NET Institute*

www.NETinst.org

Working Paper #13-19

September 2013

Common Agency and Coordinated Bids in Sponsored Search Auctions

Francesco Decarolis
Boston University and Hariri Institute

Maris Goldmanis
University of London

Antonio Penta
University of Wisconsin-Madison

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

Common Agency and Coordinated Bids in Sponsored Search Auctions

Francesco Decarolis*, Maris Goldmanis[†], Antonio Penta[‡]

October 1, 2013

Abstract

As auctions are becoming the main mechanism for selling advertisement space on the web, marketing agencies specialized in bidding in online auctions are proliferating. We analyze theoretically how bidding delegation to a common marketing agency can undermine both revenues and efficiency of the generalized second price auction, the format used by Google and Microsoft-Yahoo!. Our characterization allows us to quantify the revenue losses relative to both the case of full competition and the case of agency bidding under an alternative auction format (specifically, the VCG mechanism). We propose a simple algorithm that a search engine can use to reduce efficiency and revenue losses.

JEL: C71, D44, L41, L81, M37.

Keywords: Online Advertising, Internet Auctions, Common Agency.

*Boston University and Hariri Institute

[†]Royal Holloway, University of London

[‡]University of Wisconsin-Madison

We gratefully acknowledge financial support from the NET Institute, www.NETinst.org.

1 Introduction

Sponsored search auctions are the real time auctions used to allocate advertisement space on the results web page of search engines like Google, Microsoft Bing and Yahoo!. They represent one of the fastest growing and most economically relevant forms of online advertisement, with an annual growth of 14% and a total value of 35 billion dollars in 2011.¹ A novel development within these auctions is that advertisers are switching from individually managing their bidding campaigns to delegating them to specialized agencies known as Search Engine Marketing Agencies (SEMA).²

Outsourcing bidding to SEMA has several advantages for advertisers, which may be driven both by economies of scale and by economies of scope. As ad costs and competition in online auctions increase, due to the entry of new advertisers, using a SEMA has become key to maintain effective ad campaigns.³ Since all these motives are likely to persist in the next years, the expansion of outsourcing to agencies is expected to continue.

What the effect of this expansion will be is difficult to predict. On the one hand, agencies can help the functioning of this market by both fostering advertisers participation and improving the quality of the ads consumers receive. On the other hand, their presence can substantially alter the functioning of the auction mechanism used to sell ad space. Indeed, market data reveal that as agency usage increases and agencies become more specialized by products, it is becoming more common that a single agency bids in the same auction on behalf of multiple advertisers. This opens the door to the possibility that agencies take advantage of their control of different bidders as to lower their clients' payments. In other words, an agency can gain from implementing coordinated bidding among its clients, which will inevitably change the strategic interactions occurring in the search auctions. In this paper we focus on the latter aspect of this recent development, and analyze the impact of SEMA on the performance of the generalized second price (GSP) auction, the format used by all major search engines. In particular, we propose a theoretical model of bidding in the GSP when a subset of the advertisers delegate their bids to a common agency. The equilibrium predictions generated from this model are used to quantify revenue losses relative to the case of full competition and to propose an algorithm that a search engine can implement to reduce efficiency and revenue losses.

Given the large amount of information that bidders have about their competitors in these auctions, GSP auctions are often modeled as games with complete information (e.g., Varian (2007), Edelman Ostrovsky and Schwarz (2007, EOS hereafter), etc.).⁴ A major difficulty

¹Source: Econconsultancy (2011b).

²Indeed, a survey by the Association of National Advertisers (ANA) among 74 large U.S. advertisers indicates that about 77% of the respondents fully outsource their search engine marketing activities (and 16% partially outsource them) to specialized agencies, see ANA (2009, 2011). Analogously, a different survey of 325 mid-size advertisers by Econsultancy (EC) reveals that the fraction of companies not performing their paid-search marketing in house increased from 53% to 62% between 2010 and 2011, see EC (2011a).

³Google's publicly released results for the fourth quarter of 2010 revealed that the average cost per click on its search page had increased approximately by 5% over one year before.

⁴Athey and Nekipelov (2011) noticed that, since bids are typically held fixed across multiple auctions, bidders face uncertainty in both the set of rivals and the quality scores assigned by the search engine provider.

in the analysis of these auctions is due to the plethora of equilibria that they generate (see Borghers, Cox, Pesendorfer and Petricek, 2012). To overcome this problem, EOS and Varian (2007) put forward a particular equilibrium refinement, the ‘lowest-revenue envy-free (LREF) equilibrium’, which is both tractable and sustained by several foundations (see EOS and Edelman and Schwarz, 2011). From a positive perspective, this refinement is particularly relevant because it conforms to the guidelines on how to bid in these auctions that some of the search engines provide.⁵ For these reasons, the LREF equilibrium has become the standard solution concept to analyze GSP auctions with complete information.

As a first step towards understanding bid coordination in the GSP auction, we consider the same environment that has served as a workhorse for the early studies on the GSP auction (namely, EOS and Varian’s). We modify the baseline model introducing a SEMA, which we model as a player that chooses the bids of its clients. Hence, formally a SEMA consists of a coalition of advertisers that place their bids jointly.

Introducing the possibility of bids coordination in the GSP auction adds several layers of complexity. Besides the modeling choices regarding the SEMA, the ability to coordinate bids may exacerbate the already mentioned problem of multiplicity. In the absence of a refinement that could be applied to the environment with SEMA, the multiplicity problem in the GSP auction with coordination may be daunting. An insightful analysis of the problem thus requires a careful balance between tractability and realism of the assumptions.

