorem 1) or redundant (Theorems 2 and 3).

For some general considerations on the solution concept, first note that, 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, 2015, 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. Hence, a direct application of their concept to this setting 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.

In light of the well-known difficulties of modeling partial cartels in auction which we discussed in the introduction (cf. Hendricks et al. (2008)), the results in the next sections perhaps suggest that concepts built on RV's approach such as our RAE may prove useful to overcome some of those difficulties. In that sense, the general concept of RAE may provide a valuable methodological contribution from a broader theoretical perspective.

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 (i.e., 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.[25] If

[25] 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.

the outside options for the coalition members were defined with respect to the competitive case, for instance, bidder 2 in Ex.2 would remain in the coalition even when forced to take the lower position. But, as already discussed, 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? The recursive stability constraint therefore captures an important aspect of the environment and poses economically meaningful restrictions on the agency's freedom to manipulate the bids of its clients. This feature of our model is reflected in the proof of Theorem 1, which is based on a recursive argument showing 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 to use 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. (2018)). 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 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 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 $C \in \mathcal{C}$,

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

In words, $R^{UC}(C)$ is comprised of all bid profiles that could be observed in 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, $R^{UC}(C)$ 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.[26] 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 (and hence total revenues) as the unrestricted RAE of the VCG mechanism.

To understand the implications of this equilibrium, note 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 some qualitative insights for this complex problem.

Remark 1 *Hold the agency configuration, C, constant. Then, in both the unconstrained RAE of the VCG and in the UC-RAE of the GSP auction, the revenue losses due to agency bidding are larger if: (i) the differences $(x_{i-1} - x_i)$ associated to the agency's clients $i \in C$ are larger; or if (ii) the difference in valuations between the agency's clients and the independents immediately below them in the ranking of valuations are larger.*

To understand this remark, recall that the price-per-click for position s in the VCG, given a profile b, is equal to $\sum_{t=s+1}^{S+1} b^t (x^{t-1} - x^t)$. By Theorem 1, in the RAE of the

[26] 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.

VCG the agency lowers the bids of its members as much as possible, while preserving the efficient ranking of bids. Hence, holding C and $(v_i)_{i\in I}$ constant, it is clear that the revenue losses due agency bidding are larger if the terms $(x^{t-1} - x^t)$ associated to agency members are larger, which is part (i) of the Remark. To understand part (ii), let i be an agency member such that $i+1$ is an independent. Since independents bid truthfully in the VCG, we have $b_{i+1} = v_{i+1}$, and hence the efficient ranking can be maintained only if $b_i \geq v_{i+1}$. Hence, the lower v_{i+1}, the stronger the impact of agency bidding.

The next comparative statics refer to the agency composition. Besides the obvious statement that an agency's impact is stronger if it includes more bidders, the impact of different coalitions in general depends on the exact CTRs and valuations. To isolate the position effects from the comparative statics in Remark 1, which were driven by the differences $(x_s - x_{s+1})$ and $(v_s - v_{s+1})$, we assume that they are constant in s.

Remark 2 *Assume that $\Delta_s(x) := (x_s - x_{s+1})$ and $\Delta_s(v) := (v_s - v_{s+1})$ are constant in s. Then, in both the RAE of the VCG and in the UC-RAE of the GSP, the revenue losses due to agency bidding are larger if the agency includes members that occupy* adjacent *or* lower *positions in the ranking of valuations.*

To understand this result, note that if an agency has no two 'adjacent' members, then $i+1$ is an independent for all $i \in C$, and hence for the above explanation the lower bound to i's bid equals v_{i+1}. But if instead $i+1$ also belongs to the agency, then the lower bound drops to the valuation of the next lower independent. The rest of the Remark follows directly from the fact that a given reduction of a bid in the VCG has a larger impact if it's *lower* in the ranking, because it affects the payments for all positions above. The latter point is particularly interesting, since one might have expected that the agency would have a larger impact if she controlled the high-valuation bidders. We find that, in fact, the opposite is true when one controls for the increments $\Delta_s(x)$ and $\Delta_s(v)$.

The equilibrium characterization in Theorem 2 involves bidders' valuations, which are typically not observable. The conditions in (9), however, can be rearranged to obtain a characterization that only depends on the CTRs and the individual bids:

Corollary 2 *For any C, the UC-RAE bids $\hat{b}$ in the GSP are such that:*

- *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 for the competitive benchmark (eq. 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, if efficiency were to be preserved, 5's bid would also have to be reduced, to make the higher positions more attractive. But the reduction of 6's bid in this example is large enough that 4's undercut is sufficiently low that the coalition prefers to give up position 5. 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 exogenous restrictions that force the agency to induce efficient allocations. 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 $\mathcal{R}^{EFF} = \left\{R^{EFF}\left(C\right)\right\}_{C \in \mathcal{C}}$ be such that, for each non trivial coalition $C \in \mathcal{C}$,

$$R^{EFF}(C) := \{b \in A : \rho(i; b) = i \ \ \forall i \in I\}.$$

Definition 3 *An efficiency-constrained RAE of the GSP auction is a RAE of the GSP auction where the exogenous restrictions are given by* $\mathcal{R} = \mathcal{R}^{EFF}$.

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 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:*

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.

- *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. (2018)), 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.[27] Hence, while competition between agencies may produce the expected result of mitigating the revenue losses due to bidding coordination, it may also impair the

[27] 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.

Table 3: Competition between Agencies

Valuations	GSP (EOS)	Single Coalition: $C = \{1, 2, 4, 5\}$	Two Coalitions: $C_1 = \{1, 2\}$, $C_2 = \{4, 5\}$	Two Coalitions: $C_1 = \{1, 4\}$, $C_2 = \{2, 5\}$
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	–

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. With the formal definitions given in Appendix A.3, the following result holds.

Theorem 4 *1. If no members of different coalitions occupy adjacent positions in the ordering of valuations, 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 lower than under full competition. Moreover, both the allocation and the associated revenues are identical to those resulting in the unconstrained RAE of the VCG mechanism with the same agency configuration.*

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.

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).[28] If we split C as in the last column 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 $b_5 = 0$, and given this C_1 would ideally like 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 so 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 non-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

[28]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 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).

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 effect'. The price paid by bidder i in position $\rho(i)$ is $p_i = e^{\rho(i+1)} b^{\rho(i+1)} / e_i$.[29] 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 two-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 $J = 0$, as in the UC-RAE case, or 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 a 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 we let J_t denote the value taken by quantity J in auction t, then we can study the distribution of J_t

[29] 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.

across these auctions to assess whether bidding 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.

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

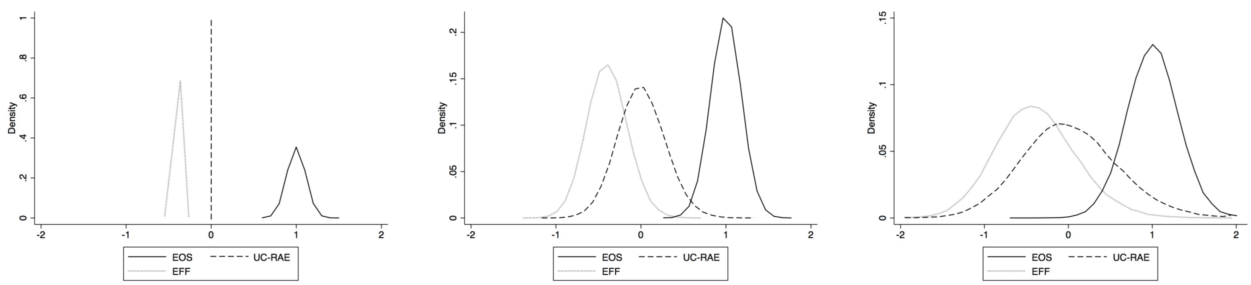

The distributions in panel (a) show that, as expected, J_t is never negative when we simulate EOS, it always equals zero when we simulate UC-RAE, and it is never positive when we simulate Eff-RAE.[30] Under the ideal conditions of the simulation, the observation of the distribution of J_t thus allows 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

[30] Detecting bids as coming from UC-RAE, in which coordinated bids were defined as 'undetectable', may strike as oxymoronic. The reason is that, by definition, 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.

bidders believe it to be $\tilde{e}_{it}$, where $\tilde{e}_{it} = d_{it} \cdot e_{it}$, where d is drawn 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 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. This is indeed what we also find in Decarolis, Goldmanis and Penta (2018) where we apply the methodology described above to data from a search engine and use it to estimate bounds on the revenue losses induced by collusion.

7 Conclusions

This is the first study to focus on the impact of marketing agencies on online ad auctions, and in particular on their role in coordinating the bids of their clients. Our results uncover a striking fragility of the GSP auction to bid coordination.[31] Aside from its theoretical interest, this is a first order finding since most of the online marketing is still passing through GSP auctions. Our findings may also provide a rationale for why Facebook has recently adopted the VCG and Google is said to be considering the transition. Shifts between one mechanism and the other are important both for the large stakes involved and because the proper functioning of this market is essential for both advertisers to reach consumers and for consumers to learn about products.

From a methodological perspective, we note that the notion of RAE has been key to obtain clear results in the complex GSP auction, and more broadly to accommodate the coexistence of competitive and coordinated bidding. This suggests that our 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.

Our results are also interesting from a market design perspective. While beyond the

[31] The empirical analysis in Decarolis, Goldmanis and Penta (2018) shows that even the small two-bidder coalitions frequently observed in the data can have large effects on revenues.

