



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Marketing Agencies and Collusive Bidding in Online Ad Auctions

Francesco Decarolis, Maris Goldmanis, Antonio Penta

To cite this article:

Francesco Decarolis, Maris Goldmanis, Antonio Penta (2020) Marketing Agencies and Collusive Bidding in Online Ad Auctions. Management Science

Published online in Articles in Advance 03 Apr 2020

. <https://doi.org/10.1287/mnsc.2019.3457>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Marketing Agencies and Collusive Bidding in Online Ad Auctions

Francesco Decarolis,^a Maris Goldmanis,^b Antonio Penta^{c,d}

^a Department of Economics and Innocenzo Gasparini Institute for Economic Research, Bocconi University, 20136 Milano, Italy; ^b Department of Economics, Royal Holloway, University of London, Egham TW20 0EX, United Kingdom; ^c Catalan Institution for Research and Advanced Studies, Universitat Pompeu Fabra and Barcelona Graduate School of Economics, 08005 Barcelona, Spain; ^d Toulouse School of Economics and Digital Center, 31015 Toulouse, France

Contact: francesco.decarolis@unibocconi.it,  <https://orcid.org/0000-0002-8547-7021> (FDC); maris.goldmanis@rhul.ac.uk (MG); antonio.penta@upf.edu,  <https://orcid.org/0000-0003-3796-8245> (AP)

Received: September 13, 2018

Revised: April 9, 2019

Accepted: July 2, 2019

Published Online in *Articles in Advance*:
April 3, 2020

<https://doi.org/10.1287/mnsc.2019.3457>

Copyright: © 2020 INFORMS

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, Microsoft Bing, and Yahoo!) and the Vickrey–Clarke–Groves (VCG) mechanism (used by Facebook). We find that despite its well-known susceptibility to collusion, the VCG mechanism outperforms the GSP auction in terms of both revenues and efficiency.

History: Accepted by Gabriel Weintraub, revenue management and market analytics.

Funding: A. Penta acknowledges financial support from the European Research Council Starting Grant [Grant 759424].

Keywords: collusion • digital marketing agencies • Facebook • Google • GSP • internet auctions • online advertising • VCG

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 United States: in 2017, for instance, the value of advertising on search engines alone exceeded \$50 billion in the United States with an annual growth of nearly 10% (PwC 2017). 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 network feeds, and so on. The most important auction formats in this market are the generalized second-price (GSP) auction (used, for instance, by Google, Microsoft Bing, Yahoo!, etc.) and the Vickrey–Clarke–Groves (VCG) mechanism (used by Facebook and Google for its contextual ads). 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. In contrast, the GSP auction—whose study was pioneered by Varian (2007) and Edelman et al. (2007) (EOS)—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. Both auction formats have been studied extensively. With few exceptions, however, existing models have largely ignored a major trend in this market: the rise of intermediaries operating on the bidding platforms.¹

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 agency networks (seven in the United States) that conduct all bidding activities through centralized agency trading desks (ATDs). In recent years, these agency networks have expanded their activities and contributed to a major increase in the market concentration, reaching global revenues that compare well even with those of tech giants, such as Google.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. This not only affects advertisers’ optimal bidding strategies, but it also has the potential to alter the very functioning of these auction formats.

This paper proposes a theoretical analysis of the impact of agency bidding on the two main auction formats. We find that the agency’s equilibrium bids are akin to implementing a certain form of collusion³ (even if none of its clients explicitly attempt it) and that, in this situation, the VCG outperforms the GSP in terms of both revenues and efficiency. This is a strong result because the VCG is typically considered 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 potential rationale for why Facebook’s recent adoption of the VCG mechanism was so successful despite the early surprise it provoked (e.g., Metz 2015) and for why the last few years have recorded a steady decline in ad prices.⁴ The fragility of the GSP auction that we uncover suggests that further changes may occur in this industry, raising important questions from both a market-design and an antitrust perspective.

A satisfactory model of agency bidding needs to satisfy at least two desiderata: first, it must allow collusive and competitive behavior to coexist because agencies in these auctions typically operate side by side with independent advertisers;⁵ second, it needs to be sufficiently tractable and amenable to direct comparisons to the existing benchmarks in the literature. To achieve these goals, we modify EOS and Varian’s baseline model by introducing a marketing agency that we model as a player choosing bids for its clients in order to maximize the total profits. Bidders that do not belong to the agency are referred to as “independents” and have the usual objectives. To overcome the curse of multiplicity in the GSP auction and ensure a meaningful comparison with the competitive benchmark, we introduce a refinement of bidders’ best responses that distills the individual-level underpinnings of EOS equilibrium and assume that independents place their bids accordingly. This stratagem enables us to maintain the logic of EOS refinement for the independents even if their equilibrium is not defined in the game with collusion. The marketing agency, in turn, makes a proposal of a certain profile of bids to its clients. The proposal is implemented if it is “recursively stable” in the sense that, anticipating the bidding strategies of others and taking into account the possible unraveling of the rest of the coalition, no client has an incentive to abandon the agency and bid as an independent. Hence, the clients’ outside options are equilibrium objects themselves and implicitly incorporate the restrictions entailed by the underlying coalition formation game.⁶

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 distinguished from EOS competitive equilibria by an external observer (e.g., the auction platform or an antitrust authority). We show that, under this constraint, the GSP auction is efficient and its revenues are identical to those obtained if the same agency bid in a VCG auction. We then relax this “undistinguishability constraint” and show that, even in the absence of allocative distortions, the GSP’s revenues are lower than those obtained in the VCG mechanism with the same agency configuration. Furthermore, once the undistinguishability constraint is lifted, efficiency is

no longer guaranteed by the GSP. Because the VCG is well known to be highly susceptible to collusion, finding that it outperforms the GSP in terms of both 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 its position or its payment. Given the rules of the GSP auction, the agency can benefit from this simple strategy only if two of its members occupy adjacent positions. But because of the GSP’s complex equilibrium effects, the agency can do more than that. For instance, suppose that the agency shades the bid of its lowest member with no direct impact on its other clients’ payments. Intuitively, if this bid is kept persistently lower, then the logic of EOS’s refinement suggests that the third highest bidder, who is an independent, would eventually lower its 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 second highest bid lowers the highest bidder’s payment) and for the *indirect effect* (lowering the fourth 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 that 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 with 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 not the typical case in the data, 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 nonexistence of pure equilibria, very much like the case of competitive (nonagency) bidding in a generalized first price (GFP) auction. But as seen in the early days of this industry, when the GFP was adopted, lack of pure equilibria may generate bidding cycles that 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 2000s, from the GFP to the GSP auction (cf. Edelman and Ostrovsky 2007). Hence, not only does agency competition not solve the problems with these auctions; it appears likely to exacerbate them.

The rest of the paper is organized as follows: Section 2 reviews the competitive benchmark. Section 3 introduces the agency model, and Section 4 presents the main results. Section 5 discusses the related literature and some extensions; Section 6 concludes.

2. Competitive Bidding in Online Ad Auctions

Online ad auctions are mechanisms to assign agents $i \in I = \{1, \dots, n\}$ to slots $s = 1, \dots, S$, $n \geq S$, where, for simplicity, we assume that $n = S + 1$ (the extension to $n \geq S$ is straightforward). In our case, agents are advertisers, and slots are positions for ads on a web page (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 $x^1 > x^2 > \dots > x^S > 0$. We also let $x^t = 0$ for all $t > S$. Finally, we let v_i denote the per-click valuation of advertiser i , and we label advertisers so that $v_1 > v_2 > \dots > v_n$. As in Varian (2007) and EOS, we maintain that valuations and CTRs are common knowledge. This complete information environment is the main benchmark for the literature on the GSP auction.⁷

2.1. Rules of the Auctions

In both VCG and GSP auctions, 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).⁸ When b is

clear from the context, we omit it and write simply $\rho(i)$. For any $t = 1, \dots, n$ and b or b_{-i} , we let b^t and b_{-i}^t denote the t -highest component of the vectors b and b_{-i} , respectively. 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. The resulting payoff, 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 bid profile $b \in \mathbb{R}_+^n$, can thus be written as $u_i^g(b) = (v_i - b^{\rho(i)+1})x^{\rho(i)}$.

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

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

2.2. Equilibria

Despite the relative complexity of its payment rule, bidding behavior in the VCG is very simple because truthful bidding (i.e., $b_i = v_i$) is a dominant strategy in this auction. 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 opponents, $b_{-i} = (b_j)_{j \neq i}$, partitions the space of i 's bids into $S + 1$ intervals of payoff-equivalent bids for bidder i . So, for each $b_{-i} \in \mathbb{R}_+^{n-1}$, let $\pi_i(b_{-i})$ denote i 's favorite position given b_{-i} .⁹ Then i 's best-response correspondence $BR_i: \mathbb{R}_+^{n-1} \rightarrow \mathbb{R}_+$ is such that, for every b_{-i} , $BR_i(b_{-i}) = (b_{-i}^{\pi_i(b_{-i})}, b_{-i}^{\pi_i(b_{-i})-1})$.