For this reason we focus exclusively on environments that resemble what we observe in the market, and restrict our attention to cases where there is a single SEMA per auction. We further assume that the SEMA’s objective is to maximize the total surplus of its clients. This is a simplifying assumption, which can be justified in a number of ways. From a theoretical viewpoint, for instance, our environment is consistent with the informational assumptions of Bernheim and Whinston (1985, 1986). Hence, as long as the SEMA is risk-neutral, this particular objective function may be the result of an underlying common agency problem that determines the agency’s incentives. More relevantly from an empirical viewpoint, however, the agency contracts most commonly used in this industry are such that the SEMA receives a lump-sum fee per advertiser and per campaign (which typically last only a few months). Thus, the SEMA ability to generate surplus for its clients is an important determinant of its long run profitability. Our modeling assumption therefore is a reasonable proxy for the agency’s objective, unless the SEMA is particularly myopic.

Finally, consistent with the observation that agency contracts typically require that advertisers fully delegate bidding to the agency for a given amount of time, it may be tempting to assume that the agency can freely set the bids of its clients. This, however, overlooks the fact that advertisers choose to join the agency, and therefore the SEMA’s freedom is limited by its clients’ willingness to be part of the coalition. Since our primary objective is to understand the strategic bearings of SEMA in GSP auctions, we take the coalition structure as given, and do not model the coalition formation process explicitly.⁶ However,

⁵One example is the Google Ad Word tutorial on Youtube where Hal Varian teaches how to maximize profits by following the LREF equilibrium bidding strategy: <http://www.youtube.com/watch?v=jRx7AMb6rZ0>.

⁶On the other hand, understanding how given coalition structures impact the market seems a necessary preliminary step to any model of the advertisers’ decision to join a SEMA.

we maintain that when clients consider whether or not to abandon the SEMA, they not only anticipate the response of the other bidders, but also take into account the incentives that the other members of the coalition may have to further abandon the coalition, and the resulting response of the SEMA. Thus, the agency can be thought of as making a binding proposal of a certain profile of bids to its (potential) 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 SEMA and bid as an independent. Our approach to modeling coordination is thus very close to the seminal work of Ray and Vohra (1997) (see also Ray, 2008), in which the outside options of the members of a coalition are equilibrium objects themselves, and implicitly incorporate the restrictions entailed by the underlying coalition formation game.⁷

Given the model of the SEMA, there remains the problem of marrying the coordinated bids of coalition members with the competitive bidding of the advertisers that are not part of the SEMA, also referred to as ‘independents’. This is a delicate matter, because the multiplicity problem of the standard non-agency setting may be exacerbated by the stability restrictions that we imposed on the model with coordination. To address this problem, we introduce a refinement of bidders’ best responses that distills the individual-level underpinnings of the LREF equilibrium.⁸ This device enables us to maintain the logic of that refinement for the independent bidders, even if the LREF equilibrium is not defined in the game with SEMA. Furthermore, because it naturally extends the LREF equilibrium (and coincides with it if there is no SEMA in the auction), our formulation also allows a meaningful comparison with the non-agency benchmark of EOS and Varian (2007).

We begin our theoretical investigation first assuming that the agency operates under the constraint that its bids cannot be detected as ‘coordinated’ by an external observer. This is a useful working hypothesis, which also has obvious intrinsic interest. Under this assumption, we show that the resulting allocation in the GSP with SEMA is efficient and the revenues are the same as those that would be generated if the same coalition structure was bidding in a VCG auction. We then relax this ‘undetectability constraint’, and we show that in that case the search engine revenues in the GSP auction are never higher, and are in fact typically strictly lower, than those obtained in the VCG mechanism with coordination. Furthermore, once the ‘undetectability constraint’ is lifted, efficiency in the allocation of bidders to slots is no longer guaranteed by the GSP mechanism. Since the VCG is famously regarded as a poor mechanism in the presence of coordination, finding that it outperforms the GSP both in terms of revenues and allocation is remarkably negative for the GSP.

We conclude the paper with the proposal of a simple algorithm that the search engine provider can use to reduce the distortions due to coordinated bids. The algorithm uses the bids of the non-agency bidders to derive bounds on the minimum acceptable agency bids. The use of the algorithm ensures that in equilibrium the allocation will be efficient and

⁷We further discuss the relations with Ray and Vohra’s approach in Section 3

⁸In particular, we show that the fixed points of our refinement of the best response correspondence coincide with the LREF equilibria. This result can thus be seen as providing yet another, individual-based, foundation to the LREF equilibrium.

the revenues will be no less than in the VCG auction with agency bidding. We argue that in the typical application most of the gain in revenues comes from containing the drop in non-agency bidder payments. This indirect effect on revenues is typically larger than the direct effect due to agency payments and explains why even small 2-bidder coalitions can substantially alter the market if the search engine takes no action against bid coordination.

This paper has important policy implications because it identifies a trend toward agency bidding that might cause a deep transformation in sponsored search auctions. Our results suggest that the auction mechanism currently used is not well suited to face this new trend. The loss of profitability for search engines might have perverse effects on their incentive to invest in innovation. Moreover, inefficiency in the allocation of the ads resulting from coordination might lower consumers welfare by increasing their search costs. For both these reasons, it would be desirable to either replace the GSP with a mechanism more robust to this type of coordination, or to set reserve prices to weaken the effects of coordination. Alternatively, if the payments to marketing agencies were incorporated into the search engine platform, a pay-per-click payment scheme could be enforced to align the incentives of agencies with those of the search engine.