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 bid coordination. This kind of intervention would thus reinforce the resilience of the GSP auction, without necessarily entailing major changes in the mechanism.

From a broader perspective, our findings complement other recent work on the evolution of bidding behavior in online auctions. For instance, Blake, Nosko and Tadelis (2015) use large scale experiments to explore how eBay could benefit from a more nuanced bidding behavior that distinguishes between brand and non-brand keyword ads. Einav, Farronato and Sundaresan (2018) document a decline in the importance of consumers’ bidding in the eBay auctions, with a progressive shift towards purchasing at posted prices. Altogether, it emerges the picture that bidding behavior in online marketing platforms is undergoing important transformations that still need careful analysis.

Implementing the type of collusion that we discussed simply requires an agency to use a bidding algorithm optimizing joint profits. The use of deep learning techniques can make such an algorithm particularly effective at bolstering the revenues of the agency’s clients. The anticompetitive potential of algorithms has not been unnoticed (OECD, 2017), but there is no consensus on what the regulatory response should be. The tacit nature of algorithmic collusion is what drives the opposite views. Scholars emphasizing the of US doctrine on communication as an important condition to define the concept of agreement advocate a more cautions approach relative to those arguing in favor of a broader and less formalistic approach to firms’ coordination that emphasizes the “economic consequences regardless of the particular manner of interactions that generate this outcome” (Kaplow, 2011). Our analysis underscores that the tacit collusion implementable through agency bidding can achieve the same outcomes of an hard core cartel. Furthermore, the collusive outcomes involve not only a different distribution of revenues between advertisers and search engines, but also efficiency losses in terms of ad rankings that might negatively impact the consumers’ search experience. This indicates the need for close scrutiny of agency bidding behavior by antitrust agencies.

Furthermore, a second feature from an antitrust perspective is that the agency behavior in our model is analogous to that of buying consortia, which have been sanctioned in the past (see footnote 4). Nevertheless, the specificities of online ad market suggest a more nuanced view of the harm to the consumers. First, although our model focuses on agencies’ role to coordinate their clients’ bids, agencies in this market have other roles which are expected to improve the efficiency of the system (e.g., in improving sellers’ ability to reach new consumers, improving advertisers’ campaigns, bringing new advertisers to the market, etc.) Second, it is likely that the degree of competition between different search engines is substantially less than that between most of advertisers. Since the lower auction prices due to agency bidding imply a reduction in the marginal cost advertisers pay to reach consumers, advertiser competition implies that some savings are passed on to consumers.

Therefore, harm to consumers would result only if the agency engages in coordinating not only auction bids, but also the prices charged to consumers. Third, bid coordination can negatively affect the quality of the service received by consumers by further exacerbating the advantage of dominant search engines relative to fringe ones. In Europe, for instance, where 90% of the searches pass via Google, agencies might be rather careful not to harm Google given the risk of being excluded from its results page. Smaller search engines cannot exert such a threat because agencies are essential to attract new customers. The shift of revenues from small search engines to marketing agencies could thus deprive the former of the essential resources needed for technology investments. Thus, to the extent that competing search engines exert pressure for quality improvements, bid coordination poses a threat to consumer welfare. All these considerations represent potentially fruitful directions for future research on this important market.

A Appendix

A.1 Technical Details

As discussed in Section 3 any generic profile $b_{-i} = (b_j)_{j \neq i}$ in the GSP auction partitions the space of i's bids, $\mathbb{R}_+$, into $S+1$ intervals: $[0, b_{-i}^S)$, $[b_{-i}^S, b_{-i}^{S-1})$, $\ldots$, $[b_{-i}^1, \infty)$. Letting $b_{-i}^0 \equiv \infty$ and $b_{-i}^{S+1} \equiv 0$, if bidder i bids $b_i \in (b_{-i}^t, b_{-i}^{t-1})$, then he obtains slot $t = 1, \ldots, S+1$ at per-click-price b^t. If b_i is placed at one extreme of such intervals, the allocation is determined by the tie-breaking rule embedded in the function ρ. The function π_i introduced in section 3 can be seen as a corresopndence $\pi_i : \mathbb{R}_+^{n-1} \rightrightarrows \{1, \ldots, S+1\}$ such that for each $b_{-i} \in \mathbb{R}_+^{n-1}$, $\pi_i(b_{-i}) = \arg\max_{t=1,\ldots,S+1} \left(v_i - b_{-i}^t\right) x^t$.[32] To allow for the possibility of ties in the bids profiles, it is necessary to generalize some of these concepts. In particular, if some of i's opponents place equal bids (i.e., $b_{-i} = (b_j)_{j \neq i}$ is such that $b_j = b_k$ for some $j \neq k$), then, depending on the tie-breaking rule embedded in ρ, some of the $S+1$ positions may be precluded to player i (e.g., if $i = 1$, and $b_2 = b_3$, if the tie-breaking rule is specified as in footnote 14, position $s = 2$ is precluded to player i). In that case, the argmax in the definition of π_i should be taken over the set of positions that are actually accessible to i. Formally: for any $b_{-i} \in \mathbb{R}_+^{n-1}$, let

$$\mathcal{S}(b_{-i}) = \{s = 1, ..., S+1 : \exists b_i \text{ s.t. } \rho(i; b_i, b_{-i}) = s\}.$$

Then, we redefine the function $\pi_i : \mathbb{R}_+^{n-1} \to \{1, ..., S+1\}$ as follows: for every $b_{-i} \in \mathbb{R}_+^{n-1}$

$$\pi_i(b_{-i}) \in \arg\max_{s \in \mathcal{S}(b_{-i})} \left(v_i - b_i^t\right) x^t.$$

[32]This correspondence is always non-empty valued, and multi-valued only if i is indifferent between two positions. We can ignore this case here (for instance, assuming that such ties are always broken in favor of the lower position) and treat $\pi_i : \mathbb{R}_+^{n-1} \to \Pi$ as a function (if not, π_i should be thought of as a selection from the correspondence above).

Since $\mathcal{S}(b_{-i})$ is always non-empty and finite, the best responses $BR_i : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$ defined in Section 3 is well-defined, and so is $BR_i^* : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$ in (1). With these changes to the definition of π_i, the rest of the analysis also extends to the case of ties in bids.

A.2 Proofs of the Main Results

All the results are proven for the case in which $n = S + 1$. The extension to the general case is straightforward but requires more cumbersome notation.

A.2.1 Proof of Lemma 1

Let $\hat{b} \in E^*(v)$. By definition, for any i , $\rho(i) = s$ implies $\pi_i\left(\hat{b}_{-i}\right) = s$ if $s \leq S$ and $\pi_i\left(\hat{b}_{-i}\right) = S + 1$ if $s > S$. Hence, $\hat{b}_{-i}^{\pi_i(b_{-i})} = \hat{b}^{s+1}$ whenever $s \leq S$. Now, for any i such that $\rho(i) \leq S$ and j s.t. $\rho(j) = \rho(i) + 1$, the following must hold:

$$\text{by the optimality of } \hat{b}_i: \quad \left(v_i - \hat{b}^{\rho(i)+1}\right) x^{\rho(i)} \geq \left(v_i - \hat{b}^{\rho(i)+2}\right) x^{\rho(i)+1}; \tag{15}$$

$$\text{by the condition in (1) for } j: \quad \left(v_j - \hat{b}^{\rho(i)+2}\right) x^{\rho(i)+1} = \left(v_j - \hat{b}^{\rho(i)+1}\right) x^{\rho(i)}. \tag{16}$$

Rearranging, we obtain

$$v_i \cdot \left(x^{\rho(i)} - x^{\rho(i)+1}\right) \geq \hat{b}^{\rho(i)+1} x^{\rho(i)} - \hat{b}^{\rho(i)+2} x^{\rho(i)+1} = v_j \cdot \left(x^{\rho(i)} - x^{\rho(i)+1}\right),$$

which implies that $v_i > v_j$ (since, by assumption, $x^s > x^{s+1}$ for all $s \leq S$ and $v_i \neq v_j$ for all $i \neq j$). Hence, in equilibrium, the top S bidders are ranked efficiently among themselves. For the others, for any i such that $\rho(i) > S$, eq. (1) requires that $0 = \left(v_i - \hat{b}_i\right) x^S$, hence $v_i = \hat{b}_i$ whenever $\rho(i) > S$. It follows that $\hat{b}^i = \hat{b}_i$ for all i (agents bids are efficiently ranked) and $\hat{b}_i = v_i$ for all $i \geq S + 1$. Equation (2) follows immediately, applying eq. (1) for all $i = 2, ..., S$ with initial condition $\hat{b}_{S+1} = v_{S+1}$. The only restriction this entails on $\hat{b}_1$ is that $\hat{b}_1 > \hat{b}_2$. Finally, note that (2) coincides with EOS' lowest envy free equilibrium (EOS, Theorem 2), and with Varian's lower-bound symmetric Nash Equilibrium (Varian, 2007, eq.9).

A.2.2 Proof of Theorem 1

We prove the statement by induction on the size of the coalition. The **induction basis** is the non-collusive benchmark (i.e., $|C| = 1$). In this case all players use their dominant strategies, $b_i = v_i$ for each i, which clearly ensures $v_i \in (b_{i+1}, v_{i-1})$ for all i, and the equilibrium bids profile is as claimed in the Theorem.

