governance & regulation

Fondation Paris-Dauphine

DAUPHINE UNIVERSITÉ PARIS | PSL

GovReg WORKING PAPER SERIES

Collusion with a rent-seeking agency in sponsored search auctions

Emmanuel Lorenzon

Issue 2018/03

RESEARCH

Collusion with a rent-seeking agency in sponsored search auctions *

Emmanuel Lorenzon[†1]

[1]University Paris-Dauphine, PSL Research University, F-75016 Paris, France

*Acknowledgments: I am grateful to Kenneth Corts, Daniel Spulber, Byung-Cheol Kim, Wilfried Sand-Zantman, Nicolas Carayol, Leslie Marx, David Ettinger, Joëlle Toledano, Claude Lacour, Iryna Topoylan, Germain Gaudin, Hervé Hocquard for their very helpful comments and discussions. I am also grateful for the comments received from the participants at: the IOEA 2018, the 2017 Economics of digitization seminar at Telecom ParisTech, the 2016 Seventh Annual Searle Center Conference on Internet Commerce and Innovation in Chicago, the 2016 43rd EARIE Annual Conference in Lisbon, the 2016 33rd JMA in Besançon, the 2015 Workshop in Industrial Organization and Marketing in Lisbon.

†*e-mail:* emmanuel.lorenzon@dauphine.fr

Abstract

We analyse the implications of delegating bids to a *bidding agency* for the revenues and efficiency of the *Generalised Second-Price* auction, the standard sales mechanism for allocating online ad space. The *agency* maximises both its own profits and the advertisers' surplus and implements collusive agreements by means of side contracts. Our results suggest that an *agency* can profitably deliver bid delegation services, which increase the advertisers' surplus and contribute to market efficiency. The agency coordinates advertisers on a unique efficient equilibrium and extracts an optimal fee on the collusive gain that implies a lower bound on the seller's revenues. We also find that bid delegation uniquely selects the *Vickrey-Clarke-Groves* outcome if side contracts are based on the *locally envy-free* stability criterion.

Keywords: Generalised second-price auction; Position auction; Bidding ring; Cartel; Bidding agency; Sponsored Search

JEL Classification: D44, C72, M3, L41

1 Introduction

Billions of simultaneous sponsored search auctions — called *Generalised Second-Price* (*GSP*) auction — are held daily to allocate bundles of online ad spaces by Google, Microsoft and Yahoo![1][2]. Each time a customer makes a search query, an automated auction is triggered putting advertisers in competition for ad space. This sales mechanism has become the main source of revenue for search engines and is the most widely used auction ever designed. In 2016, earnings from advertising represented more than 90% of Google's revenues in an online advertising market that in the same year was worth nearly $190 billion[3].

The concentration of marketing activities and the novel strategic considerations it represents for the players engaged in the online ad market have led an increasing number of advertisers to delegate their search marketing and paid search strategies to third parties — *bidding agencies* — that specialise in these activities and which also belong to the ad networks in which they conduct their bidding activities. For instance, Gartner's U.S. Digital Marketing Spending report for 2013 documents that from a pool of 243 firms more than 50% delegate their search marketing activities[4]. This means a *bidding agency* is likely to end up acting on behalf of different advertisers during the same set of keyword auctions, and so it can affect both allocative and revenue performance of the sales mechanism by coordinating advertisers' strategies. Thus, in a market, essentially driven by the Google-Facebook duopoly and which generated around 20% of the global ad spending in 2016, it is clearly important to understand how rent is allocated among players and how these trend in delegation affect market efficiency.

Despite its importance and potentially adverse outcomes, advertising intermediation — and bid delegation — has been overlooked by the main strand of research analysing sponsored search

[1]The auction works as follows: advertisers announce the maximum price they are willing to pay for each click on their ad, referred to as their bid for a click, and they pay a price (per click) equal to the bid announced by the advertiser assigned to the position just below them. In this way, the highest bidder is allocated the ad slot at the top of the page receiving the highest number of clicks, the second-highest bidder is allocated to the second position receiving the second-highest number of clicks, and so on.

[2]A recent exception to this is provided by Facebook, which adopted the *Vickrey-Clarke-Groves* (*VCG*) mechanism in 2015.

[3]See the eMarketer's Worldwide Ad Spending report and Google annual report 2016

[4]Another survey, run in 2013 by Constant Contact, found that in a pool of 1305 small firms nearly 35% outsourced their bid management activities. See also Decarolis et al. (2017) for a discussion of ad network concentrations.

auctions, as pioneered by Aggarwal et al. (2006), Varian (2007) and Edelman et al. (2007)[5]. The unique quality of ad space is that its value (the expected return on investment) is determined by consumers search behaviour, which constitutes complex and costly information for firms without the skills to assess. Against this backdrop, a *bidding agency* can offer specialised skills and a service that can internalize advertisers' search costs. Bid delegation acts as buffer between the search engine and the advertisers, obstructing rent extraction and raising questions regarding the way it is shared out between advertisers, the search engine and the *agency*. At the same time, by bidding on behalf of different advertisers during the same auction (and on the same keywords), the *agency* is at liberty to manipulate its clients final payments by coordinating individual bids, which in the spirit of classic collusion games relies on distorting the rivalry between advertisers.

In this paper, we undertake a theoretical exploration of the way in which advertiser collusion — by means of bid delegation — jeopardises the revenues and distorts rent allocation of the *GSP* auction by considering a *bidding agency* that seeks to minimise its clients' rivalry and final payments but also to maximise its own profits. We focus on a collusion enhanced by a *pre-communication* phase and in the spirit of the mechanism proposed by Graham and Marshall (1987), we introduce a collusive device to the model developed originally by Varian (2007) and Edelman et al. (2007), grounded by side payments between advertisers and we limit our attention to a simultaneous two-position game with three advertisers[6]. Our results suggest that the *agency* can profitably deliver bid delegation services. Its presence in the market increases the advertisers' surplus and contributes to market efficiency. In short, bid delegation acts as an efficient collusion constrained by the *agency*'s participation fee, which limits the adverse effect of a complete cartel on auction revenues and rent expropriation.

We consider a market organisation that has sevral features in common with menu auction and common agency games, as analysed by Bernheim and Whinston (1985, 1986a,b), in which several principals delegate their right to make a decision to a single common agent. Here, the

[5]Studies in this literature have examined such questions as the optimal design of the auction mechanism (e.g., Edelman and Schwartz (2010)), the introduction of consumer search (e.g., Athey and Ellison (2011)), search engine competition (e.g., Ashlagi et al. (2011)), budget constraints (e.g., Ashlagi et al. (2010)) and more recently, incomplete information settings (e.g., Gomes and Sweeney (2014)).

[6]We consider an all-inclusive collusion, in which reducing to the smallest set of advertisers and positions does not affect the soundness of the model and where the main insights carry can be applied to more positions and players.

bidding agency is assumed to be risk-neutral, to support no unit cost and to begin the procedure by making an individual contract proposal that includes a bid profile to be placed on the client's behalf, a management fee and the way in which side payments are to be implemented[7]. The *allocation rule* embodied in the contract proposal ensures a monetary payback to each member and the *agency* computes this according to each contribution to the collusion. Contributions are determined as the difference in the prices from which the members benefit when the collusion is in operation and are assumed to be defined *prior* to the auction and supported as a *sunk* cost by each member. The rules of the *GSP* do not permit one bid per position, all advertisers seek the top position and the incentives to defect from a weak agreement are high. This means the standard collusive strategies encountered in multi-unit second-price auctions do not hold in this context. By retaining the complete information setting proposed by Varian (2007) and Edelman et al. (2007), we eliminate the issue of private information extraction encountered in the literature on bidding rings as the *agency* has full knowledge of each member's private valuations[8]. Hence, incentive compatibility relies on finding a suitable way to implement collusive agreements that are robust to *ex-post* defections (i.e., once a contract has been settled). The collusive structure is based on two properties: (*i*) it is grounded on an allocation rule that uniformly divides collusive *spoils* and (*ii*) the *agency* implements a flat fee-based compensation policy that has to be incentive compatible.

The well-known multiplicity of — asymmetric and symmetric — non-cooperative equilibria (see Börgers et al. (2013)) makes it difficult to study collusion in the *GSP*. Collusive bids depend, by way of individual contributions and members' outside options, and on non-cooperative outcomes so that the collusive game also supports multiple collusive equilibria. We find that the *agency* (*i*) coordinates advertisers on a unique collusive equilibrium which entails surpluses that are beyond the convex hull of *Nash* equilibrium payoffs (*low-price* equilibrium) and (*ii*) extracts an optimal fee

[7]Various strategies can be adopted. One might include a bid rotation mechanism so that the agency assigns different advertisers to different keywords and rotates these assignments at each auction. Another might use "split award" agreements and hold as fixed the set of auctioned keywords for which advertisers compete, or not so that each keeps its own bundle of keywords. The first of these strategies may be technically costly to implement in a complex and dynamic environment, while the second raises the issue of which advertiser is more relevant to which keyword and how such a choice can be justified on economic grounds. We find the strategy of maintaining advertiser competition for the same set of keywords to be technically the easiest way of achieving efficient coordination and the most suitable for the specific multi-object environment considered here.