2 The GSP auction

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

In the ‘generalized second price’ (GSP) auction, each advertiser submits a bid $b_i \in \mathbb{R}_+$. Bids are ordered, and determine both the assignment and the price at which each slot is sold: the advertiser who submits the highest bid obtains the first slot and pays a price equal to the second highest bid every time his ad is clicked; the advertiser with the second highest bid obtains the second slot and pays a price-per-click equal to the third highest bid, and so on (ties are broken randomly). We denote the profile of bids 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 . When b is clear from the context, we will 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, for any b and i , $b_i \equiv b^{\rho(i)}$).

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

are common knowledge.⁹ We thus model the GSP auction as a game $\mathcal{G}(v) = (A_i, u_i^{\mathcal{G}})_{i=1, \dots, n}$ where $A_i = \mathbb{R}_+$ denotes the set of actions of player i (his bids), and the payoff functions are such that, for every i and every $b \in \mathbb{R}_+^n$,

$$u_i^{\mathcal{G}}(b) = (v_i - b^{\rho(i)+1}) x^{\rho(i)}.$$

2.1 Best Responses and Equilibria

Any profile $b_{-i} = (b_j)_{j \neq i}$ partitions the space of i 's bids, $\mathbb{R}_+$, into $S + 1$ intervals: $[0, b_{-i}^S)$, $[b_{-i}^S, b_{-i}^{S-1})$, $\dots$, $[b_{-i}^1, \infty)$. For any bid of player i , the interval that contains b_i determines the position obtained by player i . So, letting $b_{-i}^0 \equiv \infty$ and (with a slight abuse of notation) $b_{-i}^{S+1} \equiv 0$, if bidder i bids $b_i \in (b_{-i}^t, b_{-i}^{t-1})$, then he obtains slot $t = 1, \dots, S + 1$ at per-click-price b^t (slot $S + 1$ corresponds to not obtaining a position). The best response of bidder i boils down to the optimal choice of the position, which in turn determines the interval from which i should choose b_i . Note that this is the only payoff-relevant component of i 's choice: i would be indifferent among all bids that grant him the same position. Thus, for each i , let $\pi_i : \mathbb{R}_+^{n-1} \rightrightarrows \{1, \dots, S + 1\}$ denote the correspondence that assigns to each vector of bids of the opponents, the position(s) that player i would like to occupy. Formally, for each $b_{-i} \in \mathbb{R}_+^{n-1}$,

$$\pi_i(b_{-i}) = \arg \max_{t=1, \dots, S+1} (v_i - b_{-i}^t) x^t. \quad (1)$$

Clearly, π_i is always non-empty valued. It is non-single valued in case i is indifferent between two positions. We can ignore this case here (for instance, assuming that such ties are always broken lexicographically, say, 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 in (1)).

Then, the best-response correspondence of bidder i , $BR_i : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$, is defined as:¹⁰

$$BR_i(b_{-i}) = \left(b_{-i}^{\pi_i(b_{-i})}, b_{-i}^{\pi_i(b_{-i})-1} \right) \quad (2)$$

for any $b_{-i} \in \mathbb{R}_+^{n-1}$.

The set of Nash equilibria in the non-agency GSP auction is the set of fixed-points of the best-response correspondences:

$$\mathcal{EG}^0(v) = \{b \in \mathbb{R}_+^n : b_i \in BR_i(b_{-i}) \ \forall i \in I\}. \quad (3)$$

It is well-known that, even in the absence of agencies that may coordinate the bids of their members, the GSP auction admits a large set of equilibria, some of which are inefficient (see

⁹For studies of the GSP auction under incomplete information, see e.g. Athey and Nekipelov (2010) and Gomes and Sweeney (2011).

¹⁰Depending on the tie-breaking rule, one or both of the extremes of the intervals may be included. However, as already mentioned, bidder i would be indifferent among all those bids that grant him the same position. Thus, defining BR_i in terms of the open interval entails no essential loss of generality.

Borgers et al., 2008). Refining such multiplicity therefore is key to the analysis. We introduce next an equilibrium refinement based on a selection of the best response correspondence. We will show that, with no agencies, this refinement coincides with the refinements introduced by Varian (2007) and EOS, which became the standard to analyze the GSP auction.

Let $BR_i^* : \mathbb{R}_+^{n-1} \rightrightarrows \mathbb{R}_+$ be such that, for any $b \in \mathbb{R}_+^{n-1}$,

$$BR_i^*(b_{-i}) = \left\{ b_i^* \in BR_i(b_{-i}) : \left(v_i - b_{-i}^{\pi_i(b_{-i})} \right) x^{\pi_i(b_{-i})} = (v_i - b_i^*) x^{\pi_i(b_{-i})-1} \right\}. \quad (4)$$

In words, of the many $b_i \in BR_i(b_{-i})$ that would grant player i his favorite position $\pi_i(b_{-i})$, he chooses the bid b_i^* that makes him indifferent between occupying the current position and climbing up one position paying a price equal to b_i^* .

The equilibrium refinement that we propose is the set of fixed points of the BR^* correspondence:

Definition 1. *The set of marginally envy-free (MEF) equilibria of baseline GSP auction game $\mathcal{G}(v)$ is defined as*

$$\mathcal{EG}(v) = \{ b \in \mathbb{R}_+^N : b_i \in BR_i^*(b_{-i}) \text{ for all } i \in I \} \quad (5)$$

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

$$b_i = v_i - \frac{x^i}{x^{i-1}} (v_i - b_{i+1}). \quad (6)$$

Hence, the marginally envy-free equilibrium is unique, up to the highest bid. Furthermore, it is efficient: for any v , for any $b \in \mathcal{EG}(v)$ and for any $i \in I$, $i = \rho(i)$.