For the **inductive step**, suppose we have shown that the result holds for all coalitions C' such that $C' \subseteq C$. We want to show that it also holds for C. Let i be the lowest bidder in the coalition, and let r denote his position. Then, his payoff is equal to:

$$u_i = v_i x^r - \sum_{t=r+1}^{S+1} b^t \left(x^{t-1} - x^t\right).$$

It is useful to introduce notation to rank independent among themselves, based on their valuation. Let $v_{I\backslash C} = (v_j)_{j\in I\backslash C}$, and let $v_{I\backslash C}(k) = v_{I\backslash C}^{|I\backslash C|+1-k}$ denote the valuation of the k-th lowest value independent: for $k = 1$, $v_{I\backslash C}(1) = v_{I\backslash C}^{|I\backslash C|}$ is the lowest valuation among the independents, $v_{I\backslash C}(2) = v_{I\backslash C}^{|I\backslash C|-1}$ is the second lowest valuation among the independents, and so on. Now, if i is the lowest-bidding member of the coalition, all players placing lower bids are independents, and therefore bid according to their dominant strategy, $b_j = v_j$. This in turn implies that bids in positions $t = r+1, ..., S+1$ are ranked efficiently between themselves, but it does not guarantee that $b^t = v_t$ for each $t \geq r+1$, unless all $j \in C$ are such that $j \leq r$. Thus, we conclude that bids b^t for $t = r+1, ..., S+1$ are placed by the $S+1-r$ lowest-valued independents. Hence,

$$u_i = v_i x^r - \sum_{t=r+1}^{S+1} v_{I\backslash C}(S+2-t)\left(x^{t-1} - x^t\right). \tag{17}$$

Let us consider the function $\tilde{u}_i(k)$ of i's payoff, as a function of the position k he occupies, given that he is the lowest-bidder in the coalition. Let $u_i^* := \max_k \tilde{u}_i(k)$. Clearly, $u_i^* \geq u_i$. We show next that, if $i \neq \max\{j : j \in C\}$, then $u_i^* < u_i^{C\backslash\{i\}}$ (the payoff i would obtain by leaving the coalition). Hence, the coalition is stable only if the lowest bidding member is also the member with the lowest valuation.

First we show that $\tilde{u}_i$ is maximized only if i is placed efficiently with respect to the independents. That is, for any $j \in I\backslash C$, $j < i$ if and only if $\rho(j) < r$. We proceed by contradiction: suppose that there exist $j \in I\backslash C$ such that either $j < i$ and $\rho(j) > r$, or $j > i$ and $\rho(j) < r$. Consider the first case: Since independents are ranked efficiently among themselves, for any $j, l \in I\backslash C$, $l < j$ if and only if $\rho(l) < \rho(j)$. It follows that if there exists $j \in I\backslash C : j < i$ and $\rho(j) > r$, such j can be chosen so that $j = r+1$, i.e. j occupies the position immediately following i's. We next show that, in this case, i's payoff would increase if he dropped one position down. To see this, notice that

$$\begin{aligned}\tilde{u}_i(r+1) - \tilde{u}_i(r) &= v_i\left(x^{r+1} - x^r\right) + v_{I\backslash C}(S+1-r)\left(x^r - x^{r+1}\right)\\ &= \left(v_{I\backslash C}(S+1-r) - v_i\right)\left(x^r - x^{r+1}\right),\end{aligned}$$

where $v_{I\backslash C}(S+1-r) = v_{r+1}$ is the valuation of the highest independent if i occupies position r. Since, by assumption, $x^r > x^{r+1}$, it follows that

$$sign\left(\tilde{u}_i(r+1) - \tilde{u}_i(r)\right) = sign\left(v_{I\backslash C}(S+1-r) - v_i\right).$$

Under the absurd hypothesis, $v_{I\backslash C}(S+1-r) > v_i$, hence u_i increases dropping one position down. A similar argument shows that in the second case of the absurd hypothesis, i.e. if there exists $j \in I\backslash C : j > i$ and $\rho(j) < \rho(i)$, u_i could be increased climbing one position up, from r to $(r-1)$. The result obtains considering the difference

$$u_i(r) - u_i(r-1) = \left(v_{I\backslash C}(S+2-r) - v_i\right)\left(x^{r-1} - x^r\right),$$

which is negative, under the absurd hypothesis.

We have thus proved that, in equilibrium, for all $j \in I\backslash C$, $j < i$ if and only if $\rho(j) < r$. Hence, the lowest coalition bidder is placed efficiently with respect to the independents, and only independents are below him. Letting $\mathcal{J} = \{j \in C : j > i\}$ denote the set of coalition members with values lower than v_i, the lowest coalition bidder i therefore occupies position $i + |\mathcal{J}|$. (Clearly, i occupies the i-th position if and only if $\mathcal{J} = \emptyset$, i.e. if i, the lowest bidding member of the coalition, also has the lowest value in the coalition.) But then, setting $r = i + \mathcal{J}$ in eq. (17) , we have that

$$u_i^* = v_i x^{i+|\mathcal{J}|} - \sum_{t=i+|\mathcal{J}|+1}^{S+1} v_{I\backslash C}(S+2-t)\left(x^{t-1} - x^t\right). \tag{18}$$

We show next that $\mathcal{J} \neq \emptyset$ implies $u_i^* < u_i^{C\backslash\{i\}}$. For any k, let $\bar{b}_k$ denote k's bid in the equilibrium with coalition $C\backslash\{i\}$. Since, under the inductive hypothesis, the equilibrium with coalition $C\backslash\{i\}$ is efficient, $\bar{b}_k = \bar{b}^k$ for any k, and hence

$$u_i^{C\backslash\{i\}} = v_i x^i - \sum_{k=i+1}^{S+1} \bar{b}_k \left(x^{k-1} - x^k\right).$$

By the inductive hypothesis, the equilibrium with this smaller coalition is as in the Theorem's statement. Hence, $\bar{b}_k < v_{k-1}$ for all $k \in I$ (if k is an independent, because he bids $\bar{b}_k = v_k < v_{k-1}$; if he's the highest-value member of the coalition, because $\bar{b}_k \in \left(b_{k+1}^+, v_{k-1}\right)$, otherwise $\bar{b}_k = b_{k+1}^+ < v_{k-1}$). We also show that $\bar{b}_k \leq v_{I\backslash C}(S+2-k)$ for all k. To this end, observe that all $k \geq \max\{\mathcal{J}\}$ are independents (both before and after i drops out), so that for all $k \geq \max\{\mathcal{J}\}$, $\bar{b}_k = v_k = v_{I\backslash C}(S+2-k)$: these are the lowest bidding and the lowest-value bidders, hence also the lowest independents. For $k < \max\{\mathcal{J}\}$, at least one of the $S+2-k$ elements of the set $\{k, k+1, ..., S+1\}$ is a member of the coalition. It follows that the valuation of the $(S+2-k)$-th lowest independent is higher than v_k, hence $v_{I\backslash C}(S+2-k) \geq v_{k-1}$, which in turn implies $v_{I\backslash C}(S+2-k) > \bar{b}_k$. Overall, we have that $\bar{b}_k < v_{k-1}$ and $\bar{b}_k \leq v_{I\backslash C}(S+2-k)$ for all $k \in I$. Using the first inequality for $k \le i + |\mathcal{J}|$ and the second inequality otherwise, we see that if $\mathcal{J} \ne \emptyset$,

$$\begin{aligned} u_i^{C\backslash\{i\}} &= v_i x^i - \sum_{k=i+1}^{i+|\mathcal{J}|} \bar{b}_k \left(x^{k-1} - x^k\right) - \sum_{k=i+|\mathcal{J}|+1}^{S+1} \bar{b}_k \left(x^{k-1} - x^k\right) \\ &> v_i x^i - \sum_{k=i+1}^{i+|\mathcal{J}|} v_{k-1} \left(x^{k-1} - x^k\right) - \sum_{k=i+|\mathcal{J}|+1}^{S+1} v_{I\backslash C}(S+2-k) \left(x^{k-1} - x^k\right) \end{aligned} \quad (19)$$

Combining (18) and (19), we get

$$\begin{aligned} u_i^{C\backslash\{i\}} - u_i^* &> v_i \left(x^i - x^{i+|\mathcal{J}|}\right) - \sum_{k=i+1}^{i+|\mathcal{J}|} v_{k-1} \left(x^{k-1} - x\right) \\ &\ge v_i \left(x^i - x^{i+|\mathcal{J}|}\right) - v_i \left(x^i - x^{i+|\mathcal{J}|}\right) = 0, \end{aligned}$$

where the latter inequality follows because $v_{k-1} \le v_i$ for all $k \ge i+1$. Hence, whenever $\mathcal{J} \ne \emptyset$, we obtain $u_i < u_i^{C\backslash\{i\}}$: that is, the recursive stability condition (S.2) is violated for bidder i. $\mathcal{J} = \emptyset$ therefore is a necessary condition for equilibrium. Hence, in any equilibrium, the lowest coalition bidder also has the lowest valuation in the coalition. Moreover, if $\mathcal{J} = \emptyset$, $u_i^* = u_i^{C\backslash i}$ (by equations (18) and (19)), hence in equilibrium $u_i = u_i^*$ and $i = \rho(i)$:

$$u_i = v_i x^i - \sum_{k=i+1}^{S+1} v_k \left(x^{k-1} - x^k\right) = u_i^{C\backslash\{i\}}. \quad (20)$$

Furthermore, since the payment of coalition members above i is strictly decreasing in b_i and positions are independent of b_i (as long as $b_i \in (b_{i+1}, b_{i-1})$), the coalition will set b_i as low as possible to ensure i's efficient position. That is, $b_i = b_{i+1}^+ = v_{i+1}^+$.