[8]See, for instance, Graham and Marshall (1987), McAfee and McMillan (1992), Mailath and Zemsky (1991), Marshall and Marx (2007) and Lopomo et al. (2010).

that secures a minimum revenue for the seller without weakening incentives to collude. Strategic plays are a function of the commission fee, which acts as a cursor: an increase in the fee drives up collusive bids (and prices), which converge to the bids that sustain the smallest Nash equilibrium, at which allocative efficiency can no longer be guaranteed. As a result, a bid delegation organised by an absolute rent-seeking *agency* rules out the multiplicity of equilibria and the corresponding fee implies an upper bound to the seller's revenue. Therefore, in line with Bulow and Klemperer (1996), it can be argued that in the short run, the incentives for a seller to break collusion, if it has a long-term perspective and is focusing its attention on capturing market shares, are mitigated in that it might be achieved at the expense of allocative efficiency and the advertisers' surplus. Finally, a well-known result forwarded by Varian (2007) and Edelman et al. (2007) is that the *GSP* auction has a *locally envy-free* (*LEF*) equilibrium that is *Vickrey-Clarke-Groves* (VCG)-equivalent[9]. We find that if incentives to defect are consistent with the *LEF* refinement proposed by Varian (2007) and Edelman et al. (2007), considered as a stability criterion, the model uniquely selects the *VCG* outcome, independently of the fee level. In the model, the *agency* picks the prices that maximise bidders' individual profits in their names. When contemplating some different positions (and so cheating on the agreements), taking each position price as given, we find that advertisers do not have sufficient incentives to reduce their expressed demand. Coordination uniquely selects the *VCG* outcome and the *GSP* auction becomes collusion-proof.

To the best of our knowledge, only two studies can be directly related to our work. Both examine the decentralisation of advertisers' bidding strategies to a third party in sponsored search auctions, but in contrast to this paper they consider either different goals or different settings. Ashlagi et al. (2009) were the first to study a position auction game in which an incentiveless third-party manages to coordinate the bidding strategies of a pool of advertisers by means of the solution concept of a *mediated equilibrium* as first proposed by Monderer and Tennenholtz (2009). The focus of Ashlagi et al. (2009) was then to study the ability of a third-party to implement the *VCG* outcome with incomplete information. Decarolis et al. (2017) also consider bid delegation

[9] The *locally envy-free* equilibrium refinement, as proposed by Varian (2007) and Edelman et al. (2007) and on which the equivalence relies, states that bidders should choose the position that maximises their profits taking prices as given. In equilibrium, no bidder should find it profitable to swap his position with that of a bidder allocated just above them.

to a marketing *agency* that bids on behalf of different advertisers in the same keyword auction. In their model, the *agency* is incentiveless and acts as a social planner whose main goal is to maximise the advertisers' surplus.

Here, however, we consider an additional aspect in the formation of agreements between advertisers and the *agency*: namely, that relations are contractually based and so involve an exclusivity rule and a commission fee policy. To account for this, we assume that the *agency* provides a monetary payback to advertisers so as to align incentives for reaching an agreement. Hence, we consider a strategic rent-seeking *agency*, which while also operating on behalf of each advertiser, uniformly allocates spoils from coordination at the end of the auction. In our model the *agency* seeks to maximise both the advertisers' surplus and its own profits, which raises the question as to what payment scheme the *agency* should adopt. In line with common practice in this industry, we suppose the *agency* extracts a management fee for intermediation[10]. The model assumes that the *agency* implements a common flat-fee on collusive spoils in order to maximise its own revenue[11]. We are thus able to account for the effects of the *agency*'s strategic behaviour on incentives to sustain the collusive agreement and on the seller's revenue, an explanation that is absent in Decarolis et al. (2017). We can now derive the bid profiles to be implemented and an analytic representation of how the seller's revenue reacts to a the *agency*'s rent-seeking behaviour.

Since third-parties operate along with outside advertisers, coordinated and non-cooperative bids interact, which may lead to tractability issues. Decarolis et al. (2017) manage to integrate this interaction by restricting the behaviour of outside bidders to that of the locally envy-free refinement of Varian (2007) and Edelman et al. (2007). To avoid these difficulties, we consider that a monopolistic *agency* is present on the market and assume an all-inclusive collusion.[12] Such an environment is optimal as the *agency* only has to control its members' bids and does not have to anticipate the bids of potential outside advertisers. We implicitly assume here that if a profitable,

[10]Thus, an advertiser can be charged each time there is a click on the ad being managed; or the *agency* can base a mark-up on the entire amount spent (budget) on the ad; or it can also require a mark-up toward the profits that the advertiser makes on each purchase. Another interesting but more complex system, is to *pay for performance*. Here, a metric is defined *ex-ante* (in somewhat fuzzy way) and the advertiser only pays once a sale is made or by employing another metric resulting from the management process.

[11]The 16th Edition of the Association of National Advertisers' (ANA) report *Trends in Agency Compensation Survey* (2013), documents that over 80% of 98 respondents claimed to uses a fixed-fee compensation policy.

[12]This assumption is motivated by empirical observations that during the most important keyword auctions, advertisers are mostly represented by just one agency, see Decarolis et al. (2017) and Decarolis and Rovigatti (2017).

efficient collusion cannot be implemented under the all-inclusive assumption, we can reasonably expect it not to be the case if we allow for incomplete cartels.

Finally, there is a recent strand of literature, including studies by Feldman et al. (2010) and Balseiro and Candogan (2017), that focuses on bid delegation to *agencies* in the context of multi-sided ad-exchange platforms, the second main segment of the advertising market. In this market, advertisers seek to attract the attention of internet users by placing advertisements in an ad space provided by a publisher on a webpage. Ad-exchanges bring publishers and advertisers together in the same marketplace in which an ad space is sold using a single-unit Vickrey auction[13]. In practice, advertisers belong to a network handled by an *agency* that buys ad space on their behalf directly on the exchange platform. In contrast to our model, in addition to the fact that we consider a multi-unit auction, these studies analyse delegation to an *agency* that has to select only one bid from its portfolio of advertisers and does not make side payments.

The rest of this paper is organised as follows: section 2 describes the model, explains the main issues facing the *bidding agency* and presents the *non-cooperative* benchmark. Section 3 presents the main results and section 4 concludes.

2 A position auction with a *bidding agency*

We model the *GSP* auction as a game involving an all-inclusive cartel (or *ring*) of three bidders in a set $\mathcal{N} = \{1,2,3\}$ who compete for two positions in a set $\mathcal{K} = \{1,2\}$ sold simultaneously by a search engine. A third-party, taking the role of what we call the *bidding agency*, coordinates the collusion. Each position k has an associated, commonly known, expected *click-through rate* (henceforth *ctr*) denoted by $\alpha_k \geq 0$ with $\alpha_1 > \alpha_2 > \alpha_3 = 0$.

Advertisers A player's valuation expresses his willingness to pay for a click and is denoted by x_i, with $x_i \in [0, \bar{x}]$. Let $\mathbf{x} = \{x_1, x_2, x_3\}$ be the set of individual valuations, labelled in decreasing order so that $x_1 > x_2 > x_3$. Valuations are understood to be independent of the positions, the identities of

[13] See McAfee (2011) and Muthukrishnan (2009) for an approach to the practical design of this market.

the advertisers allocated and of the customers' clicking behaviour[14]. Henceforth, i's value of being allocated in position 1 or 2 is defined by the product $\alpha_k x_i$. Each player simultaneously submits a single-dimensional and non-negative bid $b_i \geq 0$ for a click as a function of his valuation. To distinguish between a *non-cooperative* (or *Nash*) bid and a *collusive* bid we denote the former by b_i and the latter by $b_i^{\mathcal{N}}$. Then, let $\mathbf{b} = \{b_1, b_2, b_3\}$ be the set of bids and $\mathbf{b}_{-i}$ the set of bids excluding b_i and let $\mathbf{b}^{\mathcal{N}} = \left\{b_1^{\mathcal{N}}, b_2^{\mathcal{N}}, b_3^{\mathcal{N}}\right\}$ be the set of possible actions for cartel members, with $\mathbf{b}_{-i}^{\mathcal{N}}$ defined in a similar fashion. The game consists of allocating positions to bidders based on the order of their bids (or *expressed demand*). According to the auction's rules, when bidder i is assigned a position k he is charged a price *per* click $p_k = \alpha_k b_{k+1}$. Thus, the net payoff utility function of player i when assigned position k is equal to $\pi_i = \alpha_k (x_i - b_{k+1})$.