Remark 1. *Notice that the recursion in 6 coincides with the recursions of the EOS' LREF equilibrium (EOS, Theorem 2), and with Varian's lower-bound symmetric Nash Equilibrium (Varian, 2007, equation 9).*

From now on, we will use the terms “marginally envy-free” (“MEF”) and “lowest-revenue envy-free” (“LREF”) interchangeably.

2.2 Comparison with the VCG mechanism

Because of its well known properties, the VCG mechanism (Vickrey, 1961, Clarke, 1971 and Groves, 1973) represents the standard benchmark in the literature on the GSP auction (see, for instance, Varian (2007, 2009) and EOS). With only one slot on sale, the GSP auction coincides with the VCG mechanism. With two slots or more, however, the two mechanisms differ.

In the VCG mechanism advertisers submit bids $b_i \in \mathbb{R}_+$. The allocation rule is the same as in the GSP auction: given bids profile b , advertiser i obtains position $\rho(i)$ if $\rho(i) \leq S$; he obtains no position otherwise. In the VCG mechanism, each advertiser pays a price equal

to the externality that his presence imposes on the others if each advertiser's bid equals his true valuation. To calculate these payments, note that if the advertiser in position s were to leave the market, those ranked above s would not be affected, while those ranked below would each gain one position. Thus the total externality that this advertiser imposes on others (and thus also this advertiser's payment in the VCG mechanism) is equal to

$$p^{\mathcal{V};s}(b) = \sum_{t=s+1}^{S+1} b^t(x^{t-1} - x^t). \quad (7)$$

The VCG mechanism therefore induces a game $\mathcal{V}(v) = (A_i, u_i^{\mathcal{V}})_{i=1,\dots,n}$, where $A_i = \mathbb{R}_+$ is the set of actions of player i , and payoff functions $(u_i^{\mathcal{V}})_{i=1,\dots,n}$ are such that, for every i and every $b \in \mathbb{R}_+^n$,

$$u_i^{\mathcal{V}}(b) = v_i x^{\rho(i)} - \sum_{t=\rho(i)+1}^{S+1} b^t(x^{t-1} - x^t). \quad (8)$$

It is well-known that bidding $b_i = v_i$ is a (weakly) dominant strategy in this game. In the resulting equilibrium, advertisers are efficiently assigned to positions (that is, advertiser i obtains position i for all $i \leq S$). Furthermore, EOS (Theorem 1) showed that the position and payment of each advertiser in the unique (weakly) dominant strategy equilibrium are equal to those that obtain in the LREF equilibrium of the GSP auction. The following result thus follows immediately from Lemma 1 and Remark 1:

Remark 2. *The total revenue to the seller in the marginally envy-free equilibrium of $\mathcal{G}(v)$ is equal to the revenue in the dominant strategy equilibrium of the VCG mechanism, $\mathcal{V}(v)$.*

3 Search Engine Marketing Agencies

Our analysis of the Search Engine Marketing Agencies (SEMA) focuses on their opportunity to coordinate the bids of different advertisers. We thus borrow the language from the literature on cooperative game theory and refer to the clients of the agency as ‘members of a coalition’ and to the remaining bidders as ‘independents’. On the other hand, institutional aspects of the market for sponsored search auctions make a straightforward application of existing models of coalition behavior problematic. We discuss these points and the related literature on bidder coalitions in Section 3.2

We model the SEMA as a player that makes proposals of binding agreements to its members. The agency's proposals consist of a collection of bids, one for each member of the coalition. The independents then play the subgame which ensues from taking the bids of the agency as given.

We assume that the agency forms beliefs about the behavior of the independents. In equilibrium, such beliefs will be required to be correct. Given such beliefs, the agency chooses proposals that maximize the coalition surplus (i.e., the sum of the profits of its

members). The agency, however, can only choose proposals that are ‘stable’ in two senses: first, they are consistent with the independents playing an equilibrium in the continuation game; second, they are stable in the sense that no individual member of the coalition has an incentive to abandon the coalition and play as an independent. When considering such deviations, the members of the coalition are ‘farsighted’ in the sense that they anticipate the impact of their deviation on both the independents and the remaining members of the coalition. The constraint for a coalition of size C thus depends on the solutions to the problems of all the subcoalitions of size $C - 1$. Therefore, the solution concept for the game with the agency will be defined recursively.

3.0.1 The Recursively Stable Agency Equilibrium

Let $G(v) = (A_i, u_i^G)_{i=1,\dots,n}$ denote the baseline game generated by the underlying mechanism (e.g., the GSP ($G = \mathcal{G}$) or the VCG ($G = \mathcal{V}$) mechanism), given the profile of valuations $v = (v_i)_{i \in I}$. For any $\mathcal{C} \subseteq I$ with $|\mathcal{C}| \geq 2$, we let $\mathcal{C}$ denote the agency, and we refer to advertisers $i \in \mathcal{C}$ as ‘members of the coalition’ and to $i \in I \setminus \mathcal{C}$ as the ‘independents’. The coalition chooses a vector of bids $b_{\mathcal{C}} = (b_j)_{j \in \mathcal{C}} \in \times_{j \in \mathcal{C}} A_j$. Given $b_{\mathcal{C}}$, the independents $i \in I \setminus \mathcal{C}$ simultaneously choose bids $b_i \in A_i$. We let $b_{-\mathcal{C}} = (b_j)_{j \in I \setminus \mathcal{C}}$ and $A_{-\mathcal{C}} = \times_{j \in I \setminus \mathcal{C}} A_j$. Finally, given profiles b or $b_{-\mathcal{C}}$, we let $b_{-i,-\mathcal{C}}$ denote the subprofile of bids of all independents other than i (that is, $b_{-i,-\mathcal{C}} = (b_j)_{j \in I \setminus \mathcal{C}: j \neq i}$).