We have determined the positions and bids of all bidders $k \ge i$. We know that the remaining coalition members are positioned above these bidders and do not affect u_i. Thus, the remaining task for the coalition is to choose bids $(b_j)_{j \in C\backslash\{i\}}$ in order to maximize $\sum_{j \in C\backslash\{i\}} u_j$, subject to the constraint that $b_j > b_i$ for all $j \in C\backslash\{i\}$. We now need to look separately at two cases: $|C| = 2$ and $|C| > 2$.

First, if $|C| = 2$, the task is simply to maximize the payoff of the other member of the coalition, j, by determining his position relative to the remaining independents. But this, by the usual argument, is achieved when j is placed efficiently with respect to these independents. This is achieved if and only if $b_j \in (b_{j+1}, v_{j-1})$.

Second, if $|C| > 2$, note that even when one of the members $j \in C\backslash\{i\}$ drops out, i still remains a non-top member of the coalition. Hence, its bid does not change. Naturally, the bids of all $k > i$ (who are independents) do not change either. Hence, the payoffs of all bidders $k < i$ both before and after one of the coalition members (other than i) drops out are shifted by the same constant relative to a game in which the bidders $k \ge i$ (and the corresponding slots) are removed: thus, the presence of these bidders has no effect on either the payoffs or the outside options. It follows that the problem we are solving at this stage is exactly equivalent to finding the equilibrium in the VCG game played between coalition $C \backslash \{i\}$ and independents $\{j \in I \backslash C : j > i\}$ with slots $x^1, ..., x^{i-1}$. This game has coalition size $C - 1$, so the solution follows by the inductive hypothesis.■

A.2.3 Proof of Theorem 2

Since the UC-restrictions imply the stability restriction (S.1), the agency's problem in the GSP auction with the feigned values restriction reduces to:

$$\begin{aligned} &\max_{b_C} u_C\left(b_C, \beta^*\left(b_C\right)\right) \\ \text{subject to} &: \text{(R) } \exists v'_C \in \mathbb{R}_+^{|C|} \text{ s.t. } \left(b_C, \beta^*(b_C)\right) \in E^*\left(v'_C, v_C\right) \\ &: \text{(S.2) } \forall i \in C, u_i\left(b_C, \beta^*(b_C)\right) \geq \bar{u}_i^{C\backslash\{i\}}. \end{aligned}$$

where the equilibrium conjectures β^* are such that,

$$\forall b_C, \beta^*\left(b_C\right) \in \left\{b^*_{-C} \in \mathbb{R}_+^{n-|C|} : \forall i \in I\backslash C, b^*_i \in BR^*_i\left(b_C, b^*_{-i,-C}\right)\right\}.$$

Let $\sim$ be an equivalence relation on $\mathbb{R}^n_+$ such that $v \sim v'$ (resp., $b \sim b'$) if and only if v and v' only differ in the highest valuation (resp., highest bid), but not in the *identity* of the highest valuation individual (bidder).[33] For any $v \in \mathbb{R}^n_+$, let $[v]$ (resp., $[v]$) denote the equivalence class of v (resp., b) under this equivalence relation, and let $\mathbb{V}^\sim$ (resp., $\mathbb{B}^\sim$) denote the set of such equivalence classes. Next, consider the competitive equilibrium correspondence $E^* : \mathbb{R}^n_+ \rightrightarrows \mathbb{R}^n_+$, which assigns to each profile $v \in \mathbb{R}^n_+$ the set $E^*(v)$ of competitive equilibria in the GSP auction. Denote the set of equivalence classes under $\sim$ on the range of E^* as $E^*(\mathbb{V}^\sim) \subseteq \mathbb{V}^\sim$, and let $E^\sim : \mathbb{V}^\sim \to E^*(\mathbb{V}^\sim)$ denote the function induced by E^*. Lemma 1 implies that $E^\sim$ is a bijection. Further note that the payoffs of all bidders in the GSP with bids $E^*(v)$ are the same as in the VCG with truthful bids:

$$\text{for all } v \in \mathbb{R}^n_+ \text{ and } i \in I,\ u_i^{\mathcal{V}}(v) = u_i^{\mathcal{G}}(E^*(v)). \tag{21}$$

Since $E^\sim$ is a well-defined function on the equivalence classes of $\sim$, the profile of valuation v'_C in the restriction (R) uniquely pins down $(b_C, b^*_{-C}) \in E^*(v'_C, v_{-C})$ up to the highest overall bid. That is, $(b_C, b^*_{-C}), (b'_C, b'_{-C}) \in E^*(v'_C, v_{-C})$ if and only if $(b_C, b^*_{-C}) \sim (b'_C, b'_{-C})$. Together with (21), this implies that $u_i^{\mathcal{G}}(b_C, b^*_{-C}) = u_i^{\mathcal{V}}(v'_C, v_{-C})$, so that also $u_C^{\mathcal{G}}(b_C, b^*_{-C}) = u_C^{\mathcal{V}}(v'_C, v_{-C})$. As a result, we can now easily recast the coalition's problem

[33]Formally: $v \sim v'$ if and only if the following two conditions hold: (1) $\arg\max_{i\in I} v_i = \arg\max_{i\in I} v'_i$; (2) $v_i = v'_i$ for all $i \neq \arg\max_{i\in I} v_i$.

as one of choosing v'_C (the coalition's 'feigned valuations'):

$$\max_{v'_C} u_C^{\mathcal{V}}\left(v'_C, v_{-C}\right)$$
$$\text{subject to : (S.2) } \forall i \in C, u_i^{\mathcal{V}}\left(v'_C, v_{-C}\right) \geq \bar{u}_i^{C\setminus\{i\}}.$$

(Notice that the restriction (R) and the restriction that $\beta^*(b_C)$ always be in the set BR^*_{-C} are both built in this formulation of the problem.)

Furthermore, $\bar{u}_i^C = \bar{u}_i^{C;\mathcal{V}}$ for all i when $|C| = 1$, and the recursion defining $\bar{u}_i^C$ is identical to that defining $\bar{u}_i^{C;\mathcal{V}}$. It follows that the coalition's problem is now equivalent to its problem in the VCG game. By Theorem 1, the solution v'^*_C is unique up to the report of the highest coalition member, $v'^*_{\min(C)}$.

Finally, by (R), the UC-RAE of the GSP satisfies $(b^*_C, \beta^*(b_{-C})) \in E^*(v'^*_C, v_{-C})$. Hence all bidders' positions and payoffs in this GSP equilibrium are the same as in the unrestricted RAE of the VCG, (v'^*_C, v_{-C}). Because the ordering of bidders in the RAE of the VCG is efficient (Theorem 1), so is the ordering of bidders in the the UC-RAE of the GSP. However, because v'^* is unique only up to the highest coalition bid, $(b^*_C, \beta^*(b_{-C}))$ is not uniquely defined: there exists a continuum of equilibria differing in the payments of all bidders above the highest coalition bidder: for each $v'^*_{\min(C)} \in (v'^*_{\min(C)+1}, v_{\min(C)-1})$, there exists one equivalence class of UC-RAE of the GSP, $[(b^*_C, \beta^*(b_{-C}))]$. Because E^* is unique only up to the highest overall bid, there also exist a continuum of equilibria yielding the same payoffs and positions, but differing in the highest overall bid, within each $[b^*]$. In this sense, the equilibrium is unique up to the highest coalition and overall bids.

A.2.4 Proof of Theorem 3

The claim about the possibility of strict ordering in revenues is proven by Example 6 in the text. Here we prove the general claims about existence, uniqueness and weak ordering. The proof is by construction, and it is based on the following intermediate result.

Lemma 2 *Fix $C \subset I$, and let $\mathcal{K}$ be a finite index set. Let $\left\{b^{(k)}\right\}_{k\in\mathcal{K}}$ be a collection of bid profiles such that, for each $k \in \mathcal{K}$, $b^{(k)}_{-C} \in BR^*_{-C}(b^{(k)}_C)$ and $\rho(i; b^{(k)}) = i$ for each $i \in I$. Define $\mathcal{L}\left(\left\{b^{(k)}\right\}_{k\in\mathcal{K}}\right) \equiv \hat{b} \in \mathbb{R}^n_+$ as follows:*

$$\hat{b}_i = \begin{cases} \hat{b}_i = \min_{k\in\mathcal{K}} b_i^{(k)} & \text{if } i \in C \\ \hat{b}_i = v_{S+1} & \text{if } i = S+1 \notin C \\ \frac{1}{x^{i-1}}\left[\sum_{j=i}^{\bar{c}(i)-1} v_j(x^{j-1} - x^j) + \hat{b}_{\bar{c}(i)} x^{\bar{c}(i)-1}\right] & \text{otherwise} \end{cases};$$

where $\bar{c}(i) := \min\{j \in C \mid j > i\}$ if $i < \max C$ and $\bar{c}(i) = S+1$ otherwise.