Consider that bidders meet before the main auction starts. If they attempt to outbid each other during this auction they will give most of the surplus to the seller. The task then is to reach an agreement whereby they can limit surplus extraction and expropriate a significant share of that surplus from the seller. The goal of the *bidding agency* is to coordinate matters in such a way that no bidder finds it profitable to break the settlement by reverting to competitive behaviour. The goal, therefore, is to draw up individual contracts $\boldsymbol{\gamma} = \left(\gamma_1, \gamma_2, \gamma_3\right)$ so that each potential member is better off than they would be in the absence of coordination. The contract comprises a bid recommendation, the side payments that are made and the flat-fee imposed by the *agency*.

Side payments We assume that the sustainability of the cartel (or *ring*) is based on the presence of side payments. This takes the form of a monetary transfer ω_k, corresponding to a bidder's contribution, from each member to the *bidding agency*. In return, the *bidding agency* makes a lump-sum transfer τ to each member comprising the collusive benefits[15].

Remark 1. *Individual contributions are based on the final allocation that is anticipated and incurred as a* sunk cost *by members.*

[14] We do not consider the question of the allocative externality generated, for instance, by a firm's reputation among its customers. A firms of high repute might imply more clicks for a firm of low repute if the former is placed at a position just above the latter. More experienced customers appear to click more carefully and also more frequently on ads placed in the median position. The quality of the ad and a firm's reputation affect their choices. Then valuations for clicks could be allowed to vary non-linearly between positions.

[15] This is done at the end of the main auction so that the *bidding agency* holds the bargaining power and necessarily controls the incentive compatibility constraints.

Contributions are computed as the difference between the price bidder i would have to pay to the seller in the absence of a *cartel* and the price he actually pays in the presence of the *cartel*. Let $\hat{p}_k$ be the payment made by member i to the seller when assigned position k in the presence of collusion. Then, $\omega_k = \max\{\alpha_k(p_k - \hat{p}_k); 0\}$. The expected total gain $\Pi_{\mathcal{N}}$ from colluding (referred to as *spoils*), and which is to be *uniformly* redistributed at the end of the auction, is the sum of each member's payment to the *bidding agency*, $\Pi_{\mathcal{N}} = \sum_{k \in \mathcal{K}} \omega_k$. At the end of the auction, the *bidding agency* exerts a credible monopoly power by setting a fee $\varepsilon \in [0,1]$ upon $\Pi_{\mathcal{N}}$. It may (i) act as a risk-neutral, incentiveless agent when $\varepsilon = 0$ or (ii) levy a fee $0 < \varepsilon < 1$ upon collusive profits, or (iii) keep the entire collusive gain with $\varepsilon = 1$ and so substitute the seller. The quantity ε is set *non-strategically* and each member receives a transfer τ equals to:

$$\tau = \frac{(1-\varepsilon)}{3} \sum_{k \in \mathcal{K}} \omega_k = \frac{1-\varepsilon}{3} \sum_{k \in \mathcal{K}} \alpha_k \left(b_{k+1} - b^{\mathcal{N}}_{k+1} \right) \tag{1}$$

We call this quantity a *uniform-ε redistribution rule*. It is assumed as being retained by the *bidding agency* if a defection occurs. Note that τ is independent of i's own allocation. Given the redistribution scheme, we denote by $p_i^{\mathcal{N}}$ the total payment made by player i to the cartel:

$$p_i^{\mathcal{N}} = \begin{cases} \omega_k - \tau & \text{if } k = \{1,2\} \\ -\tau & \text{if } k = \{\emptyset\} \end{cases} \tag{2}$$

that is, bidder i if assigned to position $k = \{1,2\}$ pays his individual contribution to the cartel and receives his individual share from the *bidding agency* in return, whereas if $k = \{\emptyset\}$, he receives his individual share but does not contribute[16]. As a result, member i's expected gain assigned to position k equals:

$$\pi_i^{\mathcal{N}} = \begin{cases} \alpha_k \left(x_i - b^{\mathcal{N}}_{k+1} \right) - p_i^{\mathcal{N}} & \text{if } k = \{1,2\} \\ -\tau & \text{if } k = \{\emptyset\} \end{cases}$$

[16] The environment of complete information simplifies greatly the functioning of such a redistribution rule as there is no need to implement a *pre-knockout* in order to make players reveal their private willingness to pay for a click. There is no adverse selection in this framework and no issue of cartel misrepresentation at the main auction.

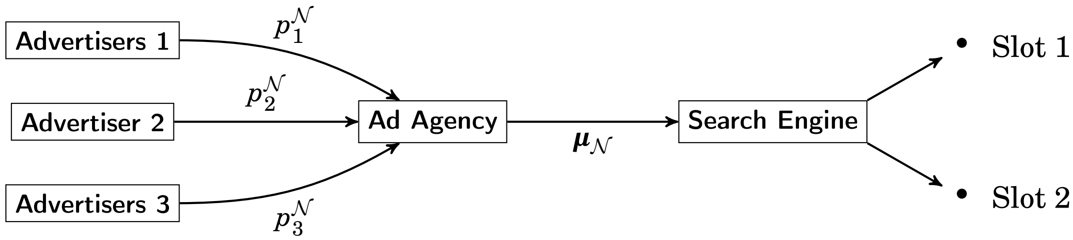

Figure 1: Market structure

Agency problem The *agency*'s profits are given by $\Gamma_{\mathcal{N}} = \varepsilon \sum_{k \in \mathcal{K}} \omega_k$. It offers individual contracts γ_i, to which it is committed, composed of a system of recommended non-negative bids $\boldsymbol{\mu}_{\mathcal{N}} = \left(b_i^{\mathcal{N}}\right)_{i \in \mathcal{N}}$ to be placed during the auction, the share $\varepsilon \in [0,1]$ it intends to keep and the side-payments required from each member $\mathbf{p}^{\mathcal{N}} = \left(p_i^{\mathcal{N}}\right)_{i \in \mathcal{N}}$. Figure 1 illustrates the market structure we consider. To illustrate the idea, the timing of the game can be described as follows:

1. Before the auction, the *agency* makes a contract proposal $\gamma_i = \left(\varepsilon, \left\{b_i^{\mathcal{N}}, p_i^{\mathcal{N}}\right\}\right)$.

2. In this *participation phase*, each member either accepts or rejects the proposal. If a member rejects the proposal then no delegation takes place and each member plays non-cooperatively; otherwise, bid delegation comes in to operation.

3. Members are then asked to pay their contributions ω_i and to refrain from bidding in their own name at the main auction. At this *deviation phase*, they can defect from the agreement and bid some $\tilde{b}_i$ in order to win a position they could not have won if all were to play non-cooperatively.

4. The seller then allocates the bidders in decreasing order of their bids and each is charged $\hat{p}_k$.

5. If no defection occurs during the auction stage, the *bidding agency* redistributes τ and nothing else.

The *agency* maximises the advertisers' total surplus, subject to incentive compatibility, $\forall i = \{1,2,3\}$, $\forall b_i^{\mathcal{N}} \in \mathbf{b}^{\mathcal{N}}$ and $\forall x_i \in \mathbf{X}$:

$$\pi_i^{\mathcal{N}}\left(b_i^{\mathcal{N}}, \mathbf{b}_{-i}^{\mathcal{N}}, x_i\right) - \tilde{\pi}_i\left(\tilde{b}_i, \mathbf{b}_{-i}^{\mathcal{N}}, x_i\right) \quad \geq \quad 0 \tag{3}$$

in which $\tilde{\pi}_i\left(\tilde{b}_i, \mathbf{b}_{-i}^{\mathcal{N}}, x_i\right)$ represents player i's surplus if he exploits the collusion with a bid $\tilde{b}_i \neq b_i^{\mathcal{N}}$ while the others play according to the strategy $b^{\mathcal{N}}$. Profits from upward and downward deviations take the following formulation:

$$\forall l > k : \tilde{\pi}_i\left(\tilde{b}_i, .\right) = \alpha_l\left(x_i - \max\left\{b_{l+1}^{\mathcal{N}}, 0\right\}\right) - \omega_k \tag{4}$$

$$\forall s < k : \tilde{\pi}_i\left(\tilde{b}_i, .\right) = \alpha_s\left(x_i - \max\left\{b_s^{\mathcal{N}}, 0\right\}\right) - \omega_k \tag{5}$$

and subject to individual rationality:

$$\pi_i^{\mathcal{N}}\left(b_i^{\mathcal{N}}, \mathbf{b}_{-i}^{\mathcal{N}}, x_i\right) \geq \pi_i\left(b_i, \mathbf{b}_{-i}, x_i\right) \tag{6}$$

and the restriction that $b_1^{\mathcal{N}} \geq b_2^{\mathcal{N}} \geq b_3^{\mathcal{N}}$.