The GSP auction has many Nash equilibria (NEs; fixed points of the $\times_{i \in I} BR_i$). 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.¹⁰ As EOS showed, such equilibria induce the same allocations and payments as truthful bidding in the VCG, and they are fully characterized

by the following conditions: $b_1 > b_2$, $b_i = v_i$ for all $i > S$, and for all $i = 2, \dots, S$,

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

But EOS's refinement is not defined when agencies are present. We, thus, consider instead a refinement of the individual best response correspondences, which distills the individual-level underpinnings of EOS refinement. Formally, for any $b_{-i} \in \mathbb{R}_+^{n-1}$, let

$$BR_i^*(b_{-i}) = \left\{ b_i^* \in BR_i(b_{-i}) : \left(v_i - b_{-i}^{\pi_i(b_{-i})} \right) x^{\pi_i(b_{-i})} = (v_i - b_i^*) x^{\pi_i(b_{-i})-1} \right\}. \quad (2)$$

In words, of the many $b_i \in BR_i(b_{-i})$ that would grant player i the player's favorite position $\pi_i(b_{-i})$, the player chooses the bid b_i^* that makes the player 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 $\times_{i \in I} BR_i^*$ correspondence, given valuations v , are denoted as $E^*(v)$.

Lemma 1. For any $v = (v_i)_{i=1, \dots, n}$, $b \in E^*(v)$ if and only if b is an EOS equilibrium.

This lemma shows that EOS's equilibrium—originally defined as a refinement of the Nash equilibrium correspondence—can be equivalently defined as the fixed point of a refinement of individual best responses. Hence, BR_i^* provides a model of individual behavior that is consistent with EOS's equilibrium and that is well defined in our setting even if EOS's equilibrium is not. The next example is used repeatedly throughout this 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 EOS's lowest-revenue *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.

3. 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 4.3.

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. 1997). Unfortunately, these concepts have no bite in the GSP auction because it can be shown that EOS's equilibrium satisfies both refinements.

We 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 that ensues from taking the bids of the agency as given. The agency's proposals, however, can only be implemented if they are *stable* in two senses: (S.1) first, if they are consistent with the independents' equilibrium behavior, and (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 (Ray and Vohra 1997). Hence, given a coalition C , the outside option for each member $i \in C$ is the member's 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 is defined recursively. We thus call it the *recursively-stable agency equilibrium* (RAE).

We also consider a third constraint (**R**), which we formalize as a set $R(C) \subseteq A$, to accommodate the possibility that the agency exogenously discards certain bids. For instance, we consider the case of an agency whose primary concern is not being identified as inducing collusion (Section 4.2.1) or to induce efficient outcomes (Section 4.2.2). In those cases, $R(C)$ would be comprised, respectively, of only those profiles that are “undistinguishable” to an external observer as collusive or efficient.

3.1. The Recursively Stable Agency Equilibrium

Let $G = (A_i, u_i)_{i=1, \dots, n}$ 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 $i \in I \setminus C$ 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 $i \in I \setminus C$ simultaneously choose bids $b_i \in A_i$. We let $b_{-C} := (b_j)_{j \in I \setminus C}$ and $A_{-C} := \times_{j \in I \setminus C} A_j$. Finally, given profiles b or b_{-C} , we let $b_{-i, -C} := (b_j)_{j \in I \setminus C; j \neq i}$ denote the subprofile of bids

of all independents other than i . We assume that the agency maximizes the sum of its members' payoffs,¹¹ denoted by $u_C(b) := \sum_{i \in C} u_i(b)$ under the three constraints (S.1, S.2, and R) discussed, which we formally introduce next.

Concerning R, $\mathcal{R} = \{R(C)\}_{C \in \mathcal{C}}$ denotes the collection of exogenous restrictions for all possible coalitions, and for each C , $R_C \subseteq A_C$ denotes the coalition bids consistent with $R(C)$:

$$R_C := \{b_C \in A_C : \exists b_{-C} \in A_{-C} \text{ s.t. } (b_C, b_{-C}) \in R(C)\}. \quad (3)$$

Concerning S.1, for any $i \in I \setminus C$, 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 (2) in the GSP). Define the *independents' equilibrium correspondence* $BR_{-C}^* : A_C \rightrightarrows A_{-C}$ as

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

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 agency's optimization and assume that the agency can only choose bid profiles from the set¹²

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

Concerning S.2, the agency forms conjectures about how its bids b_C affect the independents' bids in the continuation game. Let $\beta : S_C \rightarrow A_{-C}$ represent such conjectures and define the set of conjectures that are consistent with the independents playing an equilibrium:

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

The second stability condition requires that, given conjectures β , no coalition member $i \in C$ has an incentive to leave and bid as an independent in the game with coalition $C \setminus \{i\}$. This constraint, thus, requires a recursive definition. To this end, first let $E^* := \{b \in \mathbb{R}_+^n : b_i \in BR_i^*(b_{-i}) \text{ for all } i \in I\}$ denote the set of equilibria in the game without coalition (given refinement BR_i^*). Then, letting $E^{\mathcal{R}}(C')$ denote the set of 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 $|C'| = 1$ (i.e., 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 sub-coalitions $C' \subset C$. For each $i \in C$, and $C' \subseteq C \setminus \{i\}$, let $\bar{u}_i^{C'} = \min_{b \in E^{\mathcal{R}}(C')} u_i(b)$. Then, recursively, we have the following.

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 that¹³

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

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

$$b_C^* \in \arg \max_{b_C} u_C(b_C, \beta^*(b_C))$$

$$\text{s.t. (R) } b_C \in R_C$$

$$\text{(S.1) } b_C \in S_C$$

$$\text{(S.2) for all } i \in C, u_i(b_C, \beta^*(b_C)) \geq \bar{u}_i^{C \setminus \{i\}}.$$

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

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

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

Before moving to the general results of the next section, we first illustrate the logic of this definition in the context of a simple example. In the example as well as in some results in Section 4, equilibrium bids are sometimes such that $b_i = b_{i+1}$ for some i . Because ties are broken according to bidders' labels (cf. endnote 8), in that case, bidder i obtains the position above $i + 1$. To emphasize this, we write $b_i = b_{i+1}^+$.¹⁴

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 2, 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: because bidder 1's payment is strictly decreasing in b_5 , the optimal solution for the agency is to lower b_5 as much as possible while ensuring that bidder 1 keeps the first position. Hence, any profile $b' = (b'_1, 25, 20, 10, 0)$ such that $b'_1 > 25$ is an (unconstrained) RAE when $C' = \{1, 5\}$, and the resulting payoffs are $u' = (450, 150, 100, 10, 0)$ with a total of 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: $C'' = \{1, 2, 5\}$. In this case, the (unconstrained) RAE bids are $b'' = (b_1'', 20^+, 20, 10, 0)$, where $b_1'' > 20$, which induces payoffs $u'' = (500, 150, 100, 10, 0)$ and a total of 650 for the coalition. To see that this is an 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 before. As for 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. But, in such a profile, 2's payoff would be 145, which is lower than $u_2 = 150$, the payoff 2 could obtain if 2 bid as an independent in the game with $C' = \{1, 5\}$. Hence, lowering b_2 to the point of obtaining a lower position would increase the overall coalition payoff (by decreasing bidder 1's payment) but would violate the stability constraint (S.2) for bidder 2. Hence, the optimal b_2'' is the lowest bid that ensures that bidder 2 maintains the second position.

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 because the 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 the client to abandon the agency, the entire coalition would be disrupted and full competition restored? The recursivity of the (S.2) constraint reflects these considerations. Finally, the example also shows that RAE outcomes in general are not Nash equilibria of the baseline game, nor of the game in which the coalition is replaced by a single player. Similar to Ray and Vohra (1997, 2014) equilibrium binding agreements (EBAs, which we discuss in Section 5), the stability restrictions affect the set of equilibrium outcomes, not merely as a refinement.

4. 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 (2) for the GSP and as the dominant (i.e., truthful) strategy in the VCG.

We first present the analysis of the VCG mechanism (Section 4.1) and then proceed to the GSP auction

(Section 4.2). Our main conclusion is that the VCG outperforms the GSP in terms of both revenues and allocative efficiency, thereby uncovering a striking fragility of the GSP with respect to agency bidding.

4.1. Agency Bidding in the VCG Mechanism

Our first result characterizes the unconstrained RAE of the VCG mechanism: it shows that they are unique up to the bid of the highest coalition member and that, in all such equilibria, 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, we have the following.