Then: (i) $\rho(i; \hat{b}) = i \ \ \forall i \in I$; (ii) $u_i(\hat{b}) \geq u_i(b^{(k)})$ for all $i \in I$ and for all $k \in \mathcal{K}$, with strict inequality whenever $\hat{b}_{\bar{c}(i)} \neq b^{(k)}_{\bar{c}(i)}$; (iii) $u_C(\hat{b}) \geq u_C(b^{(k)})$ for all $k \in K$, with strict inequality whenever $\exists i \in C\setminus \min C$ such that $b^{(k)}_i \neq \hat{b}_i$; (iv) $\hat{b}_{-C} \in BR^*_{-C}(\hat{b}_C)$.

Proof of Lemma 2

We begin by noting that because for each $k \in \mathcal{K}$, $b_{-C}^{(k)} \in BR_{-C}^{*}(b_{-C}^{(k)})$ and $\rho(i; b^{(k)}) = i$ for each $i \in I$, we have that $\forall k \in \mathcal{K}, i \notin C$ s.t. $i \neq S+1$,

$$b_i^{(k)} = \frac{1}{x^{i-1}} \left[\sum_{j=i}^{\bar{c}(i)-1} v_j (x^{j-1} - x^j) + b_{\bar{c}(i)}^{(k)} x^{\bar{c}(i)-1} \right],$$

and $b_i^{(k)} = v_{S+1}$ if $i = S+1 \notin C$ ($\bar{c}(i)$ is defined in the statement in the Lemma.)

The following two key observations are now immediate:

1. For every $k \in \mathcal{K}$ and for every $i \in I$, $\hat{b}_i \leq b_i^{(k)}$: For $i \in C$, $\hat{b}_i \leq b_i^{(k)}$ by the definition of coalition bids in the statement of the lemma. For $i = S+1 \notin C$, $\hat{b}_i = v_{S+1} = b_i^{(k)}$ (the second equality is because the Lemma requires $b_{-C}^{(k)} \in BR_{-C}^{*}(b_C^{(k)})$). Finally, for $i \notin C$ s.t. $i \neq S+1$,

$$\hat{b}_i = \frac{1}{x^{i-1}} \left[\sum_{j=i}^{\bar{c}(i)-1} v_j (x^{j-1} - x^j) + \hat{b}_{\bar{c}(i)} x^{\bar{c}(i)-1} \right] \leq \frac{1}{x^{i-1}} \left[\sum_{j=i}^{\bar{c}(i)-1} v_j (x^{j-1} - x^j) + b_{\bar{c}(i)}^{(k)} x^{\bar{c}(i)-1} \right] = b_i^{(k)},$$

where the inequality follows because, by definition, $\bar{c}(i) \in C \cup \{S+1\}$ and hence $\hat{b}_{\bar{c}(i)} \leq b_{\bar{c}(i)}^{(k)}$). Note that the inequality is strict whenever $\hat{b}_{\bar{c}(i)} \neq b_{\bar{c}(i)}^{(k)}$.

2. For each $i \in I$, there exists $k \in \mathcal{K}$ such that $b_i = b_i^{(k)}$. For $i \in C$ this is immediate from the definition. For $i = S+1 \notin C$, $\hat{b}_i = v_{S+1} = b_i^{(k)}$ for all k (cf. previous point). For $i \notin C$ s.t. $i \neq S+1$, the result follows because $\bar{c}(i) \in C \cup \{S+1\}$, hence there exists $k \in \mathcal{K}$ such that $\hat{b}_{\bar{c}(i)} = b_{\bar{c}(i)}^{(k)}$, so that

$$\hat{b}_i = \frac{1}{x^{i-1}} \left[\sum_{j=i}^{\bar{c}(i)-1} v_j (x^{j-1} - x^j) + \hat{b}_{\bar{c}(i)} x^{\bar{c}(i)-1} \right] = \frac{1}{x^{i-1}} \left[\sum_{j=i}^{\bar{c}(i)-1} v_j (x^{j-1} - x^j) + b_{\bar{c}(i)}^{(k)} x^{\bar{c}(i)-1} \right] = b_i^{(k)},$$

We can now establish the lemma's results:

(i) $\rho(i; \hat{b}) = i$ for all $i \in I$: Let $i, j \in I$ be s.t. $i < j$. We show that $\hat{b}_i > \hat{b}_j$. By point 2 above, there exists $k \in \mathcal{K}$ such that $\hat{b}_i = b_i^{(k)}$. Because, by assumption, $b^{(k)}$ is ordered efficiently, $b_i^{(k)} > b_j^{(k)}$. By point 1, $b_j^{(k)} \geq \hat{b}_j$. Hence, $\hat{b}_i = b_i^{(k)} > b_j^{(k)} \geq \hat{b}_j$, as desired.

(ii) $u_i(\hat{b}) \geq u_i(b^{(k)})$ for all $i \in I$ and all $k \in \mathcal{K}$, with strict inequality if $\hat{b}_{\bar{c}(i)} \neq b_{\bar{c}(i)}^{(k)}$: Because i obtains its efficient position under both $\hat{b}$ (established in (i)) and $b^{(k)}$ (given),

$$u_i(\hat{b}) = (v_i - \hat{b}_{i+1}) x^i \geq (v_i - b_{i+1}^{(k)}) x^i = u_i(b^{(k)}),$$

where the inequality holds because $\hat{b}_{i+1} \leq b_{i+1}^{(k)}$ by point 1 above, with strict inequality if $\hat{b}_{\bar{c}(i)} \neq b_{\bar{c}(i)}^{(k)}$, as noted at the end of point 1.

(iii) $u_C(\hat{b}) \geq u_C(b^{(k)})$ for all $k \in K$, with strict inequality whenever $\exists i \in C \backslash \min C$ such that $b_i^{(k)} \neq \hat{b}_i$: The weak inequality follows immediately from part (ii). Now, suppose

$b_i^{(k)} \neq \hat{b}_i$ for some $i \in C \backslash \min C$, and let $j = \max\{k \in C | k < i\}$ be the coalition member directly above i in the ranking of valuations. Then $\bar{c}(j) = i$, so that by the strict inequality part of result (ii), $u_j(b^{(k)}) < u_j(\hat{b})$. Since $u_{j'}(b^{(k)}) \leq u_{j'}(\hat{b})$ for all other terms in the sums defining $u_C(\cdot)$, this completes the proof for strict inequality.

(iv) $\hat{b}_{-C} \in BR^*_{-C}(\hat{b}_{-C})$: The LREF condition holds by construction. We must simply prove the Nash condition, i.e., that each $i \notin C$ (weakly) prefers position i to position j for all $j \neq i$. Define $j' = j+1$ if $j > i$ and $j' = j$ if $j < i$. Note that if bidder i deviates to position $j \neq i$ under bid profile $\hat{b}$, it gets payoff $(v_i - \hat{b}_{j'})x^j$. By the observation in point 2 above, there exists some k such that $\hat{b}_{j'} = b_{j'}^{(k)}$, so that $(v_i - \hat{b}_{j'})x^j = (v_i - b_{j'}^{(k)})x^j$. Because $b_{-C}^{(k)} \in BR^*_{-C}(b_{-C}^{(k)})$ and $\rho(i; b^{(k)}) = i$, i cannot profitably deviate from position i to position $j \neq i$ under bid profile $b^{(k)}$, i.e. $(v_i - b_{j'}^{(k)})x^j \leq (v_i - b_{i+1}^{(k)})x^i$. Finally, by point 1 above, $b_{i+1}^{(k)} \geq \hat{b}_{i+1}$, so that $(v_i - b_{i+1}^{(k)})x^i \leq (v_i - \hat{b}_{i+1})x^i$. Putting these results together,

$$(v_i - \hat{b}_{i+1})x^i \geq (v_i - b_{i+1}^{(k)})x^i \geq (v_i - b_{j'}^{(k)})x^j = (v_i - b_{j'})x^j.$$

That is, bidder i cannot profitably deviate to position $j \neq i$ under bid profile $\hat{b}$, as desired. This concludes the proof of the Lemma.■

Armed with this Lemma, we can now prove Theorem 3. We begin with existence and weak ordering of revenues, using induction on the coalition's size, C. For the **induction basis**, we use $|C| = 1$. Both existence and weak order now hold trivially, as both the efficiency-constrained RAE of the GSP and the RAE of the VCG mechanism are equal to the LREF equilibrium by definition.

For the **inductive step**, we fix C and suppose that for all coalitions of size $|C| - 1$ Eff-RAE exist, then we show that Eff-RAE also exists for C, and that in each of these RAE the coalition's surplus is no lower than in any RAE of the VCG mechanism, while the auctioneer's revenue is no higher than in a corresponding RAE of the VCG mechanism.

Fix C, and let $b^{UC} \in \mathbb{R}_+^n$ be the bids in the UC-RAE of the GSP auction with the same coalition C, in which the top coalition member is placing the highest possible bid (this exists, it is efficient and unique by Theorem 2). Observe that because of the bijection between UC-RAE of the GSP auction and unconstrained RAE of the VCG mechanism (established in Theorem 2), we can use the coalition's surplus in the GSP auction with bids b^{UC} as our reference point. Next, note that, for any b_C, the beliefs $\beta^*(b_C)$ in any Eff-RAE of the GSP auction are uniquely determined by the Varian/EOS recursion. Hence, a complete Eff-RAE, $(b^*, \beta^*) \in \mathbb{R}_+^n \times B^*$, if it exists, is in fact fully determined by $b_C^* \in \mathbb{R}_+^C$. We now proceed to prove that such a b_C^* exists by constructing a candidate profile.