The incentive compatibility constraint (3) expresses the idea that once bidder i receives his recommended bid (prior to playing in the auction), he has the choice of either obeying or adjusting his bid against the bids the *bidding agency* places on behalf of his competitors[17]. If a deviation occurs, the defector does not receive his individual compensation τ but still incurs his contribution as a *sunk cost*. This constraint requires that there is no deviating bid $\tilde{b}$ that allocates bidder i any different position so that he is better off. The last constraint (6) means that player i has to find it profitable to join the cartel. His profits from accepting contract γ_i are at least as high as his outside option, which corresponds to his profits in the *Nash* game so that each outside option is assumed to be an equilibrium outcome.

Equilibrium criterion

Definition 1. *A collusive mechanism $\boldsymbol{\zeta} = \left(\boldsymbol{\mu}_{\mathcal{N}}, \mathbf{p}^{\mathcal{N}}\right)$ is an equilibrium profile if it is (i) individually rational and (ii) incentive compatible.*

In single unit auctions, a collusive device is said to implement an efficient outcome if the highest-valuing member of the collusion represents the cartel, wins the object and the cartel is able to suppress internal competition. Such an outcome naturally maximises social welfare as

[17] Assuming they behave according to the recommendation.

the good is allocated to the highest-valuing player. We modify this definition slightly to make it coherent with the *GSP* auction context.

Definition 2 (*First-best outcome*)**.** *A collusion in the* GSP *auction is an efficient mechanism if:* (i) *the final allocation is such that only the two-highest valuing members are active during the targeting auction,*(ii) *it suppresses competition from bidders not allocated to any position*, (iii) *it maximises social welfare.*

2.1 Benchmark — Non-cooperative outcome

In line with Varian (2007) and Edelman et al. (2007), we first give the set of *Nash* equilibria in which bidders do not collude. Renumber players so that x_k is the valuation of the bidder assigned to position k. An equilibrium outcome will be supported by some bid profile $\mathbf{b} = (b_1, b_2, b_3)$ if under it (i) no player can profitably deviate from the position to which he is assigned in equilibrium to any lower or higher positions; that is, $\alpha_k(x_k - p_k) \geq \alpha_l(x_k - p_{l-1})$ for $l < k$ and that $\alpha_k(x_k - p_k) \geq \alpha_s(x_k - p_{s+1})$ for $s > k$ and $p_k = \alpha_k b_{k+1}$; (ii) an assigned player cannot profitably deviate to win any position; that is, $\alpha_k(x_k - p_k) \geq 0$ [18].

Uncoordinated bids: The following strategy profile $\mathbf{b} = (b_1, b_2, b_3)$ characterises all the *Nash* equilibria of the static *GSP* auction with complete information:

$$\begin{aligned}
b_1 &\in \left[\max\left\{x_2 - \frac{\alpha_2}{\alpha_1}(x_2 - b_3), x_3\right\}; \bar{x}\right] \\
b_2 &\in \left[\max\{x_3, b_3\}; x_1 - \frac{\alpha_2}{\alpha_1}(x_1 - b_3)\right] \\
b_3 &\in \left[\max\left\{0, x_1 - \frac{\alpha_1}{\alpha_2}(x_1 - b_2)\right\}; \min\{x_1, x_2\}\right]
\end{aligned} \tag{7}$$

Each player can envision a multiplicity of best-responses to each other's equilibrium strategy. Equilibrium bids are bounded above and below by a combination of the lowest bidder's bid and the valuations of the player above and below him. Two things are worth noticing about bid profile $\mathbf{b}$.

[18]Note that the shape of prices is modified as upward deviation occurs, which is not in the spirit of usual competitive equilibrium analysis. Players can influence the price they pay at the end and these prices are not taken as given. Each bidder when contemplating an upper position does not expect to pay the same price as that already assigned to that position.

First, it does not rule out inefficient and non-assortative allocations. In fact, the only allocation that is ruled out is that in which the highest-valuing advertiser wins no position and the lowest-valuing advertisers wins the top position[19]. Secondly, overbidding could be a candidate for a *Nash* equilibrium so that it does not restrict attention to undominated strategies.

VCG outcome: According to Varian (2007) and Edelman et al. (2007), a strategy profile $\mathbf{b} = (b_1, b_2, b_3)$ is a *LEF* or a *symmetric* equilibrium if it holds that $\alpha_k x_k - p_k \geq \alpha_{k-1} x_k - p_{k-1}$ for $k = 1,2,3$ and with $p_k = \alpha_k b_{k+1}$. This criterion implies a shift in the *IC* conditions so that upward and downward deviations become symmetric. Even if advertisers were to swap each other's position, the price structure would not change. If the bidder in position k decides to undercut the bidder assigned to a position $l < k$, he can expect to pay $p_l = b_{l+1}$ and not $p_{l-1} = b_l$. Thus, under the *LEF* stability criterion, each player chooses a position that maximises his payoff taking prices as given and the process reaches equilibrium prices that then clear the market. In order to implement a stable competitive equilibrium, bidders need to be *locally indifferent*. No player has an incentive to deviate and to bid for the position just above him. The player assigned to position k has to be *locally indifferent* between winning position $k-1$, paying his own bid and position k paying the next-highest bid. This implies that $\alpha_k x_k - p_k = \alpha_{k-1} x_k - p_{k-1}$, that is $p_k = \alpha_i b_{i+1} = \sum_{k=i+1}^{3} x_k (\alpha_{k-1} - \alpha_k)$, which gives the following set of bids $\mathbf{b}^v = \left(b_1^v, b_2^v, b_3^v\right)$:

$$
\begin{aligned}
& b_1 > b_2^v \\
& b_2 = x_2 - \frac{\alpha_2}{\alpha_1}\left(x_2 - b_3^v\right) \\
& b_3 = x_3
\end{aligned} \tag{8}
$$

The bid profile $\mathbf{b}^v$ corresponds to the lowest point in the set of competitive prices among all the *locally envy-free* equilibria. Even if truth-telling is not an equilibrium strategy of the GSP auction, the outcome is strictly equivalent to what a VCG mechanism would have implemented.

[19] Indeed, if the highest-valuing player was to win no position then it should be the case that $b_3 \geq x_1 \geq x_3$ which cannot be an equilibrium profile. If the lowest-valuing player was to now win the top position then it should be the case that $p_3 = b_1 > x_3$, implying a strict loss for him.

Notations: Two equilibria, which we refer to as *Lower-Nash Equilibrium* (LE) and *Upper-Nash Equilibrium* (UE), reside at the boundaries of the *Nash* set. Each profile is denoted respectively by $\mathbf{b}^l = (b_i^l)_{i=1,2,3}$ and $\mathbf{b}^u = (b_i^u)_{i=1,2,3}$. We call the equilibrium achieved when restricting b_3 as being equal to x_3 the *Dom-Nash* equilibrium denoted by $\mathbf{b}^d = (b_i^d)_{i=1,2,3}$ and denote the VCG outcome by $\mathbf{b}^v = (b_i^v)_{i=1,2,3}$.

From a revenue perspective, any *symmetric* equilibrium (or *locally envy-free* equilibrium) induces a revenue at least equal to the *VCG* revenue (e.g., Feldman et al. (2011), Lucier et al. (2012)). As a result, if we denote by R^l and R^u the auctioneer's revenues achieved by $\boldsymbol{b}^l$ and $\boldsymbol{b}^u$, respectively we have that $R^l \leq R^v \leq R^u$.

3 Equilibrium and revenue

3.1 Collusive equilibria

Consider that the *bidding agency* acts as a credible banker to which bidders entrust their expected cash surplus from colluding *efficiently*. The *agency* computes individual contributions conditional on the efficient *Nash* allocation deduced from the set (7). We assume that contributions are *allocation-independent,* computed *ex ante* and, thus, that they are not altered by potential deviations[20]. Let $\theta = \frac{\alpha_2}{\alpha_1} \in [0,1]$ and $\eta \in [0,1]$ denote, respectively, the ratio of clicks between both positions and some given threshold value.

Fact 1. *The more positions are substitutes the less robust collusive agreements are.*

At the benchmark, in order to remain optimal, *Nash* equilibrium bids are decreasing in the click ratio θ. If the substitutability between both positions becomes perfect, bidders should bid less as the need to outbid competitors to win the highest position decreases. However, when coordination is active, if the substitutability between positions increases, incentive compatibility constraints become more stringent. The low-valuing player's profit from cheating and aiming at second position increases with θ. To see this, recall that the incentive compatibility constraints

[20] The main concern is the inability of the *bidding agency* to provide sufficient control over the bidders' bidding behaviour and to deter *shill bidding* and *fake aliases*, two concerns inherent to online auctions.

for this member are given by $IC_3 : \tau \geq \alpha_1 \left(x_3 - b_1^{\mathcal{N}}\right)$ and $IC_3^{'} : \tau \geq \alpha_2 \left(x_3 - b_2^{\mathcal{N}}\right)$. The player's payoff from deviating to position one is independent of θ. However, his payoff is strictly increasing in θ in the case of a deviation to the second position. As a result, the *bidding agency* is under a constraint to release the second-highest valuing member in order to keep the incentives aligned and to maintain the ranking, which in turn makes the internal competition difficult to contain.