Theorem 1 (RAE in the VCG). *For any C , let $E(C)$ denote the unconstrained RAE of the VCG. Then $\hat{b} \in E(C)$ if and only if*

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

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

The unconstrained RAE of the VCG mechanism, therefore, are efficient with generally lower revenues than in the VCG's competitive benchmark. The efficiency result is due to the stability restrictions in RAE, which limits the agency's freedom to place bids. Restriction (S.2), in particular, requires that the agency's proposal gives no member of the coalition an incentive to abandon it and bid as an independent. Similar to the illustrative Example 2, a recursive argument further shows that the payoff that any coalition member can attain from abandoning the coalition is bounded below by the equilibrium payoffs in the baseline (coalition-less) game, in which assignments are efficient. The "Pigouvian" logic of the VCG payments, in turn, implies that such (recursive) participation constraints can only be satisfied by the efficient assignment of positions. As shown by Example 2, the recursive stability restriction (S.2) is key to this efficiency result.

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. Because, in the VCG mechanism, lowering the i th bid affects the price paid for all slots $s = 1, \dots, \min\{S + 1, i - 1\}$, even a small coalition can have a significant impact on the total revenues. 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.¹⁵ Hence, although 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 $\hat{b} = (\hat{b}_1, 4, 2^+, 2, 1)$. The resulting revenues are 86 as opposed to 96 of the competitive benchmark.

4.2. 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 which, to an external observer, are undistinguishable from a (competitive) EOS equilibrium: the “undistinguishable (from EOS) coordination” (UC) restriction. Theorem 2 shows that the equilibrium outcomes of the GSP with this restriction are exactly the same as the unrestricted RAE of the VCG mechanism. We lift the UC restriction in Section 4.2.2 and show that the GSP’s RAE may be inefficient and induce strictly lower revenues than their VCG counterparts. Moreover, the revenue ranking holds even if the agency is restrained from inducing allocative distortions (Theorem 3).

4.2.1. Undistinguishable Coordination: A VCG-Equivalence Result. Consider the following set of exogenous restrictions: for any $C \in \mathcal{C}$,

$$R^{UC}(C) := \left\{ \in A : \exists v_C^f \in \mathbb{R}_+^{|C|} \text{ s.t. } b \in E^*(v_C^f, v_{-C}) \right\}. \quad (9)$$

In words, $R^{UC}(C)$ is comprised of all bid profiles that could be observed as part of an EOS 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’s bidding strategy could not be distinguished from an EOS equilibrium (and, hence, detected as “collusive”) even if the independents had revealed their own valuations to the external observer.¹⁶

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 , let $v_{n+1}^f = 0$, and for each $i = n, \dots, 1$, let $v_i^f := v_{i+1}^f$ if $i \in C$ and $v_i^f = v_i$ if $i \notin C$. Then, in any RAE

of the GSP auction under the undistinguishable coordination (UC) restriction, the bid profiles \hat{b} are such that, for every i ,

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

where $\hat{b}_0 := \infty$ and $x^i/x^{i-1} := 0$ whenever $i > S$. Moreover, in each of these equilibria, advertisers are assigned to positions efficiently, and advertisers’ payments are the same as in the corresponding unrestricted RAE of the VCG mechanism (Theorem 1).

Hence, the UC-RAE of the GSP are unique up to the highest bid of the coalition and up to the highest overall bid, and they are equivalent to the (competitive) EOS equilibria for some profile $(v_i^f)_{i \in I}$ of feigned valuations (which satisfy $v_i^f = v_i$ for all $i \notin C$). Though notationally involved, the idea is simple and provides clear insights into 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 advertiser in the competitive benchmark for some valuation. The agency’s problem, therefore, boils down to “choosing” a feigned valuation for each of its members and bidding 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.¹⁷ 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 bid manipulation is weaker in the GSP than in the VCG mechanism, in which the payments for all positions above i are affected. Unlike the VCG mechanism, however, manipulating the bid of coalition member i also has an indirect effect on the bids of all the independents placed above i , who lower their bids according to the recursion in (10).

Example 4. Consider the environment of Example 3 with $C = \{1, 3\}$. Then, applying the formula in (10), the UC-RAE is $\hat{b} = (\hat{b}_1, 2.9, 1.8, 1.6, 1)$, which results in revenues 86. These are the same as in 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, because of the reduction in \hat{b}_3 , is only equal to $(b_3^{EOS} - \hat{b}_3) \cdot x_2 = 5$ (where b_3^{EOS} denotes 3's bid in the nonagency 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 (10) essentially replicate the complexity of the VCG payments: once the direct and indirect effects are combined, the resulting revenue loss is the same in the two mechanisms. This result also enables us to simplify the analysis of the impact of 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.

Corollary 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 resulting from agency bidding are larger if (i) the differences $(x_{i-1} - x_i)$ associated with 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 corollary, 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 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 from agency bidding are larger if the terms $(x^{t-1} - x^t)$ associated with agency members are larger, which is part (i) of the corollary. To understand part (ii), let i be an agency member such that $i+1$ is an independent. Because 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 is 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 Corollary 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 .

Corollary 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*

resulting from 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 preceding 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 corollary 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 because one might have expected that the agency would have a larger impact if it controlled the high-valuation bidders. We find that, in fact, the opposite is true when one controls for the increments $\Delta_s(x)$ and $\Delta_s(v)$.

4.2.2. Lifting the UC Restriction: Revenue Losses and Inefficiency.

As discussed in Section 4.1, 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 because it is undistinguishable from an EOS equilibrium. The next example shows that, when the undistinguishability 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 eight bidders and seven 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, bidder 5's bid would also have to be reduced to make the higher positions more attractive. But the reduction of bidder 6's bid in this example is large enough that bidder 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 bidder 6's bid, but indirectly, by attracting bidder 4 to the lower position.

Table 1. Summary of Results in Examples

Valuations	VCG	GSP (EOS)	RAE in VCG	UC-RAE in GSP	(Eff.) RAE in GSP
5	5	b₁	b₁	b₁	b₁
4	4	3.15	4	2.9	2.8
3	3	2.3	2⁺	1.8⁺	1.6⁺
2	2	1.6	2	1.6	1.6
1	1	1	1	1	1
Revenues	96	96	86	86	82

Notes. 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.

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 $\mathcal{R}^{EFF} = \{R^{EFF}(C)\}_{C \in \mathcal{C}}$ be such that, for each nontrivial 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 an RAE of the GSP auction in which 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.¹⁸

Example 6. Consider the environment of Examples 3 and 4 with $C = \{1, 3\}$. The efficiency-constrained RAE

is $\hat{b} = (\hat{b}_1, 2.8, 1.6^+, 1.6, 1)$, which results in revenues 82, which are lower than the RAE in the 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 or her 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 1 summarizes and compares the equilibria illustrated in our running examples.)

Summing up, because—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 because of the agency’s effect on prices. Obviously, if allocative inefficiencies were introduced, they might 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).

The unrestricted RAE of the GSP are difficult to characterize and (as shown with Example 5) allow perhaps too much freedom to the agency.¹⁹ The extra restrictions enabled by the Eff-RAE or UC-RAE thus may prove to be more fruitful from an applied perspective because they represent outcomes that are computationally easier to attain for an agency.

4.3. 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 this point. In particular, for certain coalition structures, our earlier results extend to the case with multiple agencies essentially unchanged:

the revenue losses are 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. But, for other coalition structures, equilibria in pure strategies do not exist, and hence, bidding cycles are likely to emerge.²⁰ Hence, although competition between agencies may indeed mitigate the agencies’ effects on the platforms’ revenues, it may also impair the working of the current mechanisms in a more fundamental way.

For simplicity, we consider the case with two agencies (the extension to more than two agencies is cumbersome but straightforward). We also assume that agencies break indifferences over bids in the same way that independents do. This implies that the highest bidder in any coalition bids as if the bidder were an independent. With the formal definitions given in Section A.3 of the appendix, the following result holds.

Theorem 4. (i) *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.* (ii) *If nontop members of different coalitions occupy adjacent positions in the raking of valuations, 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 because of coordinated bidding. If coalitions have bidders in adjacent positions (part (ii) of the theorem), further problems arise, such as nonexistence 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 1. Table 2 reports EOS’s 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 or her 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 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 (i) 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).²¹ If we split C as in the last column of Table 2, $C_1 = \{1, 4\}$ 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 (ii) of Theorem 4 shows that this phenomenon emerges whenever two coalitions have nontop members that occupy contiguous positions in the ordering of valuations. It is interesting to note that the behavior behind this phenomenon is nearly identical to that explained by Edelman and Ostrovsky (2007) in their characterization of the original 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.

From an empirical perspective, these results suggest further directions of research because they imply that

Table 2. 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	b_1
4	3.15	2.75	3.05	b_2
3	2.3	1.5	2.1	b_3
2	1.6	0 ⁺	1.2	b_4
1	1	0	0	b_5
Revenues	96	60	88	—

Note. Bids by coalition members are in bold.

bidding cycles are more likely to be observed as agency competition spreads in this market and especially so for certain configurations of agency membership.