For each $i \in C$, let $b^{(i)}$ be the bids in an Eff-RAE with coalition $C \backslash \{i\}$ (these exist under the inductive hypothesis). Let $b^{(0)} = b^{UC}$. Let $\hat{b} = \mathcal{L}\left(\{b^{(i)}\}_{i \in C \cup \{0\}}\right)$, where $\mathcal{L}$ is as defined in Lemma 2. Now, by results (i) and (iv) of Lemma 2, we have $\rho(i; \hat{b}) = i$ for all $i \in I$ and $\hat{b}_{-C} \in BR^*_{-C}(\hat{b}_C)$. It follows that $\hat{b}_C \in R_C^{EFF}$. By result (ii) of Lemma 2, $u_i(\hat{b}) \geq u_i(b^{(k)})$ for each i. Moreover, by construction, $u_i(b^{(k)}) = \bar{u}_i^{C \backslash \{i\}}$ for each $i \in C$, hence profile $\hat{b}$ satisfies the recursive stability condition. It follows that $\hat{b}_C$ is a valid bid vector for coalition C trying to achieve an Eff-RAE and that $\hat{b}_{-C} = \beta^*(\hat{b}_C)$, where β^* are the unique beliefs consistent with Eff-RAE. Maintaining the assumption of finite bid increments, as in Theorems 1 and 2, the coalition is therefore maximizing over a non-empty, finite set of valid bid vectors, so that a maximum, b^*_C, exists. Thus, an efficiency constrained RAE for coalition C exists (and is equal to $((b^*_C, \beta^*(b^*_C)), \beta^*)$).

Now the weak ordering of coalition surplus is immediate: Result (iii) of Lemma 2 implies $u_C(\hat{b}) \geq u_C(b^{UC})$, and clearly the optimal bid profile $(b^*_C, \beta^*(b^*_C))$ must satisfy $u_C(b^*_C, \beta^*(b^*_C)) \geq u_C(\hat{b})$. It follows that $u_C(b^*_C, \beta^*(b^*_C)) \geq u_C(b^{UC})$.

Next, we establish the ordering for the auctioneer's revenues. We first show that, in the Eff-RAE (b^*, β^*), the bid of coalition members other than the highest-valuation is weakly lower than in $\hat{b}$. To this end, suppose that there exists some $i \in C \backslash \min C$ such that $b^*_i > \hat{b}_i$. Let $b' = \mathcal{L}(\{b^*, \hat{b}\})$. By part (i) of Lemma 2, b'_C is still a valid bid vector for the coalition, whereas part (iii) implies $u_C(b'_C, \beta^*(b'_C)) > u_C(b^*_C, \beta^*(b^*_C))$ which contradicts the optimality of b^*_C. We thus conclude that $b^*_i \leq \hat{b}_i$ for all $i \in C \backslash \min C$.

Because the independents' bids are fixed by the recursion under both $\hat{b}$ and b^*, we know that in fact $b^*_i \leq \hat{b}_i$ for all $i > \min C$. Because by construction $\hat{b}_i \leq b^{UC}_i$ for all $i \in I$, we thus have $b^*_i \leq b^{UC}_i$ for all $i > \min C$. If $\min C = 1$, this completes the proof that the auctioneer's revenues are weakly lower under b^* than under b^{UC}. If $\min C > 1$, we need to show that even the top coalition bidder in b^* cannot bid more than this bidder's maximum possible UC-RAE bid. Because $b^{UC}_{\min C}$ is the maximum bid that the top coalition bidder can place in a UC-RAE, it is equal to (cf. Theorem 2)

$$b^{UC}_{\min C} = v_{\min C-1} - \frac{x^{\min C}}{x^{\min C-1}} \left(v_{\min C-1} - b^{UC}_{\min C+1}\right).$$

If $b^*_{\min C} > b^{UC}_{\min C}$, then the independent above the top coalition member obtains a payoff

$$U_0 = (v_{\min C-1} - b^*_{\min C}) x^{\min C-1} < (v_{\min C-1} - b^{UC}_{\min C}) x^{\min C-1} = (v_{\min C-1} - b^{UC}_{\min C+1}) x^{\min C},$$

where the last inequality follows by substituting in the expression for $b^{UC}_{\min C}$ from above.

Dropping one position down this independent would obtain

$$U' = (v_{\min C-1} - b^*_{\min C+1}) x^{\min C} \geq (v_{\min C-1} - b^{UC}_{\min C+1}) x^{\min C} > U_0,$$

where the first inequality follows because $b^*_i \leq b^{UC}_i$ for all $i > \min C$, as established above. Thus this independent has a profitable deviation; a contradiction. We conclude that $b^*_{\min C} \leq b^{UC}_{\min C}$. But then, by the independents' recursion, we also have $b^*_i \leq b^{UC}_i$ for all $i \leq \min C$. Because we already knew that the $b^*_i \leq b^{UC}_i$ for all $i > \min C$, we have established that all bids in b^* are weakly lower than in b^{UC}, which completes the claim about the auctioneer's revenues.

Next, we show that the Eff-RAE is unique up to the highest coalition bid. To this end, fix some coalition $C \subseteq I$ and let b^{R1} and b^{R2} be two (possibly equal) Eff-RAE for C. Let $\hat{b} := \mathcal{L}(\{b^{R1}, b^{R2}\})$. By results (i), (iii) and (iv) of Lemma 2, $\hat{b}$ is still efficiently ordered and $\hat{b}_{-C} \in BR^*_{-C}(\hat{b}_C)$, so that $\hat{b}_C$ is in the set of permitted bids for the coalition in the efficiency-constrained problem without the recursive stability restriction, with $\hat{b}_{-C} \in \beta^*(\hat{b}_C)$. Furthermore, by result (ii) of Lemma 2, each coalition member is at least as well off under $\hat{b}$ as under b^{R1} and b^{R2}. Therefore, the fact that b^{R1} and b^{R2} satisfy the recursive stability condition implies that so does $\hat{b}$. The optimality of b_C^{R1} and b_C^{R2} in this set therefore implies that $u_C(\hat{b}) \leq u_C(b^{Rk})\ \forall k \in \{1,2\}$. But result (iii) of Lemma 2 then implies that $\hat{b}_i = b_i^{R1} = b_i^{R2}$ for all $i \in C \setminus \min C$.

Combining these results yields $b_i^{R1} = b_i^{R2} = \hat{b}_i$ for all $i \in C \setminus \min C$. Because coalition bids also uniquely determine independents' bids, the Eff-RAE is thus unique up to the highest coalition bid. This completes the proof.

A.3 Multiple Agencies

A.3.1 Formal definition

We consider the case with two SEMAs, which coordinate the bids of subsets $C_1, C_2 \subseteq I$ of bidders, s.t. $C_1 \cap C_2 = \emptyset$. Similar to the baseline notion with a single SEMA, the definition of RAE with multiple agencies is recursive, with the outside option of coalition member $i \in C_1$ being defined as his equilibrium payoff in the game with coalitions $(C_1 \setminus \{i\}, C_2)$. Hence, the recursion in the RAE with multiple coalitions involves, for every C_g, a recursion similar to the one for the single SEMA, but with initial condition set by the RAE in which C_{-g} is the only coalition.

Let $G(v) = (A_i, u_i)_{i=1,\ldots,n}$ denote the baseline game (e.g., GSP or the VCG), given the profile of valuations $v = (v_i)_{i \in I}$. For any $C_1, C_2 \subseteq I$ with $|C_g| \geq 2$ and $C_1 \cap C_2 = \emptyset$, we let $C := C_1 \cup C_2$. For each $g = 1, 2$, coalition C_g chooses a vector of bids $b_{C_g} = (b_j)_{j \in C_g} \in \times_{j \in C_g} A_j$, and let $b_C = (b_{C_1}, b_{C_2})$. Given b_C, independents $i \in I \backslash C$ simultaneously choose bids $b_i \in A_i$. We let $b_{-C} := (b_j)_{j \in I \backslash C}$ and $A_{-C} := \times_{j \in I \backslash C} A_j$. Given profiles b or $b_{-\mathcal{C}}$, we let $b_{-i,-C} := (b_j)_{j \in I \backslash C : j \neq i}$. As above, each SEMA maximizes the sum of the payoffs of its members, $u_{C_g}(b) := \sum_{i \in C_g} u_i(b)$, under the three constraints from the single-agency model, given conjectures about both the independents and the other coalition.