Consider contract $\gamma_i = \left(0, \{b_i^{\mathcal{N}}, p_i^{\mathcal{N}}\}\right)$, where the *agency* is perfectly benevolent. The *agency* maximises the advertisers' surplus and acts as a social planner by not extracting benefits from the collusive surplus. We can begin by stating that in the absence of commission fees, the *agency* implements the most subversive outcome in relation to the seller's revenue. It coordinates advertisers over the smallest price equilibrium, at which they capture the entire market rent.

Proposition 1. *For $\theta \geq \eta$, contract $\gamma_i = \left(0, \{b_i^{\mathcal{N}}, p_i^{\mathcal{N}}\}\right)$ maximises individual profits and implements an efficient outcome, whereby all the rent is captured by advertisers.*

Proof. The proof follows from that of proposition 2 setting $\varepsilon = 0$. □

As in the *non-cooperative* framework, *cartel* members can envision multiple best-replies in response to their respective equilibrium collusive strategies, which makes the study of coordination difficult. The underlying incentive compatibility constraints allow for symmetric/asymmetric equilibria (see set 1 in the Appendix). This proposition states that if the degree of substitutability between positions is sufficiently high, the *agency* coordinates advertisers in a first-best outcome, which rules out the multiplicity of collusive equilibria. In equilibrium, the rivalry between bidders is suppressed, thus blocking rent appropriation from the seller. It is fully captured by the advertisers. Individual payments are minimised and competition from the middle-valuing advertiser is just high enough to secure the second slot and to deter defections from the lowest-valuing advertiser.

Let us now assume a strategic *agency* that seeks to maximise its own revenue by charging advertisers a uniform commission fee, i.e., $\varepsilon > 0$. Consider the optimal contract $\gamma_i^* = \left(\varepsilon, \{b_i^{\mathcal{N}}, p_i^{\mathcal{N}}\}\right)$ such that bidders are allocated in decreasing order of their valuations with τ maximised given both the bid profile $\mathbf{b}^g = \left(b_i^g\right)_{i=1,2,3}$ and the fee ε and such that the corresponding profile $\boldsymbol{\mu}_{\mathcal{N}}^*$ defines the smallest efficient collusive equilibrium. The following proposition shows that a whole bundle of fees exists so that such a contract can be implemented by a strategic *agency*, and if the degree of substitutability between both positions is sufficiently high, the mechanism results in a constrained *low-price* outcome.

Proposition 2. *$\gamma_i^* = \left(\varepsilon, \{b_i^{\mathcal{N}}, p_i^{\mathcal{N}}\}\right)$ is an equilibrium contract. The equilibrium bid profile $\boldsymbol{\mu}_{\mathcal{N}} = \left(b_i^{\mathcal{N}}\right)_{i=1,2,3}$ is monotonically non-decreasing in ε and decreasing in $\mathbf{b}^g = \left(b_i^g\right)_{i=1,2,3}$ with $g = l, v, u$. There exist parameter values $(\theta, \eta, \delta) \in [0,1]^3$ such that $\forall \theta \geq \eta \wedge \forall \varepsilon \leq \delta$, a constrained first-best outcome is achieved and $\forall \varepsilon > \delta$ internal bidder's rivalry cannot be constricted.*

Proof. See Appendix A.1. □

In equilibrium, assigned members' bids are set equal to a quantity that is strictly below the valuation of the non-assigned member, who refrains from bidding. The profile of bids $\boldsymbol{\mu}_{\mathcal{N}}^* = \left(b_i^{\mathcal{N}}\right)_{i=1,2,3}$ is such that $\forall \varepsilon \leq \delta \wedge \theta \geq \eta$ equilibrium bids are characterised by:

$$\begin{aligned}
b_1^{\mathcal{N}} &= x_3 - \frac{(1-\varepsilon)}{3\alpha_1} \sum_{i=1}^{3} \alpha_i \left(b_{i+1} - b_{i+1}^{\mathcal{N}}\right) \\
b_2^{\mathcal{N}} &= \frac{1}{\lambda} \left(x_3 - \frac{(1-\varepsilon)}{3\alpha_2} \sum_{i=1}^{3} \alpha_i b_{i+1}\right) \\
b_3^{\mathcal{N}} &= 0
\end{aligned} \tag{9}$$

with $\lambda = \frac{1}{1-(1-\varepsilon)\frac{\alpha_1}{\alpha_2}}$, whereas $\forall \varepsilon \leq \delta \wedge \theta < \eta$ the game results in an inefficient allocation such that, $b_1^{\mathcal{N}} \in \left[b_2^{\mathcal{N}}, \bar{x}\right]$ and $b_2^{\mathcal{N}} = b_3^{\mathcal{N}} = 0$, and finally $\forall \varepsilon > \delta$ inner competition cannot be constricted. Equilibrium bids jump to the corresponding outside option levels and are defined by relations (13) and (14) in Appendix A.1.

Fact 2. *Collusive bids $b_i^{\mathcal{N}}$ negatively depend on the outside option and spoils increase with that option. As a result, the auction revenue is higher as players are assumed to bid low if playing* non-cooperatively[21].

Proposition 2 asserts that the bid delegation to the *bidding agency* takes the form of an efficient coordination according to definition 2. The highest-valuing player is allocated first position and

[21]Consider the social planner situation in which $\varepsilon = 0$ (the relation holds for $\varepsilon > 0$). Denote by $R^{\mathcal{N}}$ and R the auction revenue with and without coordination respectively. The seller's revenue under collusion drops to $R^{\mathcal{N}} = \frac{\alpha_1}{3\alpha_2 - \alpha_1} 3\alpha_2 x_3 - \frac{\alpha_1}{3\alpha_2 - \alpha_1} R$ and decreases with R.

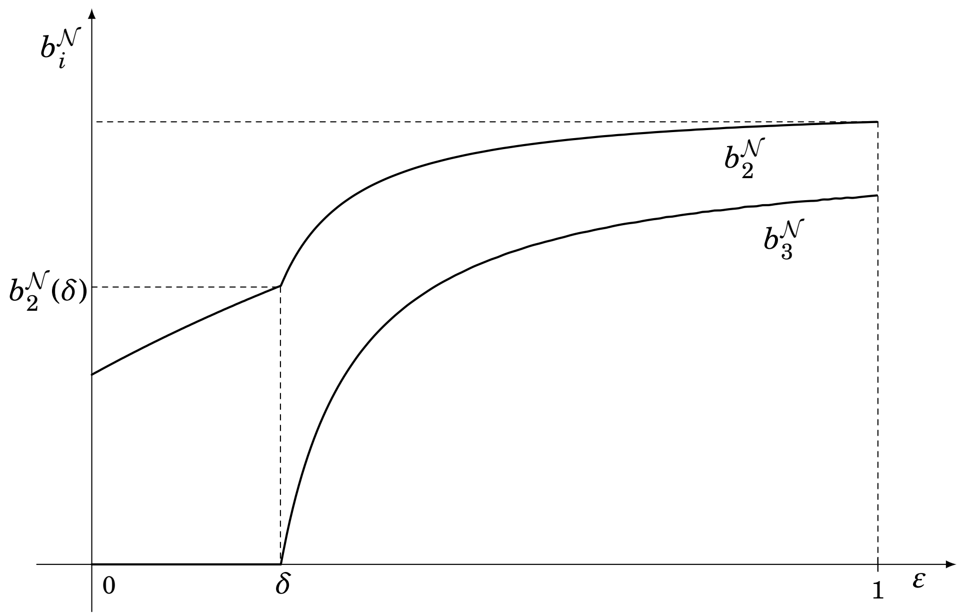

Figure 2: Collusive equilibrium bidding functions.

the second-valuing player the second, whereas the lowest-valuing player refrains from bidding. Hence, bid delegation stabilizes the auction outcome to a unique efficient allocation constrained by the fee level, which adversely affect the auctioneer's revenues.