5. Discussion: Related Literature and Extensions

5.1. RAE and EBA

Our notion of RAE is closely related to the EBAs of Ray and Vohra (1997, RV hereafter). Given a certain coalition structure, RV postulate that binding agreements are possible within a coalition. The objective is to endogenize the collection of agreements such that no subcoalition has an incentive to break the agreement and separate from the original coalition. Moreover, when considering such deviations, the subcoalition is farsighted in the sense that it does not take the behavior of the other coalitions as given, nor does it assume that the remaining members of the coalition will band together. Instead, it tries to predict the coalition structure and the agreements that would ultimately arise as a result of its deviation. In equilibrium, such predictions are required to be correct. Because of the farsightedness assumption, RV's equilibrium is defined recursively, as is our RAE.²²

RV's and our approaches share the same fundamental philosophy. Like RV, we also maintain that binding agreements are only possible within the coalition, but the interaction between the agency and the independents as well as among the independents is fully noncooperative. As in RV, the agency in our model is a proposer of a binding agreement, subject to certain stability constraints, which crucially incorporate RV's farsightedness assumption. 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 noncooperative interaction is based on Nash equilibrium), our definition also allows for refinements. As already explained, this is crucial here, especially for the analysis of GSP auction.

5.2. Bidding Rings and Partial Cartels

The literature on bidding rings in auctions (e.g., Graham and Marshall 1987, Mailath and Zemsky 1991, McAfee and McMillan 1992, Hendricks et al. 2008) also addresses related phenomena but from a very different perspective. In particular, the main focus of this literature is on whether members of the coalition may be incentivized to share their private information so as to implement collusion, a moot

point under EOS's complete information assumption. Maintaining EOS and Varian's complete information setting, we implicitly abstract away the information-extraction problem within the coalition. Furthermore, we don't allow transfers between members of the coalition.²³ Other mechanisms for collusion have been considered, for instance, by Harrington and Skrzypacz (2007, 2011). More importantly, a key feature of our setting is the copresence of coordinated and independent bidding. Combining cooperative and noncooperative interaction is a well-known challenge in this literature, which either considered mechanisms in which noncooperative behavior is straightforward (e.g., second price auctions with private values as in Mailath and Zemsky (1991)) or has assumed that the coalition includes all bidders (as in the first price auctions of McAfee and McMillan (1992) and Hendricks et al. (2008) or in the dynamic auctions of Ortner and Chassang (2018) in a different setting). The notion of RAE enables us to combine cooperative and noncooperative interaction in general mechanisms even if noncooperative behavior is complex. The results perhaps suggest that the general concept of RAE (or other concepts based on RV's approach) may provide a valuable methodological contribution from a broader theoretical perspective to overcome some of the difficulties involved with modeling partial cartels in auctions.

5.3. Alternative Competitive Benchmarks

The prior literature has shown that, without using the spitemove refinement concept of EOS, equilibria in the GSP auction can be worse in terms of both revenue and efficiency than in the VCG auction. Varian (2007) presents the EOS equilibrium as a lower bound of a class of NEs that he refers to as "symmetric" NEs. He shows that the lower bound on revenues among all NEs is generally less than the revenue bound for the symmetric NEs (i.e., the EOS revenues). Borgers et al. (2013) go a step further and show that inefficient NEs typically exist. Given these negative results on the GSP outside the EOS equilibrium, one might wonder whether the poor performance of the GSP auction is due to collusive agency bidding or simply because the equilibrium concept changed somewhat and the spite refinement has no bite anymore.

The response to this concern is that collusive agency bidding is the main driver of the poor performance of the GSP. First, the behaviors associated with the RAE equilibrium in the GSP typically lead to coalition bids that are below what could be sustained in the revenue-minimizing equilibrium of the corresponding competitive game. This is immediately clear when considering, for instance, the situation of coalition bidders occupying adjacent positions: their bids are typically so low that the lowest among these coalition

members, if it were to act as an independent, would find it individually profitable to raise its bid and jump to a higher position. The associated revenue loss, therefore, is directly a result of agency bidding and not merely to the possibility of low revenues among the Nash equilibria other than EOS's particular refinement. Second, as we showed in Lemma 1, our notion of RAE maintains the same individual-level underpinning of EOS refinement and, in particular, the features of EOS concepts that make the GSP work well. The EOS allocation and payments are embedded in our model as the end point of the recursion that defines the outside options of the coalition members. From this viewpoint, the fact that the VCG outperforms the GSP under the same conditions that—absent agency coordination—make the GSP perform better than it would if other refinements were considered, strengthens the result on the GSP's fragility. Obviously, our characterizations do exploit the specific properties of our EOS-based refinement. Results based on alternative competitive benchmarks would require altogether different proofs.²⁴ This, we think, would be an interesting enterprise for future research but seems beyond the scope of the present paper. Nonetheless, as we explain in Section 3, the general notion of RAE lends itself to this kind of exercise because it provides a tool to study the effect of agency bidding using different baseline refinements as plug-in (see also endnote 16).

5.4. Endogenous Participation and Alternative Approaches

An obvious extension 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. Once again, however, the structure of the GSP auction raises nontrivial 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.²⁵ More empirical work is needed on this subject.

Independent of these modeling choices, however, the cost of joining the agency would ultimately have to be traded off against the benefit, which, in turn, presumes

solving for the equilibrium for a *given* coalition structure. Our work can, thus, be seen as a necessary first step in developing a full-blown model of agency formation. Exploring different specifications of such costs and empirically assessing their relative merits is, thus, an important direction for future research in this area.

Our formulation of the agency problem is also related to the literature on mediators in games, introduced by Monderer and Tennenholtz (2009) for complete information and extended by Ashlagi et al. (2009) to incomplete information with an application to position auctions. Within this context, the issue of participation has been discussed, for instance, by Kalai (2010) and Roth and Shorrer (2018). Finally, a different approach to agency bidding in the GSP auction is offered in Lorenzon (2018), which considers a complete information setting in which the agency consists in the grand coalition of bidders.

5.5. Quality Scores

In the variant of the GSP auction run by Google or Microsoft Bing (but not, for instance, by Taobao), "quality scores" concur in determining the assignment of advertisers to slots and prices: advertisers are ranked by the product of their bid and quality score and pay a price equal to the minimum bid consistent with keeping that position.²⁶ EOS and Varian (2007) showed how to extend their equilibrium characterization when quality scores are introduced and assuming that they are common knowledge.²⁷ Quality scores could be introduced in our model of collusive bidding in a way similar to EOS's and Varian's (2007), delivering analogous characterizations of the preceding results. Such an extension is pursued in Decarolis et al. 2018, who develop a criterion to detect various forms of collusion based on the variables that are typically contained in the data sets available to the auction platforms.

6. Conclusions

This is the first study, to our knowledge, to focus on the impact of coordinated bidding through intermediaries in search auctions. It, therefore, contributes to the growing need of understanding both how firms operate on the platforms on which online ad space is sold and how these platforms should be designed. Our results uncover a striking fragility of the GSP auction to bid coordination.²⁸ Aside from its theoretical interest, this is a first-order finding because 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 noncooperative 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. Although beyond the scope of this paper, our analysis suggests some possible guidelines for research in this area. For instance, our analysis of the GSP auction with undistinguishable 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. The design of auction formats more robust to collusion is a challenging task, and the new formats that might emerge could have profound implications on the profitability of one of today's most important industries.

Finally, our results have implications for competition policy. For competition authorities, ad auctions might be worth investigations for potential violations of the antitrust laws, especially in those jurisdictions in which price coordination is a violation per se regardless of any welfare implication. In fact, the multiple activities that DMAs undertake beyond bid coordination make a priori ambiguous their overall effects on consumers' welfare. Furthermore, an additional complication is that, in the context of the ad auctions, bid coordination by a DMA simply requires it to use bid algorithms that optimize joint profits of its clients without the need of any explicit communication. This poses a challenge for those authorities operating under jurisdictions that only sanction explicit (as opposed to tacit) collusion. In this respect, our analysis offers a clear application of the novel problems that algorithmic pricing poses for the enforcement of competition policy.

Acknowledgments

The authors are grateful for the comments received from Susan Athey, Jean-Pierre Benoit, Yeon-Koo Che, Kerry Gabrielson, Joseph Harrington, Ken Hendricks, Jon Levin, Massimo Motta, Marco Ottaviani, Rob Porter, Marc Rysman, Matthew Shum, Andy Skrzypacz, and Steve Tadelis as well as from participants at the seminars at Berkeley-Haas Business School, Boston University, Columbia University, Center for Research in Economics and Statistics-Paris, European Commission Directorate General Competition, Facebook Research Laboratory, Harvard-Massachusetts Institute of Technology joint IO seminar,

London Business School, Microsoft Research Laboratory, Princeton University, Stanford University, University of British Columbia, University of California Davis, University of Chicago, University of Bologna, University of Toronto, and University of Wisconsin-Madison.

Appendix.