Stability-1: (Stability w.r.t. Independents) For any $i \in I \backslash C$, let $BR_i^* : A_{-i} \rightrightarrows A_i$, $BR^*_{-C} : A_C \rightrightarrows A_{-C}$ and S_C be defined as in the single-agency case (except now $C = C_1 \cup C_2$.) For each agency C_g, we let

$$S_{C_g} = \left\{ b_{C_g} \in A_{C_g} : \exists b_{C_{-g}} \in A_{C_{-g}} \text{ s.t. } \left(b_{C_g}, b_{C_{-g}}\right) \in S_C \right\},$$

Stability-2: ((Recursive) Stability w.r.t. Coalition Members) Let B^* be defined as in the single-agency case. Letting $E^{\mathcal{R}}(C_1, C_2)$ denote the set of *Recursively Stable Agency Equilibrium (RAE)* outcomes of the game with coalitions C_1 and C_2, given restrictions $\mathcal{R}$ (and refinement BR_i^*), we initialize the recursion setting $E^{\mathcal{R}}\left(C'_g, C_{-g}\right) =$

$E^{\mathcal{R}}(C_{-g})$ if $|C_g'| = 1$ (that is, if an agency controls only one bidder, then the RAE are the same as when there exists only the other agency). Suppose next that $E^{\mathcal{R}}(C_g', C_{-g})$ has been defined for all subcoalitions $C_g' \subset C_g$. For each $i \in C_g$, and $C_g' \subseteq C_g \backslash \{i\}$, let $\bar{u}_i^{C_g', C_{-g}} = \min_{b \in E^{\mathcal{R}}(C_g', C_{-g})} u_i(b)$. The second stability requirement therefore requires $u_i \geq \bar{u}_i^{C_{-g} \backslash \{i\}, C_g}$. Finaly, we define the set of 'Rational Conjectures' about the Opponent Coalition as $B_g^* = \left\{ \beta_g \in (A_{C_{-g}})^{\overline{S}_{C_g}} : \beta_g(b_{C_g}) \in BR_{-g}^C(b_{C_g}) \text{ for all } b_{C_g} \in \overline{S}_{C_g} \right\}$, where $\overline{S}_{C_g} = \left\{ b_{C_g} \in S_{C_g} : BR_{-g}^C(b_{C_g}) \neq \emptyset \right\}$, and

$$\begin{aligned}
BR_{-g}^C(b_{C_g}) = &\arg\max_{b_{C_{-g}}} u_{C_{-g}}\left(b_{C_g}, b_{C_{-g}}, \beta\left(b_{C_g}, b_{C_{-g}}\right)\right) \\
\text{subject to} &: \text{(R) } (b_{C_g}, b_{C_{-g}}) \in R_C \\
&: \text{(S.1) } (b_{C_g}, b_{C_{-g}}) \in S_C \\
&: \text{(S.2) for all } i \in C_{-g},\ u_i\left(b_{C_g}, b_{C_{-g}}, \beta\left(b_{C_g}, b_{C_{-g}}\right)\right) \geq \bar{u}_i^{C_{-g} \backslash \{i\}, C_g}
\end{aligned}$$

Definition 4 *A* Recursively Stable Agency Equilibrium (RAE) *of the game G with coalition structure (C_1, C_2), given restrictions $\mathcal{R}$ and independents' equilibrium refinement BR^*, is a profile of bids and conjectures $(b^*, \beta^*, \beta_1^*, \beta_2^*) \in A_C \times B^* \times B_1^* \times B_2^*$ such that:*

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

2. *The conjectures of the agencies are correct and consistent with the exogenous restrictions: $\beta^*(b_C^*) = b_{-C}^*$, and, for each $g \in \{1,2\}$, $\beta_g^*\left(b_{C_g}^*\right) = b_{C_{-g}}^*$, and $\left(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*\left(b_{C_g}, \beta_g^*(b_{C_g})\right)\right) \in R(C)$ for all $b_{C_g} \in R_{C_g}$.*

3. *Each agency best responds to the conjectures β^* and β_g^*, given the exogenous restrictions (R) and the stability restrictions about the independents and the coalition members (S.1 and S.2, respectively): For each $g = 1, 2$*

$$\begin{aligned}
b_{C_g}^* \in &\arg\max_{b_{C_g}} u_{C_g}\left(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*\left(b_{C_g}, \beta_g^*(b_{C_g})\right)\right) \\
\textit{subject to} &: \textit{(R) } \left(\left(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*\left(b_{C_g}, \beta_g^*(b_{C_g})\right)\right)\right) \in R_C \\
&: \textit{(S.1) } \left(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*\left(b_{C_g}, \beta_g^*(b_{C_g})\right)\right) \in S_C \\
&: \textit{(S.2) for all } i \in C_g,\ u_i\left(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*\left(b_{C_g}, \beta_g^*(b_{C_g})\right)\right) \geq \bar{u}_i^{C_g \backslash \{i\}, C_{-g}}
\end{aligned}$$

The set of RAE outcomes *for the game with coalitions (C_1, C_2) (given BR^* and R_C) is:*

$$E^{\mathcal{R}}(C_1, C_2) = \left\{ b^* \in A : \exists \beta^*, \beta_1^*, \beta_2^* \textit{ s.t. } (b^*, \beta^*, \beta_1^*, \beta_2^*) \textit{ is a RAE} \right\}. \tag{22}$$

Note that the definition above does not uniquely pin down the the bid of the top bidder of the "lower" coalition. To remove this ambiguity, in the following we break these ties by making this coalition member bid as if it were an independent, whenever such bids are still in the optimal set.

A.3.2 Proof of Theorem 4

We prove the theorem by providing a precise characterization of the RAE in the VCG and the UC-RAE of the GSP. That is, we show that with two coalitions, C_1 and C_2, the following statements hold:

1. If in the overall value ranking no member of one coalition is adjacent to a member of the other coalition, then:

   (a) There exists a unique RAE of the VCG mechanism. In this equilibrium, the bid profile $\hat{b}^V$ is such that

   $$\hat{b}_i^V = \begin{cases} v_i & \text{if } i \in (I \backslash C) \cup \min C_1 \cup \min C_2; \\ (\hat{b}_{i+1}^V)^+ & \text{if } i \in C \backslash \{\min C_1 \cup \min C_2\} \text{ and } i \leq S; \end{cases} \tag{23}$$

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

   (b) There exists a unique UC-constrained RAE of the GSP auction. In this equilibrium, for every i,

   $$\hat{b}_i^G = v_i^f - \frac{x^i}{x^{i-1}} \left(v_i^f - \hat{b}_{i+1} \right),$$

   where v_i^f is equal to bidder i's bid (reported value) in the VCG mechanism (as described in Statement 1 above): $v_i^f = \hat{b}_i^V$.

2. If in the overall value ranking a non-top member of one coalition is directly above a non-top member of the other coalition (i.e., there exist i and $i+1$, such that $i \in C_j$, $i+1 \in C_{j'}$, $j \neq j'$, $i \neq \min C_j$, and $i+1 \neq \min C_{j'}$), then no unconstrained RAE of the VCG and no UC-RAE of the GSP exist.

Below, we prove the results for the VCG (statement 1(a) and the VCG part of statement 2 above). The proofs of the GSP results are analogous.

First we show that, regardless of whether there are or are not adjacencies in the value rankings, an arrangement like that in statement 1(a) is the only possible RAE of the VCG. We then show that this candidate is in fact an equilibrium when there are no adjancies, but not when there are adjacencies involving non-top bidders.[34]

[34]Compared to the single agency case, the part of the proof that parallels Theorem 1 has two complications. First, the placement of the highest bidder of the the coalition that does not have the top overall bidder requires some additional technicality, as this placement is not only relative to independents but also relative to the other coalition's bidders. Second, the candidate equilibrium produced by the recursion still needs to be verified, because the recursive procedure does not guarantee that a coalition's bidders are best-responding to those bidders of the rival coalition that are placed below them. It is precisely this verification step that will yield the fundamental difference between the cases with and without members from different coalitions that are adjacent in the value ranking.

Before proceeding to the proof, it pays to make two observations about the best-response correspondences BR_g^C:

Observation 1: The best-response function of any coalition requires that each non-top member of the coalition bid just above the bid below. Formally, let $i \in C_g \setminus \min C_g$ and let b be such that $b_{C_g} \in BR_g^C(b_{C_{-g}})$. Then $b_i = (b^{\rho(i)+1})^+$.

Proof of Observation 1: Suppose $b_i \neq (b^{\rho(i)})^+$, and let $\delta = b_i - b^{\rho(i)}$. Now note that in the definition of BR_g^C, coalition g takes the bids of the other coalition (and the independents) as fixed. Thus, lowering b_i to $b^{\rho(i)+1} + \delta/2 < b_i$ does not change the allocation, but reduces the payments of all higher-ranked members of C_g by $(\delta/2)(x^{\rho(i)-1} - x^{\rho(i)}) > 0$, and is therefore a profitable deviation for C_g, a contradiction. $\square$

Observation 2: The best-response function of any coalition requires that no member of the coalition (top or non-top) be placed above a bidder bidding higher than this member's value. Formally, if $i \in C_g$ and $b_{C_g} \in BR_g^C(b_{C_{-g}})$, then $v_i \geq b^{\rho(i)+1}$.

Proof of Observation 2: Suppose $v_i < b^{\rho(i)+1}$), and consider the deviation where C_g b_i to $(b^{\rho(i)+1})^-$. Note that this deviation improves i's individual payoff by $(b^{\rho(i)+1} - v_i)(x^{\rho(i)} - x^{\rho(i)+1}) > 0$. Also observe that the deviation decreases the payments of higher-ranked coalition members (if any) by $(b_i - b^{\rho(i)+1})(x^{\rho(i)+1} - x^{\rho(i)}) > 0$. Thus, the deviation is unambiguously profitable for the coalition. $\square$

With these observations in hand, we proceed to the proof of Theorem 4.