Remark 2. *The click ratio constraints η and ρ are meaningless when the* bidding agency *sets $\varepsilon > \delta$. Indeed, the threshold couple (η, ρ) makes no sense since (i) relations (13) and (14) hold $\forall \varepsilon \in [0, \delta]$. Thus, the second-valuing member's bid is strictly positive once ε reaches the value δ and is, thus, also strictly positive for $\varepsilon = \delta + \sigma$. Then, (ii) $\forall \varepsilon \leq \delta$ if $\theta > \eta$ his bid is also strictly positive. Finally, once $\varepsilon > \delta$ we have that $\frac{\partial(b_2^{\mathcal{N}})}{\partial \varepsilon} \geq 0$, which implies that $b_2^{\mathcal{N}} > 0$ independently of the value of θ and that $b_3^{\mathcal{N}}$ is also necessarily positive. Therefore the thresholds η and ρ become unbinding $\forall \varepsilon > \delta$.*

Equilibrium bid functions are illustrated in Figure 2, in which we have set $b_1^{\mathcal{N}} = b_2^{\mathcal{N}} + \epsilon$ since affects neither payments nor revenues. The bid profile $\boldsymbol{\mu}_{\mathcal{N}} = (b_2^{\mathcal{N}}, b_3^{\mathcal{N}})$ increases in the fee and both functions are non-differentiable at threshold δ from which within-competition strictly increases. The lowest-valuing member's bid becomes strictly positive. The non-differentiability at threshold δ of the equilibrium bidding functions is a straight implication of two correlated effects. (i) Once the *bidding agency* decides to set its fee ε at a higher level, bidding functions consequently increase, which mechanically involves a decrease in the available surplus. In return, (ii) the *bidding agency* needs to break the non-assigned member's incentives to defect, which implies an increase in the overall collusive bid functions. To see why, recall that individual payoff utility function is equal to:

$$\pi_i^{\mathcal{N}}\left(x_i;b_i^{\mathcal{N}}\right) = \alpha_k\left(x_i-b_{k+1}^{\mathcal{N}}\right)-\left(\alpha_k b_{k+1}-\alpha_k b_{k+1}^{\mathcal{N}}\right)+\frac{1-\varepsilon}{3}\sum_{j=1}^{2}\alpha_j b_{j+1}-\frac{1-\varepsilon}{3}\sum_{j=1}^{2}\alpha_j b_{j+1}^{\mathcal{N}}$$

it is thus obvious to see that $\frac{\partial^2\pi_i}{\partial\varepsilon\partial b_i^{\mathcal{N}}}>0$.

Fact 3. *The marginal profit loss due to an increase in the fee is compensated by a marginal increase in the coordinated bids, and this compensation is more likely to increase as coordinated bids increase.*

Thus, greed on the part of the *bidding agency* generate more aggressive behaviour from members in order to maintain incentive aligned and to compensate the downward shifting of the collusive gain in terms of individual contributions. These two correlated effects corroborate the intuitive idea that the monopoly power of the *bidding agency* has a substantial pervasive effect upon the durability and sustainability of the collusion.

Therefore, proposition 2 also relates to the sustainability of the *non-cooperative* outcome insofar as it represents a limit case for any collusive mechanism. We know from relation (7) that the complexity of the *GSP* auction implies a multiplicity of *Nash* equilibria. In our framework, the *agency* has the power to push the equilibrium prices upward so that bidders reverse to *non-cooperative* behaviour. Consider that the *bidding agency* endorses the role of an exogenous equilibrium ”perturbator” (as a mediator in Ashlagi et al. (2009), Monderer and Tennenholtz (2009)), up to which point can this be done?

Corollary 1. *The* non-cooperative *equilibrium is a sustainable collusive outcome with a uniform-ε redistribution rule. As $\varepsilon\mapsto 1$ the* bidding agency *coordinates bidders over the lowest* Nash *equilibrium bid profile $\mathbf{b}^l=\left(b_i^l\right)_{i=1,2,3}$.*

Proof. The result is immediate if we set $\varepsilon=1$ into the equilibrium collusive bidding functions defined in relations (13) and (14). We obtain $b_2^{\mathcal{N}}=x_3=b_2^L$ and $b_3^{\mathcal{N}}=x_1-\frac{\alpha_1}{\alpha_2}(x_1-x_3)=b_3^L$. The objective of the *agency* is thus now to maximise the function $SP=\sum_{i=1}^{3}(\alpha_i(x_i-b_{i+1}))$ under the same *IC* constraints of the *non-cooperative* benchmark. This entails $\mathbf{b}^l$ as a natural equilibrium outcome. $\square$

The simplest *IC* collusive mechanism that is *always feasible* turns out to be the *non-cooperative* mechanism in which players set their optimal bids consistent with the *Nash* equilibrium criterion. Obviously, collusive payoffs are no lower than in their corresponding *non-cooperative* counterpart. However, by incrementally increasing its expropriation abilities, the *agency* is able to evict all other equilibria, so that the *lowest Nash* equilibrium outcome $\mathbf{b}^l = \left(b_i^l\right)_{i=1,2,3}$ becomes the unique ending point. The intuition of the corollary is as follows. The *bidding agency* acts on behalf of each member and if it acts according to proposition 2 no bidder will finds it profitable to leave the cartel. Then, if $\varepsilon = 1$, it can bid in two different ways. It can pick any bid vector compatible with the *Nash* equilibrium criterion so that it implements any *Nash* outcome without being interested in the bidders' individual welfare. Otherwise, it can decide to maximise the latter. By doing so, it picks the price vector that maximises each individual's profits, which is $\mathbf{b}^l = \left(b_i^l\right)_{i=1,2,3}$. This implies that in the limit case, individual payoffs from coordination are equal to the corresponding equilibrium payoffs with $\mathbf{b}^l$.

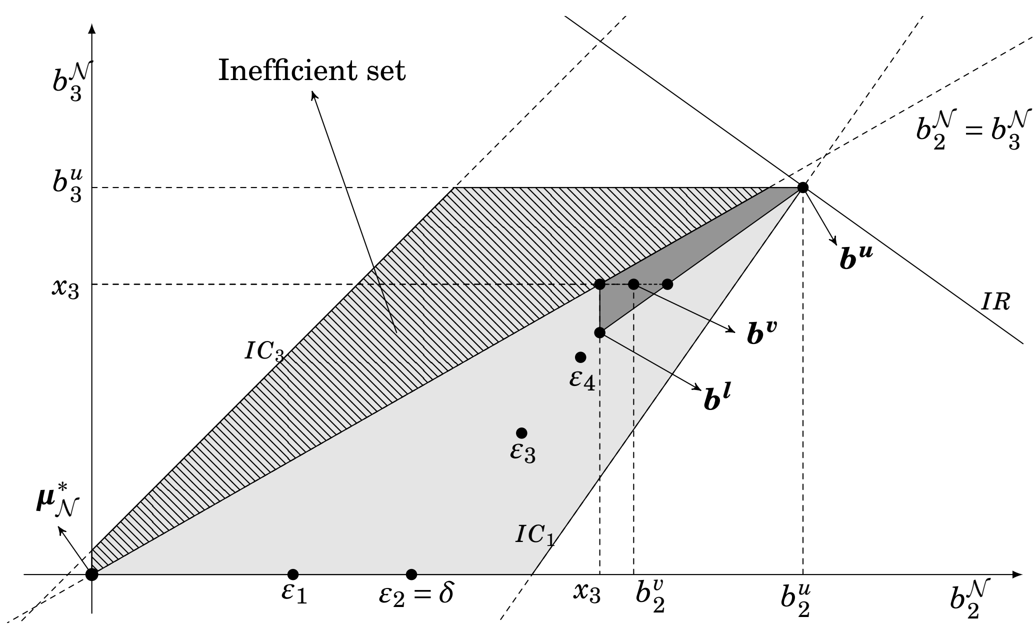

Figure 3: Set of collusive bids in light grey with $\varepsilon_0 = 0, \varepsilon_1 = 0.3, \varepsilon_2 = \delta, \varepsilon_3 = 0.6$ and $\varepsilon_4 = 0.9$.