A.1. Technical Details

As discussed in Section 2, 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: $[0, b_{-i}^S)$, $[b_{-i}^S, b_{-i}^{S-1})$, ..., $[b_{-i}^1, \infty)$. 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 bidder i obtains slot $t = 1, \dots, S + 1$ 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 2 can be seen as a correspondence $\pi_i : \mathbb{R}_+^{n-1} \rightrightarrows \{1, \dots, S + 1\}$ such that, for each $b_{-i} \in \mathbb{R}_+^{n-1}$, $\pi_i(b_{-i}) = \arg \max_{t=1, \dots, S+1} (v_i - b_{-i}^t) x^t$.²⁹ To allow for the possibility of ties in the bid 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 endnote 8, 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

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

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

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

Because $\mathcal{S}(b_{-i})$ is always nonempty and finite, the best responses $BR_i : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$ defined in Section 2 are well defined, and so is $BR_i^* : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$ in (2). 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 $\pi_i(\hat{b}_{-i}) = s$ if $s \leq S$ and $\pi_i(\hat{b}_{-i}) = S + 1$ if $s > S$. Hence, $\hat{b}_{-i}^{\pi_i(\hat{b}_{-i})} = \hat{b}^{s+1}$ 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:

$$\begin{aligned} \text{By the optimality of } \hat{b}_i: & (v_i - \hat{b}^{\rho(i)+1}) x^{\rho(i)} \\ & \geq (v_i - \hat{b}^{\rho(i)+2}) x^{\rho(i)+1}; \end{aligned} \quad (\text{A.1})$$

By the condition in (2) for j :

$$(v_j - \hat{b}^{\rho(i)+2}) x^{\rho(i)+1} = (v_j - \hat{b}^{\rho(i)+1}) x^{\rho(i)}. \quad (\text{A.2})$$

Rearranging, we obtain

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

which implies that $v_i > v_j$ (because, 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 $\rho(i) > S$, Equation (2) requires that $0 = (v_i - \hat{b}_i)x^S$; 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 $i \geq S + 1$. The equilibrium bid, $b_i = v_i - \frac{x^i}{x^{i-1}}(v_i - b_{i+1})$, then follows immediately, applying Equation (2) for all $i = 2, \dots, S$ 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 the equilibrium bid coincides with EOS's lowest envy free equilibrium (EOS, theorem 2) and with Varian's lower-bound symmetric Nash equilibrium (Varian 2007, equation 9). \square

A.2.2. Proof of Theorem 1. We prove the statement by induction on the size of the coalition. The induction basis is the noncollusive 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 its position. Then, its payoff is equal to

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

It is useful to introduce notation to rank independents among themselves based on their valuation. Let $v_{\Lambda C} = (v_j)_{j \in \Lambda C}$, and let $v_{\Lambda C}(k) = v_{\Lambda C}^{|\Lambda C|+1-k}$ denote the valuation of the k th lowest value independent: for $k = 1$, $v_{\Lambda C}(1) = v_{\Lambda C}^{|\Lambda C|}$ is the lowest valuation among the independents, $v_{\Lambda C}(2) = v_{\Lambda C}^{|\Lambda C|-1}$ 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 $t = r + 1, \dots, S + 1$ 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 $t = r + 1, \dots, S + 1$ are placed by the $S + 1 - r$ lowest-valued independents. Hence,

$$u_i = v_i x^r - \sum_{t=r+1}^{S+1} v_{\Lambda C}(S + 2 - t) (x^{t-1} - x^t). \quad (\text{A.3})$$

Let us consider the function $\tilde{u}_i(k)$ of i 's payoff as a function of the position k that i occupies, given that i 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 $u_i^* < u_i^{C \setminus \{i\}}$ (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 $j \in \Lambda C$, $j < i$ if and only if $\rho(j) < r$. We proceed by contradiction: suppose that there exist $j \in \Lambda C$ such that either $j < i$ and $\rho(j) > r$ or $j > i$ and $\rho(j) < r$. Consider the first case: because independents are ranked efficiently among themselves, for any $j, l \in \Lambda C$, $l < j$ if and only if $\rho(l) < \rho(j)$. It follows that if there exists $j \in \Lambda C : j < i$ and $\rho(j) > r$, such j can be chosen so that $j = r + 1$; that is, j occupies the position immediately following i 's. We next show that, in this case, i 's payoff would increase if i dropped one position down. To see this, notice that

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

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

$$\text{sign}(\tilde{u}_i(r + 1) - \tilde{u}_i(r)) = \text{sign}(v_{\Lambda C}(S + 1 - r) - v_i).$$

Under the absurd hypothesis, $v_{\Lambda C}(S + 1 - r) > v_i$; hence, u_i increases, dropping one position down. A similar argument shows that in the second case of the absurd hypothesis, that is, if there exists $j \in \Lambda C : j > i$ 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

$$\begin{aligned} u_i(r) - u_i(r - 1) &= (b^{r-1} - v_i) (x^{r-1} - x^r) \\ &\leq (v_{\Lambda C}(S + 2 - r) - v_i) (x^{r-1} - x^r), \end{aligned} \quad (\text{A.4})$$

which holds because all bids for positions from $\rho(j)$ down are no higher than $b_j = v_{\Lambda C}(S + 2 - r)$. The final expression is negative under the absurd hypothesis.

We thus have proved that, in equilibrium, for all $j \in \Lambda C$, $j < i$ 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 it. 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 Equation (A.3), we have that

$$u_i^* = v_i x^{i+|\mathcal{J}|} - \sum_{t=i+|\mathcal{J}|+1}^{S+1} v_{\Lambda C}(S + 2 - t) (x^{t-1} - x^t). \quad (\text{A.5})$$

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

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

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 k bids

$\bar{b}_k = v_k < v_{k-1}$; if k is the highest-value member of the coalition, because $\bar{b}_k \in (b_{k+1}^+, v_{k-1})$; otherwise, $\bar{b}_k = b_{k+1}^+ < v_{k-1}$. We also show that $\bar{b}_k \leq v_{\Gamma_C}(S+2-k)$ for all k . To this end, observe that all $k \geq \max\{\mathcal{F}\}$ are independents (both before and after i drops out) so that, for all $k \geq \max\{\mathcal{F}\}$, $\bar{b}_k = v_k = v_{\Gamma_C}(S+2-k)$: these are the lowest bidding and the lowest-value bidders; hence, also the lowest independents. For $k < \max\{\mathcal{F}\}$, at least one of the $S+2-k$ elements of the set $\{k, k+1, \dots, S+1\}$ 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, $v_{\Gamma_C}(S+2-k) \geq v_{k-1}$, which, in turn, implies $v_{\Gamma_C}(S+2-k) > \bar{b}_k$. Overall, we have that $\bar{b}_k < v_{k-1}$ and $\bar{b}_k \leq v_{\Gamma_C}(S+2-k)$ for all $k \in I$. Using the first inequality for $k \leq i + |\mathcal{F}|$ and the second inequality otherwise, we see that, if $\mathcal{F} \neq \emptyset$,

$$\begin{aligned} u_i^{C \setminus \{i\}} &= v_i x^i - \sum_{k=i+1}^{i+|\mathcal{F}|} \bar{b}_k (x^{k-1} - x^k) - \sum_{k=i+|\mathcal{F}|+1}^{S+1} \bar{b}_k (x^{k-1} - x^k) \\ &> v_i x^i - \sum_{k=i+1}^{i+|\mathcal{F}|} v_{k-1} (x^{k-1} - x^k) \\ &\quad - \sum_{k=i+|\mathcal{F}|+1}^{S+1} v_{\Gamma_C}(S+2-k) (x^{k-1} - x^k). \end{aligned} \quad (\text{A.6})$$

Combining (A.5) and (A.6), we get

$$\begin{aligned} u_i^{C \setminus \{i\}} - u_i^* &> v_i (x^i - x^{i+|\mathcal{F}|}) - \sum_{k=i+1}^{i+|\mathcal{F}|} v_{k-1} (x^{k-1} - x^k) \\ &\geq v_i (x^i - x^{i+|\mathcal{F}|}) - v_i (x^i - x^{i+|\mathcal{F}|}) = 0, \end{aligned}$$

where the latter inequality follows because $v_{k-1} \leq v_i$ for all $k \geq i+1$. Hence, whenever $\mathcal{F} \neq \emptyset$, we obtain $u_i < u_i^{C \setminus \{i\}}$; that is, the recursive stability condition (S.2) is violated for bidder i . $\mathcal{F} = \emptyset$, 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 $\mathcal{F} = \emptyset$, $u_i^* = u_i^{C \setminus \{i\}}$ (by Equations (A.5) and (A.6)); hence, in equilibrium, $u_i = u_i^*$ and $i = \rho(i)$:

$$u_i = v_i x^i - \sum_{k=i+1}^{S+1} v_k (x^{k-1} - x^k) = u_i^{C \setminus \{i\}}. \quad (\text{A.7})$$