We assume that the agency maximizes the sum of the payoffs of its members, denoted by $u_{\mathcal{C}}(b) = \sum_{i \in \mathcal{C}} u_i(b)$, under three constraints. Two of these constraints are stability restrictions: the agency can only submit bids that induce an equilibrium among the independents and that give no member of the coalition an incentive to leave. The third constraint, which we formalize as a set $R_{\mathcal{C}} \subseteq A_{\mathcal{C}}$, allows us to accommodate the possibility that the agency may exogenously discard certain bids (this restriction is vacuous if $R_{\mathcal{C}} = A_{\mathcal{C}}$). For instance, we will consider the case of an agency whose primary concern is to not be identified as inducing coordinated bids. In that case, $R_{\mathcal{C}}$ would be comprised of only those bids that are ‘undetectable’ to an external observer as coordinated (see below). We denote the collection of exogenous restrictions for all possible coalitions as $R = (R_{\mathcal{C}'})_{\mathcal{C}' \subseteq I: |\mathcal{C}'| \geq 2}$.

For any $i \in I \setminus \mathcal{C}$, we let $BR_i^G : A_{-i} \rightrightarrows A_i$ denote some refinement of player i ’s best response correspondence in $G(v)$ (e.g., BR_i^* in $\mathcal{G}(v)$ or truthful bidding in $\mathcal{V}(v)$). We define the *independents’ equilibrium correspondence* $BR_{-\mathcal{C}}^G : A_{\mathcal{C}} \rightrightarrows A_{-\mathcal{C}}$ as

$$BR_{-\mathcal{C}}^G(b_{\mathcal{C}}) = \{b_{-\mathcal{C}} \in A_{-\mathcal{C}} : b_j \in BR_j^G(b_{\mathcal{C}}, b_{-j,-\mathcal{C}}) \ \forall j \in I \setminus \mathcal{C}\}. \quad (9)$$

In general, this correspondence may be empty-valued for some bids of the agency. The first stability restriction for the agency requires that it places bids that induce a stable outcome among the independents. If a proposal is not consistent with the equilibrium behavior of the independents (as specified by $BR_{-\mathcal{C}}^G$), then the proposal does not induce a stable agreement. We thus incorporate this stability constraint into the decision problem of the agency, and assume that the agency can only choose bids profiles from the set $S_{\mathcal{C}}$,

defined as follows:

$$S_{\mathcal{C}} = \{b_{\mathcal{C}} \in A_{\mathcal{C}} : \exists b_{-\mathcal{C}} \text{ s.t. } b_{-\mathcal{C}} \in BR_{-\mathcal{C}}^G(b_{\mathcal{C}})\}. \quad (10)$$

Clearly, the strength of this constraint in general depends on the particular equilibrium correspondence $BR_{-\mathcal{C}}^G$ that is chosen to model the behavior of the independents, and on the structure of the underlying game $G(v, x)$. Although conceptually important, this restriction is not important for most of the results that follow, because the corresponding constraint would be either redundant or vacuous. In particular, this will be the case in the context of Theorem 1, where the constraint is vacuous, and in the context of Theorem 2, where it is redundant.

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

$$B^* = \{\beta \in A_{-\mathcal{C}}^{S_{\mathcal{C}}} : \beta(b_{\mathcal{C}}) \in BR_{-\mathcal{C}}^G(b_{\mathcal{C}}) \text{ for all } b_{\mathcal{C}} \in S_{\mathcal{C}}\}. \quad (11)$$

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

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

$$E^1(BR^G) = \{b \in \mathbb{R}_+^n : b_i \in BR_i^G(b_{-i}) \text{ for all } i \in I\}. \quad (12)$$

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

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

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

Definition 2. A RAE of the game G with coalition $\mathcal{C}$, given restrictions $R_{\mathcal{C}}$ and refinement BR^G , is a profile of bids and conjectures $(b^*, \beta^*) \in A_{\mathcal{C}} \times B^*$ such that:

1. The agency best responds to the conjectures β^* , given the exogenous restrictions (R) and the stability restrictions about the independents and the coalition members (S.1

and $S.2$, respectively):

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

2. The independents play a mutual best response: for all $i \in I \setminus \mathcal{C}$, $b_i^* \in BR_i^G(b_{-i}^*)$.
3. The conjectures of the agency are correct: $\beta^*(b_{\mathcal{C}}^*) = b_{-\mathcal{C}}^*$.

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

$$E^{\mathcal{C}}(BR^G, R) = \{b^* \in B^* : \exists \beta^* \text{ s.t. } (b^*, \beta^*) \text{ is a RAE}\}. \quad (13)$$

In the following we will apply the RAE to study the impact of a SEMA on the GSP auction, and compare it to the benchmark VCG mechanism. The RAE in the GSP and the VCG mechanism are obtained from definition 2, once the game G , the correspondence BR^G and the exogenous restrictions R are accordingly specified:

Definition 3. *The RAE-outcomes in the GSP and VCG mechanism are defined as follows. Given constraints R :*

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

For both mechanisms, we will refer to the case where $R_{\mathcal{C}} = A_{\mathcal{C}}$ as the ‘unconstrained case’.

Notice that, under this definition, the RAE-outcomes for the coalitions of size one (equation 12) coincide precisely with the non-agency equilibria introduced in Section 2, namely, the marginally envy-free equilibria in the GSP (Definition 1) and the dominant strategy equilibrium in the VCG.