Example 1. *Let the set of valuations be* $\boldsymbol{x} = (5,4,3)$ *and the value of both click-through rates be* $\boldsymbol{\alpha} = (10,8)$. *Suppose that as an outside option, the* bidding agency *conjectures that bidders will* *play according to the bid profile* $\boldsymbol{b}^u = (b_1^u, b_2^u, b_3^u)$*, defined as the upper bound of each interval of equilibrium bids in* (7)*. This gives* $b_3^u = 4, b_2^u = 4.2$ *and* $b_1^u = 5$ *and the corresponding collusive set is pictured in Figure 3. In Figure 3, we represent equilibrium outcomes for different fee level, ranging from a pure benevolent* agency *at* $\varepsilon = 0$ *to an* agency *that is an "absolute rent seeker" at* $\varepsilon = 1$*. In this environment* $\delta \simeq 0.51$ *and* $\eta = 0.84$ *so* $\theta < \eta$*. When the* agency *does not seek to extract rent the equilibrium proposal* $\boldsymbol{\mu}^*_{\mathcal{N}} = (b_1^{\mathcal{N}}, b_2^{\mathcal{N}}, b_3^{\mathcal{N}})$ *for* $\varepsilon = 0$ *is given by* $b_3^{\mathcal{N}} = 0, b_2^{\mathcal{N}} = 0$ *and* $b_1^{\mathcal{N}} > 0$*. The mechanism means both players 2 and 3 refrain from bidding, which results in a suboptimal allocation (with both players being randomly assigned position two). This constitutes the worst scenario for the seller as the* non-cooperative *game results in a revenue of* $R^u = \alpha_1 b_2 + \alpha_2 b_3 = 74$*, whereas under bid coordination, revenues are* $R_{\boldsymbol{\mu}_{\mathcal{N}}} = 0$*. This corresponds to the standard* low-revenue property *encountered in multi-object auctions with active collusion. As shown by the corollary, the collusive outcome converges to* $\boldsymbol{b}^l$ *as* $\varepsilon \mapsto 1$*. At this point,* $\boldsymbol{\mu}_{\mathcal{N}} = (b_1^{\mathcal{N}}, b_2^{\mathcal{N}}, b_3^{\mathcal{N}}) = (b_2^{\mathcal{N}} + \varepsilon, 3, 2.5)$*, generating a revenue* $R_{\boldsymbol{\mu}_{\mathcal{N}}} = 50$*. Consider that the outside bid profile is set to* $\boldsymbol{b}^l = \left(b_1 > b_2, x_3, x_1 - \frac{\alpha_1}{\alpha_2}(x_1 - x_3)\right)$*, then* $R^l = \alpha_1 b_2 + \alpha_2 b_3 = 50$ *and* $R_{\boldsymbol{\mu}_{\mathcal{N}}} = 50$ *if* $\varepsilon = 1$*.*

This last corollary sheds light on the fact that an efficient collusion implemented by a strategic *agency* cannot result in a full rent expropriation. Greedy behaviour strictly increases within competition, and as the results in the next subsection show, maximising the commission fee, i.e., setting an optimal fee level, secures a minimum level of rent for the seller that mitigates the adverse effect of bid delegation on the auction revenues.

3.2 The auctioneer and the *agency* revenue

A profitable coordination can then be achieved even if the *bidding agency* increases its monopoly power over the *cartel* members. The *agency* still keeps track of the payoff dominance of the collusion over the *non-cooperative* game and secures allocative efficiency. However, introducing the ability to extract a positive fee creates a tension between expropriation of surplus, constriction of bidders' rivalry and a *low-price* outcome. A higher fee level increases the *agency*'s revenue but drives up equilibrium prices, i.e., advertisers rivalry, and lowers joint profits. Both objectives are aligned in opposite ways and a rent-seeking behaviour introduces a trade-off that the *bidding* *agency* has to cope with. We begin by characterising the optimal fee level that solves the tension between surplus expropriation and the constriction of rivalry.

Lemma 1. *An optimal threshold δ is set equal to the point at which the lowest member is indifferent to being active or staying inactive. It is a non-decreasing function of the* Nash *bid profile* $\boldsymbol{b}^g = \left(b_i^g\right)_{i=1,2,3}$ *with* $g = l, v, u$.

The threshold implies a more stringent condition over individual contributions ω_i if the *Nash* bids are assumed to be played according to $\mathbf{b}^l$. A small increase in the fee makes the low-valuing member sets $b_3^{\mathcal{N}} > 0$, which increases within-competition. The *IC* constraints become binding and members' incentives are reversed. As a result, the *first-best outcome* property can no longer be maintained if the *bidding agency* were to become too greedy with respect to the surplus extraction. This causes the collusion to be broken down from the inside and the *low-price* property to vanish. This lemma, together with corollary 1, provides the following result:

Proposition 3. *The* bidding agency*'s revenue is non-monotonic in ε. An optimal incentive-compatible fee at which an* agency *maximises its revenue is set equal to the point $\varepsilon^* = \delta$.*

Proof. See Appendix A.2. □

The profit functions are depicted in Figure 4. Although bidding functions are monotonically increasing, the *bidding agency* is able to increase its profits. The increase in fees sufficiently compensates for the increase in collusive prices. As a result, biasing the redistribution of the surplus becomes a *plausible and valuable* strategy as it can be achieved without deterring the collusive incentives. The threshold δ is the critical point at which ε becomes unsustainable and from which a greedy attitude becomes pervasive. The loss in the bidders' surplus needs to be compensated for by an increase in the respective collusive bids, which implies a shift in the *bidding agency*'s profits. A rational behaviour that mechanically implies an increase in the internal competition level, so that collusion can no longer be constrained to an efficient level according to definition 2.

In fact, the latter is, in many ways, a "common sense" observation as anyone could have expected this aggressive result. The level of expropriation from the *bidding agency* makes the incentive conditions more stringent for the bidders, preventing them from showing any self-interest

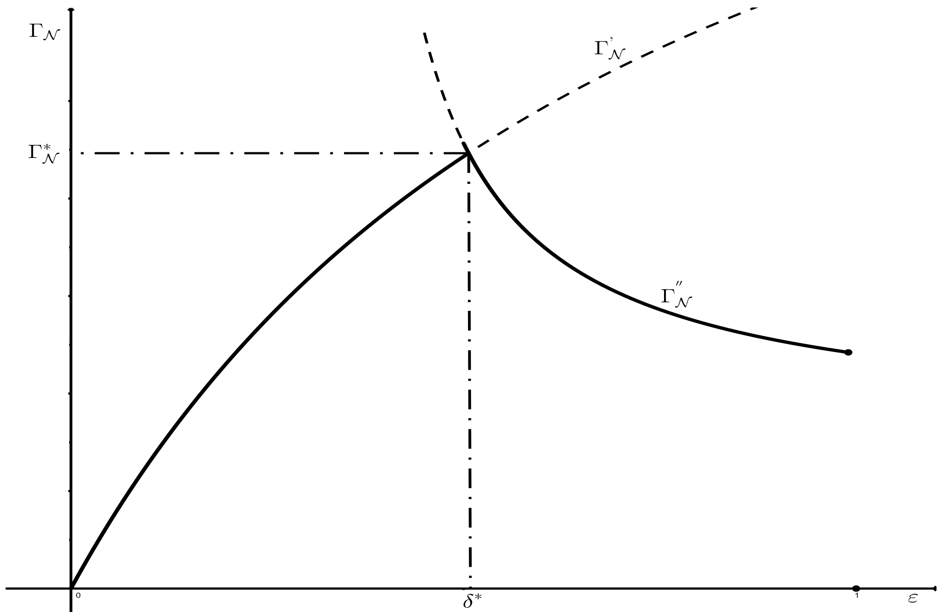

Figure 4: The *bidding agency*'s profits as a function of ε

in the cartel formation. Interestingly, such greedy behaviour from the *bidding agency* does not lead to antinomy with the auctioneer and may be beneficial to him as $\boldsymbol{\mu}_{\mathcal{N}} = \left(b_i^{\mathcal{N}}\right)_{i=1,2,3}$ is strictly increasing in ε.

Consider now that the auctioneer has the power to affect the click-through rate couple $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$. This is a strong assumption as the α_is are the result of the customers' search strategies. But let us abstract this feature away and let us suppose that, indeed, the auctioneer was to set them arbitrarily. We can therefore make the following claim:

Claim 1. *An auctioneer cannot use the ratio θ in order to induce levels of revenue higher than those of the lowest* Nash *equilibrium.*

Proof. See Appendix A.3. □

Together with the fact that $\lim_{\varepsilon \mapsto 1} b_i^{\mathcal{N}} = b_i^l$, proposition 2, which says that the bid profile $\boldsymbol{\mu}_{\mathcal{N}}$ is decreasing in the bid profile $\mathbf{b}$ and proposition 3, which says that it is optimal for the *bidding agency* to set $\varepsilon = \delta$, we can state the following:

Proposition 4. *Greedy behaviour on the part of the* agency *benefits the auctioneer. The auctioneer's revenue is monotonically increasing in ε and the maximum quantity of surplus it can extract is no higher than in the corresponding non-cooperative equilibrium bid profile $\mathbf{b}^l$.*