Furthermore, because 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 sets 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 $k \geq i$. 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 $(b_j)_{j \in C \setminus \{i\}}$ in order to maximize $\sum_{j \in C \setminus \{i\}} u_j$, subject to the constraint that $b_j > b_i$ for all $j \in C \setminus \{i\}$. 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 its 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 $j \in C \setminus \{i\}$ drops out, i still remains a nontop 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 $k \geq i$ (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 $C \setminus \{i\}$ and independents $\{j \in I \setminus C : j > i\}$ with slots x^1, \dots, x^{i-1} . This game has coalition size $C-1$, so the solution follows by the inductive hypothesis. \square

A.2.3. Proof of Theorem 2. Because the UC restrictions imply the stability restriction (S.1), the agency's problem in the GSP auction with the feigned values restriction reduces to

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

where the equilibrium conjectures β^* are such that

$$\forall b_C, \beta^*(b_C) \in \{b_{-C}^* \in \mathbb{R}_+^{n-|C|} : \forall i \in I \setminus C, b_i^* \in BR_i^*(b_C, b_{-i,-C}^*)\}.$$

Let \sim be an equivalence relation on \mathbb{R}_+^n such that $v \sim v'$ (respectively, $b \sim b'$) if and only if v and v' only differ in the highest valuation (highest bid) but not in the identity of the highest valuation individual (bidder).³⁰ For any $v \in \mathbb{R}_+^n$, let $[v]$ ($[b]$) denote the equivalence class of v (b) under this equivalence relation, and let \mathbb{V}^\sim (\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 $E^\sim : \mathbb{V}^\sim \rightarrow E^*(\mathbb{V}^\sim)$ 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:

$$\text{For all } v \in \mathbb{R}_+^n \text{ and } i \in I, u_i^{v'}(v) = u_i^{v'}(E^*(v)). \quad (\text{A.8})$$

Because 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 (A.8), this implies that $u_i^{v'}(b_C, b_{-C}^*) = u_i^{v'}(b'_C, b'_{-C})$ so that also $u_i^{v'}(b_C, b_{-C}^*) = u_i^{v'}(v'_C, v_{-C})$. As a result, we can now easily recast the coalition's problem as one of choosing v'_C (the coalition's feigned valuations):

$$\begin{aligned} \max_{v'_C} u_C^{v'}(v'_C, v_{-C}) \\ \text{s.t. (S.2)} \forall i \in C, u_i^{v'}(v'_C, v_{-C}) \geq \bar{u}_i^{C \setminus \{i\}}. \end{aligned}$$

(Notice that the restriction (R) and the restriction that $\beta^*(b_C)$ is always in the set BR_{-C}^* are both built in this formulation of the problem.) In the following, we let \bar{u}_i^C denote bidder i 's

payoff when the coalition is C in the GSP game being studied, and $\bar{u}_i^{C,V}$ denotes the same object in the corresponding VCG game. With this in mind, note that $\bar{u}_i^C = \bar{u}_i^{C,V}$ for all i when $|C| = 1$, and the recursion defining \bar{u}_i^C is identical to that defining $\bar{u}_i^{C,V}$. 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 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. \square

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. \square

Lemma A.1. Fix $C \subset I$, and let \mathcal{K} be a finite index set. Let $\{b^{(k)}\}_{k \in \mathcal{K}}$ be a collection of bid profiles such that for each $k \in \mathcal{K}$, $b_{-C}^{(k)} \in BR_{-C}^*(b_C^{(k)})$ and $\rho(i; b^{(k)}) = i$ for each $i \in I$. Define $\mathcal{L}(\{b^{(k)}\}_{k \in \mathcal{K}}) \equiv \hat{b} \in \mathbb{R}_+^n$ 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 $\hat{b}_{\bar{c}(i)} \neq b_{\bar{c}(i)}^{(k)}$, (iii) $u_C(\hat{b}) \geq u_C(b^{(k)})$ for all $k \in \mathcal{K}$ with strict inequality whenever $\exists i \in C \setminus \min C$ such that $b_i^{(k)} \neq \hat{b}_i$, and (iv) $\hat{b}_{-C} \in BR_{-C}^*(\hat{b}_C)$.

Proof of Lemma A.1. 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. $i \neq S+1$,

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

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

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

and 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, 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 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 $k \in \mathcal{K}$ with strict inequality whenever $\exists i \in C \setminus \min C$ 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 $i \in C \setminus \min C$, and let $j = \max\{k \in C \mid k < i\}$ 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), $u_j(b^{(k)}) < u_j(\hat{b})$. Because $u_j(b^{(k)}) \leq u_j(\hat{b})$ 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, that is, 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, 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 $b^{(k)}$, that is, $(v_i - b_j^{(k)})x^j \leq (v_i - b_{i+1}^{(k)})x^i$. Finally, by point 1, $b_{i+1}^{(k)} \geq \hat{b}_{i+1}$ 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, and 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 $C \setminus \{i\}$ (these exist under the inductive hypothesis). Let $b^{(0)} = b^{UC}$. Let $\hat{b} = \mathcal{L}(\{b^{(i)}\}_{i \in C \cup \{0\}})$, where \mathcal{L} is as defined in Lemma A.1. Now, by results (i) and (iv) of Lemma A.1, 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 A.1, $u_i(\hat{b}) \geq u_i(b^{(k)})$ for each i . Moreover, by construction, $u_i(b^{(k)}) = \bar{u}_i^{C \setminus \{i\}}$ 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 nonempty, 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 A.1 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 $i \in C \setminus \min C$ such that $b_i^* > \hat{b}_i$. Let $b' = \mathcal{L}(\{b^*, \hat{b}\})$. By part (i) of Lemma A.1, 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 $i \in C \setminus \min C$.

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, $\hat{b}_i \leq b_i^{UC}$ for all $i \in I$, we, thus, have $b_i^* \leq b_i^{UC}$ 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 $b_{\min C}^{UC}$ is the maximum bid that the top coalition bidder can place in a UC-RAE, it is equal to (cf. Theorem 2)

$$b_{\min C}^{UC} = v_{\min C-1} - \frac{x^{\min C}}{x^{\min C-1}} (v_{\min C-1} - b_{\min C+1}^{UC}).$$

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

$$\begin{aligned} U_0 &= (v_{\min C-1} - b_{\min C}^*)x^{\min C-1} < (v_{\min C-1} - b_{\min C}^{UC})x^{\min C-1} \\ &= (v_{\min C-1} - b_{\min C+1}^{UC})x^{\min C}, \end{aligned}$$

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

Dropping one position down, this independent would obtain

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

where the first inequality follows because $b_i^* \leq b_i^{UC}$ for all $i > \min C$ as established. Thus, this independent has a profitable deviation: a contradiction. We conclude that $b_{\min C}^* \leq b_{\min C}^{UC}$. But then, by the independents' recursion, we also have $b_i^* \leq b_i^{UC}$ for all $i \leq \min C$. Because we already knew that the $b_i^* \leq b_i^{UC}$ 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 A.1, \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 A.1, 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 A.1 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. \square

A.3. Multiple Agencies

A.3.1. Formal Definition. We consider the case with two DMAs, 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 DMA, the definition of RAE with multiple agencies is recursive with the outside option of coalition member $i \in C_1$ being defined as its 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 DMA, but with initial condition set by the RAE in which C_{-g} is the only coalition.

Let $G(v) = (A_i, u_i)_{i=1, \dots, n}$ 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 $i \in I \setminus C$ simultaneously choose bids $b_i \in A_i$. We let $b_{-C} := (b_j)_{j \in I \setminus C}$ and $A_{-C} := \times_{j \in I \setminus C} A_j$. Given profile b or b_{-C} , we let $b_{-i, -C} := (b_j)_{j \in I \setminus C; j \neq i}$. As before, each DMA 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.

A.3.1.1. Stability 1 (Stability with Respect to Independents).

For any $i \in I \setminus C$, let $BR_i^* : A_{-i} \rightarrow A_i$, $BR_{-C}^* : A_C \rightarrow A_{-C}$, 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

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

A.3.1.2. Stability 2: ((Recursive) Stability with Respect to Coalition Members).

Let B^* be defined as in the single-agency case. Letting $E^{\mathcal{R}}(C_1, C_2)$ denote the set of 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 $E^{\mathcal{R}}(C'_g, C_{-g}) = E^{\mathcal{R}}(C_{-g})$ if $|C'_g| = 1$ (i.e., 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 $C'_g \subseteq C_g \setminus \{i\}$, 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 $u_i \geq \bar{u}_i^{C_{-g} \setminus \{i\}, C_g}$. Finally, we define the set of rational conjectures about the opponent coalition as $B_g^* = \{\beta_g \in (A_{C_{-g}})^{\bar{S}_{C_g}} : \beta_g(b_{C_g}) \in BR_{C_g}^C(b_{C_g}) \text{ for all } b_{C_g} \in \bar{S}_{C_g}\}$, where $\bar{S}_{C_g} = \{b_{C_g} \in S_{C_g} : BR_{C_g}^C(b_{C_g}) \neq \emptyset\}$, and

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

Definition A.1. A 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 $i \in I \setminus C$, $b_i^* \in BR_i^*(b_{-i}^*)$.