3.0.2 Discussion

Our notion of RAE is closely related to the ‘Equilibrium Binding Agreements’ that Ray and Vohra (1997, RV hereafter).¹¹ Given a certain coalition structure, RV’s model postulates

¹¹See also Ray (2008) for a thorough introduction to the approach.

that binding agreements are only possible within a coalition. The objective is to endogenize the collection of agreements based on stability considerations: the agreements within each coalition must be such that no subcoalition has incentives to break the agreement and separate from the original coalition. 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 she 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, a consistency condition requires such predictions to be correct. Similarly to RAE, and because of the ‘farsightedness assumption’, RV’s definition of the equilibrium binding agreements is also recursive.

RV’s and our approach share the same fundamental philosophy. Similarly to 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 non-cooperative. In the language of RV, the agency in our model is a proposer of a binding agreement, subject to similar constraints as in RV. There are two main differences though. On the one hand, the stability restriction concerning the members of the coalition that we consider only refers to the possibility that the agency’s proposal is ‘blocked’ by individual members. In the solution concept of RV, the proposal can be instead blocked by the joint deviation of a subset of the coalition. That advertisers can make binding agreements outside of the domain of the SEMA does not seem realistic in this market. A straightforward application of RV’s solution concept would entail unrealistic assumptions in this specific context. For this reason, both the interaction with the independents and the ‘blocking’ of the agency’s proposal are fully non-cooperative in our model.

The second difference, relative to RV, is in the model of non-cooperative interaction (among the independents, and between them and the coalition). Being concerned with a general theory of coalition formation, the baseline non-cooperative solution concept in RV is Nash equilibrium. We can accommodate this letting BR_i^G in definition 2 coincide with the standard best-response correspondence. In the following, however, we will adopt refinements of the best-response correspondence for the independents, hence the underlying non-cooperative solution concept will be a refinement of Nash Equilibrium. There are two main reasons for this, which we already discussed in the introduction: first, to overcome the pervasive multiplicity of equilibria in the GSP auction; second, to enable us to make a meaningful comparison with the non-agency benchmarks of EOS and Varian (2007). For these reasons, the underlying solution concept in the GSP auction will be a natural extension of the LREF equilibrium.

3.1 Results

3.1.1 The VCG mechanism

As anticipated in Definition 3, in studying the impact of the SEMA on the VCG mechanism we apply RAE to the game $\mathcal{V}(v)$, letting R_C equal to A_C and assuming that the independents play the weakly dominant strategy, ‘truthful bidding’. Notice that, under this

specification of the independents' equilibrium correspondence, $S_C = A_C$, hence constraint (S.1) in Definition 2 could be dropped.

We next characterize the RAE-outcomes of the VCG mechanism (Definition 3, part 2). For equilibrium to exist in this game, we must assume that only bids on a (fine) discrete grid can be submitted, so that each set of available bids that is bounded below contains a minimum. With this assumption, for any $b_i \in \mathbb{R}_+$ we can define b_i^+ to be the lowest possible bid higher than b_i . Now, the following result obtains:

Theorem 1. *For any $\mathcal{C}$, the unrestricted RAE of the VCG is unique up to the bid of the highest coalition member. In this equilibrium, advertisers are assigned to positions efficiently, independents' bids are equal to their valuations and all the coalition members (except possibly the highest) bid the lowest possible value that ensures their efficient position. Formally, in any RAE of the VCG mechanism, the bids profile b^* is such that*

$$b_i^* \begin{cases} = v_i & \text{if } i \in I \setminus \mathcal{C}; \\ = b_{i+1}^+ & \text{if } i \in \mathcal{C} \setminus \min(\mathcal{C}) \text{ and } i \leq S; \\ \in (b_{i+1}^+, v_{i-1}) & \text{if } i = \min(\mathcal{C}) \text{ and } i > S. \end{cases}$$

where we denote $v_0 := \infty$ and $b_{n+1}^* := 0$.

The restrictions entailed by RAE are key to obtain this result. As the following example shows, without such restrictions, inefficiency is possible in the VCG with coordination:

Example 1. *Let $n = 4$, and $v_1 = 40$, $v_2 = 25$, $v_3 = 20$ and $v_4 = 10$. Let the CTRs be such that $x_1 = 20$, $x_2 = 10$, $x_3 = 9$ and $x_4 = 0$, and let the coalition be $\mathcal{C} = \{1, 2\}$. Notice that the coalition's payoff when it takes its efficient positions is bounded above by 630. However, the coalition could achieve a higher payoff by letting bidder 2 shift one position down. For example, if it sets $b_2 = 11$, the coalition's total payoff increases to 634. This deviation, however, is ruled out by RAE, because it violates the stability restriction (S.2) for the lowest coalition member. In this case, $u_2^C = 135$, which is less than agent 2's payoff in the equilibrium of the VCG without coalitions, equal to 140.*

3.1.2 The GSP auction

We turn next to the GSP auction. According to the refinement of the best responses introduced above, we set BR_i^G in equation 9 equal to the 'marginally envy-free' best response correspondence BR_i^* (eq. 4). The resulting correspondence $BR_{-\mathcal{C}}^G$ therefore assigns, to each profile b_C in the set of exogenous restrictions R_C , the set of independents profiles that are fixed points of the BR_i^* correspondence for all $i \notin \mathcal{C}$. We begin the analysis considering the case in which the agency faces the following exogenous constraint:

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

Notice that, for R_C thus defined, $BR_{-\mathcal{C}}^*(b_C) \neq \emptyset$ for all $b_C \in R_C$. Hence, $R_C \subseteq S_C$, so that constraint (S.1) in Definition 2 is redundant.