As in Theorem 1, the proof is by recursion on the overall size of the coalition, $|C| = |C_1| + |C_2|$. The induction basis is the case of no coalitions ($|C| = 2$, i.e., $|C_1| = |C_2| = 1$), for which the result holds trivially, by EOS. For the inductive step, we first look at the overall lowest placed coalition bidder, i. The same argument as in the proof of Theorem 1 shows that, due to the recursive stability condition, this bidder is in fact the lowest-valued bidder among all coalition bidders ($i = \max(C_1 \cup C_2)$) and that it must occupy its efficient position ($\rho(i) = i$). The rationale is the same as in Theorem 1: because there are only independents below this bidder, j cannot be compensated by the rest of the coalition for taking an inefficient position (which the individually bidder prefers). Furthermore, by Observation 1 above, $b_i = v_{i+1}^+$.

Just as in the proof of Theorem 1, after fixing the lowest coalition bidder's bid, we can essentially remove this bidder and all lower-valued independents from the analysis and proceed to the next-lowest placed coalition bidder. Unless this bidder is the top bidder of a coalition, the same argument as in the proof of Theorem 1 again applies to show the bidder is placed in its efficient position. In addition, by Observation 1, it is bidding just above the value of the bidder just below. We then move to the next-lowest-placed coalition bidder.

Now, suppose we reach the top bidder of some coalition, bidder i. As in Theorem 1, this bidder must simply set its bid so as to maximize its own payoff (as there are no other coalition members above, whose payoffs it would affect). As in Theorem 1, this bidder cannot be placed directly above a higher-valued independent or directly below a lower-valued independent, by the standard EOS argument (e.g., when placed directly above j with $v_j > v_i$, i can increase its payoff by $\Delta x(v_j - v_i) > 0$ if it drops one position down). Unlike Theorem 1, however, this does not necessarily guarantee the efficient placement of i, as i could be placed directly below a lower-valued member of the other coalition (i cannot be placed above a higher-valued member of the other coalition, because, by construction, i is the lowest-placed remaining member of C, with all previous members placed in their efficient positions).

To rule out this remaining possibility, suppose i is placed directly below the other coalition's member j, with $v_i > v_j$. By Observation 1, this means that $b_i < v_j < v_i$. But consider the deviation where bidder i's bid is changed to $b_i' = v_j^+ > b_i$ (note also that $b_i' < v_i$ because $v_i > v_j$). By Observation 2, this deviation causes the other coalition to move bidder j (and any other members with with values below b_i') below bidder i, reducing their bids to no more than b_i'. Consequently, bidder i gains at least one position, which happens at a price that is less than v_i. Therefore, bidder i's payoff increases by (at least) $(v_i - b_i'^{\rho(i)-1} - x^{\rho(i)}) > 0$. The deviation is thus profitable.

This completes the proof that the top bidder of each coalition must occupy its efficient position and will therefore bid its true value (by the assumed equilibrium selection).

We now can repeat the above arguments for all remaining coalition bidders until all of their bids are fixed. We have thus proved that the only possible equilibrium has all bidders placed efficiently, with bids as specified in the theorem statement.

We next verify that this candidate is in fact an equilibrium when no members of different coalitions are adjacent. Note that for the top bidders of both coalitions this is equivalent to checking that they do not have any individually profitable deviations (because their bids and positions relative to bidders outside of their coalition do not affect the payoffs of the other members off their coalition), and for non-top bidders any deviation must also be weakly profitable individually, as they are already held to their outside options in the candidate equilibrium. Also, because inefficient reversals within a coalition are never profitable for the coalition, we need to consider only deviations that preserve ranking within a coalition. Now, for deviations upward consider any coalition bidder i such that the bidder directly above is not a member of the same coalition. If i is its coalition's top bidder, then $b_i = v_i$ and hence $b_j > b_i = v_i$ for all bidders above i. Then the standard EOS argument shows that i does not have a profitable deviation upwards. If i is not a top bidder, then, by assumption, the bidder directly above i (that is, bidder $i-1$) is a higher-valued independent, so $b_{i-1} = v_{i-1} > v_i$, and again $b_j \geq b_{i-1} > v_i$ for all bidders above i. The standard EOS argument again shows that i does not have a profitable deviation upwards. For deviations downward consider any coalition bidder i such that the bidder directly below is not a member of the same coalition. If i is its coalition's top bidder, then $b_i = v_i$ and hence $b_j < b_i = v_i$ for all bidders below i. Then the standard EOS argument shows that i does not have a profitable deviation downwards. If i is not a top bidder, then, by assumption, the bidder directly below i (that is, bidder

$i+1$) is a lower-valued independent, so $b_{i+1} = v_{i+1} < v_i$, and again $b_j \leq b_{i+1} < v_i$ for all bidders below i. The standard EOS argument again shows that i does not have a profitable deviation downwards. This completes the proof of the theorem.

Finally, we show that there is no equilibrium if there are any cases where non-top members of different coalitions are adjacent to each other. That is, suppose that $v_i \in C_j$ and $v_{i+1} \in C_k \neq C_j$, with $v_i \neq \min C_j$ and $v_{i+1} \neq \min C_k$. By the first part of the proof, we know that the only candidate equilibrium has i and $i+1$ placed in their efficient positions, with $b_{i+1} = b^+_{i+2} < v_{i+1}$ and $b_i = b^+_{i+1} < v_{i+1}$ (recall that the statement about the magnitudes of the bids follows from Observation 1 about the best-response functions). However, it is obvious that b_{i+1} is not a (static) best response to b_i: if, holding b_j fixed, C_k deviates to setting $b'_{i+1} = b^+_i$, $i+1$'s individual payoff increases by $(v_{i+1} - b_i)(x^i - x^{i+1}) > 0$, without perceptibly increasing the payoff of other members of C_k. Thus, $b_{i+1} \notin BR_k(b_{C_j})$, i.e., we are not in a RAE.

References

1. Aghion, P., P. Antràs, E. Helpman, (2007). "Negotiating free trade," *Journal of International Economics*, 73, 1–30.

2. Association of National Advertisers (ANA), (2011). *Trends in Digital Agency Compensation*, 4th Edition, ANA publishing.

3. Athey, S., D. Nekipelov, (2012). "A Structural Model of Sponsored Search Advertising Auctions," mimeo.

4. Bachrach, Y., (2010), "Honor Among Thieves - Collusion in Multi-Unit Auctions," *Proc. of 9th int. Conf. on Autonomous Agents and Multiagent Systems* van der Hoek, Kaminka, Luck and Sen (eds.), Toronto Canada.

5. Balseiro, S. r., O. Besbes, G. Weintraub, (2015). "Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design," *Management Science*, 61(4), pp. 864-884.

6. Bernheim, B. D., M. D. Whinston, (1985). "Common Marketing Agency as a Device for Facilitating Collusion," *The RAND Journal of Economics*, 16(2), 269–281.

7. Bernheim, B. D., M. D. Whinston, (1986). "Common Agency," *Econometrica*, 54(4), 923–942.

8. Blake, T., C. Nosko, S. Tadelis, (2015). "Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment," *Econometrica*, 83(1):155-174.

9. Borgers, T., I. Cox, M. Pesendorfer, V. Petricek, (2013). "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence," *American Economic Journal: Micrma oeconomics*, 5, 163-187

10. Chassang, S. and J. Ortner (2016). "Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement," *mimeo*

11. Chassang, S. and J. Ortner (2017). "Making Corruption Harder: Asymmetric Information, Collusion, and Crime," *Journal of Political Economy*, forthcoming

12. Che, Y.K., Choi, S. and J. Kim, (2017). "An Experimental Study of Sponsored-Search Auctions," *Games and Economic Behavior*, 102, 20–43..

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

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

15. Celis, L. E., G. Lewis. M. M. Mobius, H. Nazerzadeh, (2015). "Buy-it-now or Take-a-chance: A New Pricing Mechanism for Online Advertising," *Management Science*, 60(12), pp.2927-2948.

16. Decarolis, F., M. Goldmanis, A. Penta, (2018). "Common Agency and Coordinated Bids in Sponsored Search Auctions, mimeo.

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

18. Edelman, B., M. Ostrovsky, M. Schwarz, (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242-259.

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

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

21. Einav, L., C. Farronato, J. Levin, N. Sundaresan, (2018). "Auctions versus Posted Prices in Online Markets, " *Journal of Political Economy*, 126(1), pp. 178-215.

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

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

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

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

26. Hendricks, K., R. Porter and G. Tan, (2008). "Bidding Rings and the Winner's Curse," *RAND Journal of Economics*, 39, 1018-1041.

27. Kaplow, L. (2011). "On the Meaning of Horizontal Agreements in Competition Law, " *California Law Review*, 99(3), pp. 683-818.

28. Levin, J., A. Scrkypacz, (2016). "Properties of the Combinatorial Clock Auction, "*American Economic Review* 106(9), pp. 2528-2551.

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

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

31. Milgrom, P., J. Mollner, (2018). "Equilibrium Selection in Multi-Player Games with Auction Applications," *Econometrica*, 86, 1, pp. 219-261

32. OECD (2017). "Algorithms and Collusion: Competition Policy in the Digital Age," `http://www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age`

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

34. 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`

35. Ray, D., R. Vohra, (1997). "Equilibrium Binding Agreements," *Journal of Economic Theory*, 73, pp. 30-78.

36. Ray, D., R. Vohra, (2014). "The Farsighted Stable Set," *Econometrica*, 83(3), pp. 977-1011.

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

38. Varian, H., C. Harris, (2014). "The VCG Auction in Theory and Practice " *American Economic Review: Papers & Proceedings*, 104(5), pp. 442-445.

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

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