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

$$\begin{aligned} b_{C_g}^* &\in \arg \max_{b_{C_g}} u_{C_g}(b_{C_g}, \beta_g^*(b_{C_g}), \beta^*(b_{C_g}, \beta_g^*(b_{C_g}))) \\ \text{s.t. (R)} & ((b_{C_g}, \beta_g^*(b_{C_g}), \beta^*(b_{C_g}, \beta_g^*(b_{C_g}))) \in R_C \\ \text{(S.1)} & (b_{C_g}, \beta_g^*(b_{C_g}), \beta^*(b_{C_g}, \beta_g^*(b_{C_g}))) \in S_C \\ \text{(S.2)} & \text{for all } i \in C_g, u_i(b_{C_g}, \beta_g^*(b_{C_g}), \\ & \cdot \beta^*(b_{C_g}, \beta_g^*(b_{C_g}))) \geq \bar{u}_i^{C_g \setminus \{i\}, C_{-g}}. \end{aligned}$$

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

$$\begin{aligned} E^{\mathcal{R}}(C_1, C_2) \\ = \{b^* \in A : \exists \beta^*, \beta_1^*, \beta_2^* \text{ s.t. } (b^*, \beta^*, \beta_1^*, \beta_2^*) \text{ is a RAE}\}. \end{aligned} \quad (\text{A.9})$$

Note that this definition does not uniquely pin down 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

$$\hat{b}_i^V = \begin{cases} v_i, & \text{if } i \in (I \setminus C) \cup \min C_1 \cup \min C_2, \\ \left(\hat{b}_{i+1}^V \right)^+, & \text{if } i \in C \setminus \{\min C_1 \cup \min C_2\} \text{ and } i \leq S, \end{cases} \quad (\text{A.10})$$

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 ,

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

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

2. If, in the overall value ranking, a nontop member of one coalition is directly above a nontop 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.

We prove the results for the VCG (statement 1(a) and the VCG part of statement 2). 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 adjacencies, but not when there are adjacencies involving nontop bidders.³¹

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

Observation A.1. The best-response function of any coalition requires that each nontop 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 A.1. Suppose that $b_i \neq (b^{\rho(i)})^+$, 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. □

Observation A.2. The best-response function of any coalition requires that no member of the coalition (top or nontop) 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 A.2. Suppose that $v_i < b^{\rho(i)+1}$ and consider the deviation in which C_g lowers 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. □

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 ($|C| = 2$, that is, $|C_1| = |C_2| = 1$), 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 because of 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 A.1, $b_i = v_{i+1}^+$.

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 A.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 that 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 A.1, this means that $b_i < v_j < v_i$. But consider the deviation in which 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 A.2, this deviation causes the other coalition to move bidder j (and any other members 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, therefore, bid its true value (by the assumed equilibrium selection).

We now can repeat these 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 nontop 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 upward. 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 upward. 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 downward. 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 downward. This completes the proof of the theorem.

Finally, we show that there is no equilibrium if there are any cases in which nontop 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 A.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, $b_{i+1} \notin BR_k(b_C)$, that is, we are not in an RAE.

Endnotes

¹ Mansour et al. (2012) first pointed at the potential risk of collusive bidding that intermediaries posed for online ad auctions. That paper focused on the ad exchanges used for display ads. Balseiro and Candogan (2017) is an important attempt to look at optimal contracts for intermediaries, showing that, under such contracts, intermediaries can bolster revenues for budget-constrained advertisers while also increasing the overall market efficiency. Still in the context of display ad auctions, Allouah and Besbes (2017) provide conditions under which bid coordination by a common intermediary can either improve or reduce the advertisers' surplus. McAfee (2011) studies how intermediaries help solve problems of limited information in ad exchanges. The importance of information and learning in display auction is also stressed in Despotakis et al. (2017).

² In 2016, the total revenues of these seven agency networks amounted to one third of those of Alphabet.

³ The use of the word "collusion" in this paper is unrelated to any assessment of the legal implications of agencies or advertisers behavior under the competition laws of the United States or other countries.

⁴ Google, for instance, reports passing from a positive growth rate in its average cost per click of about 4% 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%. Source: 10-k filings of Alphabet, Inc.

⁵ The problem of partial cartels is acknowledged as a major difficulty in the literature (e.g., Hendricks et al. 2008). We discuss this point and the connection with that literature in Section 5.

⁶ This approach, which involves both equilibrium and recursive stability restrictions, is closely related to the equilibrium binding agreements of Ray and Vohra (1997), further discussed in Section 5.

⁷ 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. Also, similar to EOS, our baseline model abstracts from quality scores, which, in practice, are often used to adjust advertisers' bids in determining their position and payments. Athey and Nekipelov (2014) introduced uncertainty over quality scores in a model with competitive bids. We discuss quality scores in Section 5.

⁸ 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.

⁹ Allowing ties in individuals' bids or nongeneric indifferences complicates the notation without affecting the results and the main insights. See Section A.1 of the appendix for details on this.

¹⁰ A Nash equilibrium $(b_i)_{i \in I}$ is locally envy-free if $x^{\rho(i)}(v_i - b^{\rho(i)+1}) \geq x^{\rho(i)-1}(v_i - b^{\rho(i)})$ for every i . EOS refinement is the lowest-revenue Nash equilibrium that satisfies this condition. This refinement is especially important because it conforms with the search engines' tutorials on how to bid in these auctions. 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>.

¹¹ 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.

¹² The strength of constraint (5) clearly depends on the underlying game G and on the particular correspondence BR_C^* . This restriction is conceptually important and needed to develop a general framework for arbitrary mechanisms, but it plays no role in our results because (5) is either vacuous (Theorem 1) or a redundant constraint (Theorems 2 and 3). In particular, under the VCG mechanism (Theorem 1), we have $\mathcal{F}^C = A_C$, thus making constraint (5) vacuous. As for the the GSA, in the two theorems we consider, it is always the case that the set of exogenous restrictions (i.e., $R^{LC}(C)$ for Theorem 2 and $R^{Eff}(C)$ for Theorem 3) are always a subset of \mathcal{F}^C , thereby making constraint (5) redundant in the agency's optimization problem for those theorems.

¹³ 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.

¹⁴ Without the tie-breaking rule embedded in ρ (endnote 8), 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} , that is, $b_i = b_{i+1} + \varepsilon$. All our results would hold in such a discrete model once the equilibrium bids in the theorems are interpreted as the limit of the equilibria in the discrete

model, letting $\varepsilon \rightarrow 0$ (the notation b_{i+1}^+ is thus reminiscent of this alternative interpretation as the right-hand limit $b_{i+1}^+ := \lim_{\varepsilon \rightarrow 0} (b_{i+1} + \varepsilon)$).

¹⁵ This property also ensures that $S_C = A_C$. Hence, constraint (S.1) in Definition 1 plays no role in the result.

¹⁶ This formulation of the UC constraint is consistent with our choice to use the EOS equilibrium as the competitive benchmark as it has become standard in the literature. The definition in (9), however, may easily accommodate alternative benchmarks too. As explained in Section 3, alternative models of competitive behavior could be accommodated in the definition of RAE by replacing (2) with the corresponding refinement of individual best responses. In that case, the set $E^*(v'_C, v_{-C})$ in (9) would consist of the fixed points of such individual best responses, that is, the Nash equilibria of the GSP taken as a benchmark of competitive bidding. Hence, whatever refinement of Nash equilibria is taken as a model of competitive behavior in the GSP—and, hence, embedded in the definition of RAE—the set R^{UC} denotes the set of bids profile that cannot be distinguished from that competitive benchmark.

¹⁷ The reason is similar to that discussed for Theorem 1, only here it is more complicated because, in the GSP auction, the bids of the agency alter the bids placed by the independents.

¹⁸ Because the UC-RAE induce efficient allocations, it may seem that Theorem 3 follows immediately from the efficiency constraint being weaker than the UC restriction. This intuition is incorrect for two reasons. First, the UC constraint requires the existence of feigned valuations, which can rationalize the observed bid profile but does not require that they preserve the ranking of the true valuations. Second, when the exogenous restrictions $\mathcal{R} = (R_C)_{C \in \mathcal{C}}$ 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.

¹⁹ Short of a characterization, one could consider whether there are clear revenue rankings for the GSP's unrestricted RAE. Meaningful revenue comparisons, however, require normalizing the criterion to break the agency's indifference over the highest bid, which does not affect the agency's payoff (beyond the position it ensures), but does affect the revenues. This indeterminacy does not create problems for the preceding results because it is preserved uniformly across the mechanisms and restrictions we consider, but when all restrictions are lifted, the mere breaking of indifference may impact revenues asymmetrically under different mechanisms and restrictions. Standard tie-breaking criteria, however, do enable natural revenue comparisons. For instance, if one applied the same logic used to refine the independents' indifference (i.e., the locally envy-free criterion implicit in (2)), or if one considered the lowest-revenue selections among the bids that maximize the agency's payoff, then it could be shown that the GSP's revenues in an inefficient RAE are never higher (and typically lower) than in the Eff-RAE.