We refer to this constraint as the ‘feigned valuations restriction’. We introduce it mainly for analytical convenience. This constraint, however, has a clear interpretation: the set $R_{\mathcal{C}}$ is comprised of all bid profiles of the agency that could be observed as part of a MEF equilibrium in the non-agency GSP auction in which independents’ valuations are $(v_j)_{j \in I \setminus \mathcal{C}}$.

For instance, consider a SEMA whose primary interest is not being detectable for inducing equilibria with coordination by an external observer. The external observer (e.g., the search engine or the anti-trust authority) can only observe the bids profile, but not the valuations $(v_i)_{i \in \mathcal{C}}$. Then, $R_{\mathcal{C}}$ characterizes the set of bids that the agency can be sure would not be detected as coordinated bids, even if the independents had revealed their own valuations to the external observer.

Besides being useful for analytical purposes, the next theorem therefore may be of independent interest for at least two reasons: one, it characterizes the set of REA in a market in which ‘not being detectable’ is a primary concern of the SEMA; second, it provides an upper bound to the seller’s revenue loss that the SEMA may induce without being detectable by an analyst that approaches the market through the lens of EOS and Varian’s LREF equilibrium:

Theorem 2. *For any $\mathcal{C}$, an RAE of the GSP auction with the ‘feigned-valuations restriction’ is unique up to the highest coalition bid and the highest overall bid. In this equilibrium, advertisers are assigned to positions efficiently, and each advertiser’s payment equals that advertiser’s payment in the RAE of the VCG mechanism (Theorem 1).*

It is well known that the performance of the VCG mechanism with bidder coalitions is rather poor. Hence, Theorem 2 gives a rather negative outlook on the sellers’ revenues in the presence of SEMA, under the feigned valuations restriction. Clearly, if the latter constraint is relaxed, for any given $\mathcal{C}$, while maintaining the constraints for all its subcoalitions, then the coalition’s payoff in this ‘relaxed RAE’ would be (weakly) higher than their counterparts in the RAE of the VCG mechanism. Hence, with the ‘relaxed RAE’, the seller’s revenue in the VCG mechanism actually constitutes an upper bound to the revenues in the GSP.

However, if the feigned valuations restriction is relaxed not only for the optimization problem of coalition $\mathcal{C}$, but also for all subcoalitions $\mathcal{C}' \subseteq \mathcal{C}$, the result that the seller’s revenues would decrease is not immediate. Indeed, whereas constraint (R) is relaxed for the coalition $\mathcal{C}$, the constraint (S) may become more stringent, as a result of the fact that constraint (R) is relaxed for the subcoalitions as well. The next result, however, shows that overall direct the effect of relaxing (R) dominates.

Theorem 3. *For any $\mathcal{C}$, in any RAE of the GSP auction, the agency’s payoff is at least as high and the auctioneer’s revenue no higher than their counterparts in the RAE of the VCG mechanism. Furthermore, (i) there exist parameter values under which the ordering is strict and (ii) there exist parameter values under which the allocation is inefficient.*

Example 2. *As an example where the allocation is inefficient consider the following case:*

$$\begin{aligned} v &= (12, 10.5, 10.4, 10.3, 10.2, 10.1, 10, 1); \\ x &= (50, 40, 30.1, 20, 10, 2, 1, 0); \\ \mathcal{C} &= \{5, 6\}, \end{aligned}$$

the equilibrium is 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:

$$\{(b_2, b_3, b_4, b_5, b_6, b_7, b_8) = (9.91, 9.76, 9.12, 9.5, 7.94, 5.5, 1); b_1 > 9.91\}.$$

The inefficiency arises as follows. The agency drastically reduces its lower-valued member's bid to benefit the other member. This, however, creates incentives for independents with values above those of the higher-valued coalition member to move down to the position just above the lower-valued member, thus stealing some of the rents generated by the bid reduction. In order to prevent these independents from doing so, the higher-valued coalition member's bid must also be reduced to make the higher positions more attractive. In this example, the optimal reduction is actually large enough to cause the higher-valued bidder lose a position.

4 Guide for Empirical Application

The model developed in the previous section can be used to develop a method that the search engine provider can use to limit the harm created by common agency bidding. In this section, we first discuss how to extend the model for an environment with quality scores. Then, we describe our proposed method.

4.1 Quality Scores

All the major search engines, including Google and Microsoft-Yahoo!, rank advertisers not by their bids but by bids adjusted for quality scores. In particular, the search engine assigns to advertiser i a score, e_i , related to the advertiser's probability of being clicked, independently from the slot to which it is assigned. The ranking of advertisers follows the product of scores and per-click bids. Since bidder i in position $\rho(i)$ is required to pay the minimum bid to keep that position, his payment when bidder k sits in position $\rho(i) + 1$ is:

$$p^{\rho(i)} = \frac{e_k b_k^{\rho(i)+1}}{e_i} \quad (14)$$

Quality weights capture the idea that advertisers in the same position might get a different number of clicks because consumers have a different perception about the likelihood that their ad is worth clicking. The average probability that a consumer clicks on a particular ad can be expressed as a function of both position and advertiser specific effects as: $c_i^j = \gamma_i x^j$. The search engine scores are perfectly aligned to the consumers tastes only when $\gamma_i = e_i$. With quality scores and non-coordinated bids, the indifference condition of the LREF is:

$$x^{\rho(i)-1}(e_i v_i - e_i b^{\rho(i)}) = x^{\rho(i)}(e_i v_i - e_k b_k^{\rho(i)+1}) \quad (15)$$

When an agency is present, it seems plausible to assume that it can directly control only its bids, while it cannot affect γ and e . Under this assumption, we now describe how the search

engine can limit the harm due to coordinated bids.