²⁰ As already mentioned, bidding cycles are indeed considered to be one of the main reasons why the GFP auction, which was adopted in the early days of this industry, was eventually abandoned in favor of the GSP. (For a discussion of bidding cycles in the Overture's GFP, see Edelman and Ostrovsky 2007; Ottaviani 2003 provides an early assessment of the transition from the GFP to the GSP.)

²¹ 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 its bid, its 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).

²² The idea of farsightedness in coalition formation is further explored in Ray and Vohra (1997, 2014).

²³ Allowing transfers would relax constraint (S.2) in the definition of RAE and affect our results (for instance, it may induce inefficiencies even in the VCG mechanism; cf. Example 2). That different advertisers make side payments to each other seems implausible in this market. If indirect transfers could be implemented through dynamic effects (e.g., swapping bids for some of its members) or across different keywords, distinct strategic issues might arise, which would best be studied considering a richer model.

²⁴ Varian's (2007) upper bound symmetric NE has also a recursive structure, and hence it could be plugged in as an alternative to the EOS in our analysis. However, although Varian (2007) offers insights on the type of individual logic that might support behaviors leading to this upper bound, this logic is less compelling than the one upon which the spite move of EOS is based. The upper bound would be reached if everyone thinks defensively of squeezing the profit of the player placed right above but only up to the point at which the player above does not prefer to jump down one position. Other refinements need not have that structure and, hence, may entail significantly different strategy of proofs. However, aside from the experimental results in Che et al. (2017), there is still limited understanding on the behaviors in the GSP auction so that our preference is for the modeling approach presented in the text.

²⁵ Moreover, costs need not be symmetric, and hence, it may be that an advertiser is willing to join the coalition but current members are better off without it. Whereas the decision to abandon an agency is unilateral, the decision to join it is not, raising further modeling questions.

²⁶ Although an in-depth discussion of potential applications to data are beyond the scope of this paper, we refer to Ghose and Yang (2009) for an empirical model of the Google or Microsoft-Bing type of search auctions and to Hsieh et al. (2018) for the case of Taobao auctions. Neither these two papers nor others we are aware of in the literature develop methods for estimating an empirical model of the search auctions in the presence of collusive bidding. This type of analysis is presented in a paper related to ours (De Carolis et al. 2018).

²⁷ Competitive bidding with quality scores has also been studied by Athey and Nekipelov (2014), who introduced introducing uncertainty over quality scores.

²⁸ The empirical analysis in De Carolis et al. (2018) shows that even the small two-bidder coalitions frequently observed in the data can have large effects on revenues.

²⁹ This correspondence is always nonempty valued and multivalued 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 $\pi_i : \mathbb{R}_+^{n-1} \rightarrow \Pi$ as a function (if not, π_i should be thought of as a selection from the correspondence above).

³⁰ Formally, $v \sim v'$ if and only if the following two conditions hold: (a) $\arg \max_{i \in I} v_i = \arg \max_{i \in I} v'_i$; (b) $v_i = v'_i$ for all $i \neq \arg \max_{i \in I} v_i$.

³¹ Compared with 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 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 yields the fundamental difference between the cases with and without members from different coalitions that are adjacent in the value ranking.

References

- Allouah A, Besbes O (2017) Auctions in the online display advertising chain: A case for independent campaign management. Research Paper No. 17-60, Columbia Business School, New York.
- Ashlagi I, Monderer D, Tennenholtz M (2009) Mediators in position auctions. *Games Econom. Behav.* 67(1):2–21.
- Athey S, Nekipelov D (2014) A structural model of sponsored search advertising auctions. Working paper, Stanford University, Stanford, CA.
- Aumann RJ (1959) Acceptable points in general cooperative n-person games. Tucker A, Luce R, eds. *Contributions to the Theory of Games IV* (Princeton University Press, Princeton, NJ).
- Ausubel LM, Milgrom P (2006) The lovely but lonely Vickrey auction. *Combinatorial Auctions* (MIT Press, Cambridge, MA).
- Balseiro S, Candogan O (2017) Optimal contracts for intermediaries in online advertising. *Oper. Res.* 65(4):878–896.
- Bernheim BD, Whinston MD (1985) Common marketing agency as a device for facilitating collusion. *RAND J. Econom.* 16(2):269–281.
- Bernheim BD, Whinston MD (1986) Common agency. *Econometrica* 54(4):923–942.
- Bernheim BD, Peleg B, Whinston MD (1987) Coalition-proof Nash equilibria I. Concepts. *J. Econom. Theory* 17:1–12.
- Borgers T, Cox I, Pesendorfer M, Petricek V (2013) Equilibrium bids in sponsored search auctions: Theory and evidence. *Amer. Econom. J. Microeconomics* 5(4):163–187.
- Che Y-K, Choi S, Kim J (2017) An experimental study of sponsored-search auctions. *Games Econom. Behav.* 102:20–43.
- Decarolis F, Goldmanis M, Penta A (2018) Common agency and coordinated bids in sponsored search auctions. Working paper, Bocconi University, Milan, Italy.
- Despotakis S, Hafalir I, Ravi R, Sayed A (2017) Expertise in online markets. *Management Sci.* 63(11):3895–3910.
- Edelman B, Ostrovsky M (2007) Strategic bidder behavior in sponsored search auctions. *Decision Support Systems* 43(1):192–198.
- Edelman B, Ostrovsky M, Schwarz M (2007) Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *Amer. Econom. Rev.* 97(1):242–259.
- Ghose A, Yang S (2009) An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* 55(10):1605–1622.
- Gomes R, Sweeney K (2014) Bayes Nash equilibria of the generalized second-price auction. *Games Econom. Behav.* 86:421–437.
- Graham DA, Marshall RC (1987) Collusive bidder behavior at single-object second-price and English auctions. *J. Political Econom.* 95(6):1217–39.
- Harrington JE, Skrzypacz A (2007) Collusion under monitoring of sales. *RAND J. Econom.* 38(2):314–331.
- Harrington JE, Skrzypacz A (2011) Private monitoring and communication in cartels: Explaining recent collusive practices. *Amer. Econom. Rev.* 101(6):2425–2449.
- Hendricks K, Porter R, Tan G (2008) Bidding rings and the winner's curse. *RAND J. Econom.* 39(4):1018–1041.
- Hsieh Y-W, Shum M, Yang S (2018) To score or not to score? Estimates of a sponsored search auction model. Research Paper No. 15-09, University of Southern California, Los Angeles.
- Kalai AT (2010) A commitment folk theorem. *Games Econom. Behav.* 69(1):127–137.
- Levin J, Skrzypacz A (2016) Properties of the combinatorial clock auction. *Amer. Econom. Rev.* 106(9):2528–2551.
- Lorenzon E (2018) Collusion with a rent seeking agency in sponsored search auctions. Working paper, Paris Dauphine University, Paris.
- Mailath GJ, Zemsky P (1991) Collusion in second price auctions with heterogeneous bidders. *Games Econom. Behav.* 3(4):467–486.
- Mansour Y, Muthukrishnan S, Nisan N (2012) Doubleclick ad exchange auction. Preprint, submitted April 2, <http://arxiv.org/abs/1204.0535>.
- McAfee R, McMillan J (1992) Bidding rings. *Amer. Econom. Rev.* 82(3):579–599.
- McAfee RP (2011) The design of advertising exchanges. *Rev. Indust. Organ.* 39(3):169–185.
- Metz C (2015) Facebook doesn't make as much money as it could—on purpose. *Wired* (September 21), <https://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose>.
- Monderer D, Tennenholtz M (2009) Strong mediated equilibrium. *Artificial Intelligence* 173(1):180–195.
- Ortner J, Chassang S (2018) Making corruption harder: Asymmetric information, collusion, and crime. *J. Political Econom.* 126(5):2108–2133.
- Ottaviani M (2003) Overture and Google: Internet pay-per-click (PPC) advertising auctions. Case Study-03-022, London Business School, London.
- PwC (2017) Global entertainment media outlook. Accessed September 13, 2018, <https://www.pwc.com/gx/en/industries/tmt/media/outlook.html>.
- Ray D, Vohra R (1997) Equilibrium binding agreements. *J. Econom. Theory* 73(1):30–78.
- Ray D, Vohra R (2014) The farsighted stable set. *Econometrica* 83(3):977–1011.
- Roth B, Shorrer RI (2018) Making it safe to use centralized marketplaces: Dominant individual rationality and applications to market design. Working paper, Harvard Business School, Boston.
- Varian H (2007) Position auctions. *Internat. J. Indust. Organ.* 25(6):1163–1178.