4.2 Correcting Bids against Bid Coordination

Suppose that the search engine observes data on bids, quality scores, click through rates and the identity of both the advertisers and the SEMA, if any, bidding for them. Then, the method that we describe below allows the search engine to prevent the worst case scenario of agency bidding without the feigned valuations restriction. Our algorithm sets a minimum bid from the low-value agency bidder such that it induces the same revenues and allocation of the case with feigned valuations.

We begin our discussion from the case of a 2-bidder coalition. Notice that in this case both the low-value agency-bidder and all the non-agency bidders below it are placed efficiently. For this latter subset of non-agency bidders, equation (15) can be applied recursively to uniquely recover their valuations using data on bids, quality scores and click through rates. Then, if i is the low-value agency bidder and k the first non-agency bidder below it, it must be that this agency bidder has a value $e_i v_i \geq e_k v_k^+$. By replacing $e_k v_k^+$ for $e_i v_i$ into equation (15) and solving for b_i we obtain the minimum bid that should be asked to the low-value agency bidder, $\underline{b}_i$. Since $\underline{b}_i$ corresponds to this advertiser's bid under the feigned valuation restriction, it is an equilibrium for the other agency bidder to behave as in that case. Thus, the search engine provider can achieve both efficiency and an amount of revenue not lower than the VCG.

Furthermore, we can assess whether the data are compatible with agency bidding with or without the feigned valuation restriction: A positive difference between $\underline{b}_i$ and the b_i observed in the data is compatible with an agency coordinating bids under no feigned valuation restriction, while a zero difference is compatible with bidding under feigned valuations. A negative difference, instead, is a rejection of the type of bid coordination strategies that we analyzed. Imposing a minimum bid is useful only when there is a positive difference between $\underline{b}_i$ and b_i . In this case, the search engine revenues increase through the higher price paid by both the high-value agency bidder and the non-agency bidders placed above the low-value agency bidder. This stresses that, even with coalitions that are small (2-bidder) and even if they occupy low positions, preventing the agency from behaving as in the case of the equilibrium without feigned valuations can have large effects on the search engine revenues due to the effect on the payments from non-agency bidders placed above the coalition members. Finally, efficiency is restored through our proposed correction because the high-value agency bidder will prefer to be efficiently placed given that all remaining bids follow the LREF formulation.

For coalitions larger than two, the search engine provider can apply the minimum bid correction iteratively to the lowest agency bid not compliant with the feigned valuation restriction. However, this solution might not be effective for two reasons. First, if the agency knows about this correction, it might respond to it using more sophisticated dynamic strategies that are not captured by our current model. Second, when $\gamma \neq e$, our procedure does not necessarily restore efficiency. In fact, the agency might choose to order inefficiently

its advertisers above the low-value one to exploit the mismatch between the quality scores and the consumer’s click probabilities. Thus, the effectiveness of our proposed correction for coalitions larger than two depends on the extent to which quality scores match click probabilities and on the agency’s knowledge of the procedure the search engine uses to undue common agency bidding.

5 Conclusions

This paper represents the first study of common agency bidding in the GSP auction. Our analysis intends to show the potentially perverse effects on revenues and efficiency that the GSP format can have when a single marketing agency bids for multiple clients. This phenomenon has potentially large economic consequences given the growing importance of SEMA for the GSP auctions that search engine providers use to sell advertisement space worth billions of dollars. Our analysis reveals that even the small 2-bidder coalitions can have large effects. Indeed, lowering competition between any two bidders will typically induce lower competition faced by *all* bidders, hence reducing prices across the board.

These findings open interesting avenues for both empirical and theoretical research. As regards the latter, the most important development would likely regard devising an alternative market design for the sale of advertisement space that would be less vulnerable to bid coordination. Furthermore, it would be interesting to analyze the implications of coordinated bidding in the GSP under the incomplete information formulation proposed by Athey and Nekipelov (2011). For empirical research, it would be important to evaluate the trade-off posed by the presence of SEMA. Indeed, while we intentionally focused only on the role of coordinated bidding, agencies certainly play other roles as well, like fostering entry and improving ad quality. Our tractable model of agency bidding might be a fruitful starting point for these future research developments.

6 References

- Association of National Advertisers (ANA), (2011). *Trends in Digital Agency Compensation*, 4th Edition, ANA publishing.
- Athey, S., D. Nekipelov, (2010). “A Structural Model of Sponsored Search Advertising Auctions,” mimeo.
- Bernheim, B. D., M. D. Whinston, (1985). “Common Marketing Agency as a Device for Facilitating Collusion,” *The RAND Journal of Economics*, 16(2), 269–281.
- Bernheim, B. D., M. D. Whinston, (1986). “Common Agency,” *Econometrica*, 54(4), 923–942.
- Borgers, T., I. Cox, M. Pesendorfer, V. Petricek, (2012). “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” *American Economic Journal: Microeconomics*,

forthcoming.

Econconsultancy (EC), (2011a). “State of Search Marketing Report 2011,” accessed on January 2012 at: <http://econsultancy.com/us/reports/sem-po-state-of-search>

Econconsultancy (EC), (2011b). “Search Marketing Statistics Compendium,” 2011

Edelman, B., M. Ostrovsky, M. Schwarz, (2010). “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97(1), 242–259.

Edelman, B., M. Schwarz, (2007). “Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions,” *American Economic Review*, 100(2), 597–602.

Ray, D., (2008). *A Game-Theoretic Perspective on Coalition Formation*, New York: Oxford University Press.

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

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