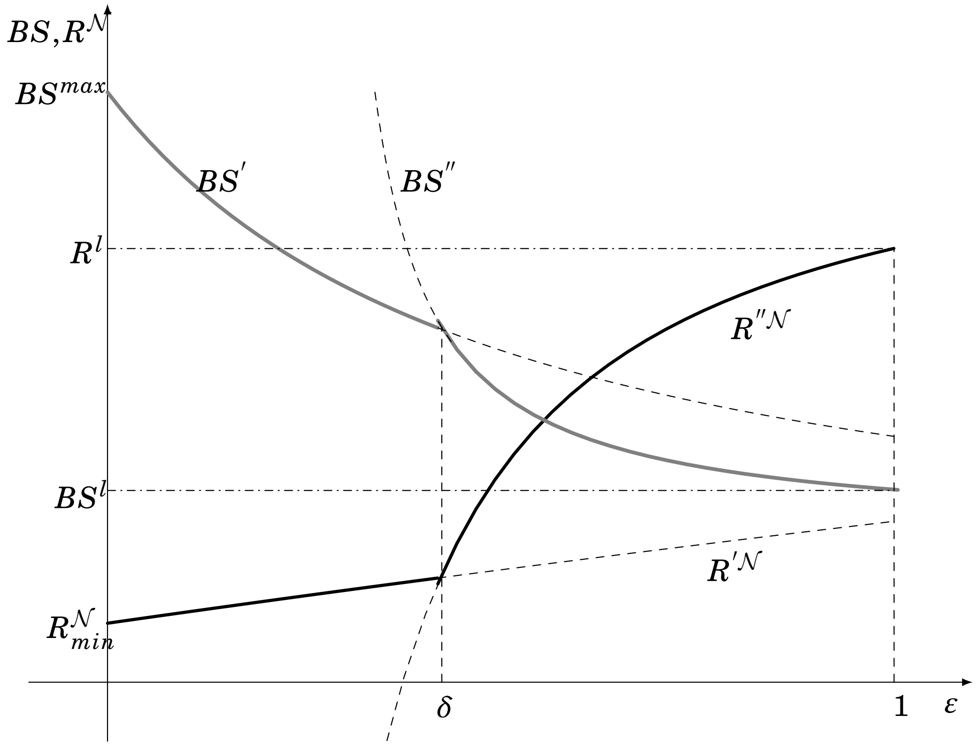

Figure 5: The auctioneer's revenue $R^{\mathcal{N}}$ and bidders' surplus BS as a function of ε

An illustrative representation of the auctioneer's revenue $R^{\mathcal{N}}$ (black lines) as a function of ε against the overall bidders' surplus BS (grey lines) is given in Figure 5. If $\varepsilon = 0$, the overall surplus will be entirely divided between members and the auctioneer; nothing is expropriated by it. The auctioneer obtains the minimum revenue that the presence of the efficient cartel generates and the members obtain the maximum surplus they can retain from him. In Figure 5 we denote by BS^l the overall surplus achieved under outcome $\mathbf{b}^l$, and by R^l the auctioneer's maximum revenue.

Proposition 3 coupled with corollary 1 shows that the *agency*'s strategic behaviour secures a lower bound in the seller's surplus.

Corollary 2. *The optimal commission fee mitigates the adverse effect of coordination and secures a minimum revenue for the seller.*

Now, given that the auctioneer cannot artificially or indirectly drive up the collusive bids by claim 1, and given that it strictly benefits from an increase in the fee, we would naturally assume that it could be in its best interests to enter into a tacit agreement with the *agency*. However, it is not in the interests of the *agency* to set $\varepsilon > \delta$. Therefore, the auctioneer could handle sufficient compensations so as to make this strategy profitable for the *agency*. If the *bidding agency* sets the maximum fee (if it intends to act as a sort of *non-cooperative* equilibrium decentralizer), i.e., $\varepsilon = 1$, $\boldsymbol{\mu}_{\mathcal{N}}$ converges to $\mathbf{b}^{l}$, so does the auctioneer's revenue. In order to maintain a sufficient degree of spoils, members of the cartel need to push the equilibrium prices forward, which is mechanically beneficial to the seller despite the fact that this surplus now has to be divided between the players, the seller and the *bidding agency*.

3.3 Substitutability between positions and competition

Finally, we compare the link between the click ratio θ as a measure of the degree of substitutability of positions and the *demand-reduction* phenomenon observed in the *GSP* auction and, more generally, in multi-object auctions. Bid shading occurs in situations in which players' bids affect the price they have to pay at the end of the auction which relies on the assumption that bidders have multi-unit demands. As a result, this should disappear whenever the same player asks for one object (one position) at the most.

In the *GSP* auction, advertisers can ask for only one position. Notice the special feature in the competition nature and equilibrium predictions suggested by the bid profile $\mathbf{b}^{v}$. In equilibrium, it is optimal for each player to bid strictly below their own valuation (except for the non-assigned player). In fact, this phenomenon can be extended or broken down with the differentiation in clicks between positions. This points to the existence of a close link between *ctr* and the nature of the competition in the *GSP* auction.

Remark 3. *The relation between the ratio $\theta = \frac{\alpha_2}{\alpha_1}$, the equilibrium non-cooperative bids and payments is summarized in the following observations:*

(i) *As $\theta \mapsto 0$, bid shading disappears. The game converges essentially to a standard second-price auction.*

(ii) *As $\theta \mapsto 1$, bid shading increases. The game essentially converges to a* Bertrand *competition.*

Take the set of *Nash* equilibrium bids in (7) and observe that the bids are convex combinations weighted by the ratio θ. Without loss of generality, let us consider the refinement whereby in the *non-cooperative* game the non-assigned bidder (bidder 3) uses a dominant strategy, i.e., $b_3 = x_3$.

We obtain $b_1 \in \left[x_2 - \frac{\alpha_2}{\alpha_1}(x_2 - x_3)\,;\bar{x}\right]$, $b_2 \in \left[x_3\,; x_1 - \frac{\alpha_2}{\alpha_1}(x_1 - x_3)\right]$ and $b_3 = x_3$. It is now obvious to see that when $\theta = 0$, i.e., the first position is the only object worth winning (the most attractive one), bids are bounded by the next-to-top and the next-to-bottom valuations. That is we obtain $b_1 \in [x_2\,;\bar{x}]$, $b_2 \in [x_3\,; x_1]$ and $b_3 = x_3$. The link is even closer at the *symmetric* equilibrium, where indeed, bidding one's valuation is the only possible bid, as it is in standard second-price auctions.

Now, for (ii), note that, as $\theta = 1$, both positions are substitutes and the middle-valuing player's equilibrium bid is equal to $\max\{b_3, x_3\}$. There is no opportunity for the highest-valuing player to undercut him and the lowest-valuing advertiser sets the market-clearing price to his private valuation. We have that $b_1 \in [x_3\,;\bar{x}]$, $b_2 = x_3$ and $b_3 = x_3$. This situation highlights the analogy with a Bertrand competition and the well-known *demand-reduction* phenomenon in standard multi-object auctions[22]. The lowest-valuing bidder's valuation determines the shape of the market-clearing price. Nevertheless, there is still an opportunity for the auctioneer to make the top position the only one worth having by destroying incentives to shade in the *non-cooperative* play.

Coordinated bids: However, there are no clear tendencies for the link between θ and the nature of the competition with an active cartel.

Remark 4. *The relation between the ratio $\theta = \frac{\alpha_2}{\alpha_1}$, the equilibrium coordinated bids and payments is summarized in the following observations:*

(i) *As $\theta \mapsto 0$, bid shading increases. The game essentially converges to a collusive implementation in a single unit second-price auction.*

(ii) *As $\theta \mapsto 1$, bid shading persists but there is no clear tendency in its magnitude. It depends on the* agency*'s objective and on the equilibrium bid profile* $\mathbf{b}$ *fixed as outside option.*

If the substitutability between positions was to decrease, that is $(\alpha_1 - \alpha_2)$ increases, then if the *bidding agency* behaves as in proposition 2 and set $\varepsilon = \delta$, the coordinate bids of player 2 would equal zero and the link would be reversed. As a result, we would nest the standard result of collusion in second-price auctions with multi-unit objects and bid shading is exacerbated[23].

[22]See for instance Engelbrecht-Wiggans and Kahn (1998), Ausubel et al. (2014).

[23]e.g., Graham and Marshall (1987), McAfee and McMillan (1992), Milgrom (2000), Brusco and Lopomo (2002)

However, if $\theta = 1$, which means that $(\alpha_1 - \alpha_2)$ increases, we obtain $b_2^{\mathcal{N}} = \frac{3x_3-(1-\varepsilon)(b_2+b_3)}{2+\varepsilon}$ and $b_3^{\mathcal{N}} = 0$, which depends on the bid profile $\mathbf{b}$. If we assume that the *bidding agency* seeks to maximise the bidders surplus, then it coordinates the outside option towards $\mathbf{b}^l = \left(b_2^l = x_3, b_3^l = x_1 - \frac{\alpha_1}{\alpha_2}(x_1 - x_3)\right)$ and when $\varepsilon = \delta$ we get $b_2^{\mathcal{N}} = \frac{1}{2}x_3$ and $b_3^{\mathcal{N}} = 0$.

For $\varepsilon > \delta$, then $b_2^{\mathcal{N}} = x_3 = b_3^{\mathcal{N}}$. Thus, if $(\alpha_1 - \alpha_2)$ increases, then for $\varepsilon = \delta$ the bid placed by player 2 equals half the valuation of the lowest-valuing bidder, whereas $\forall \varepsilon > \delta$ $b_2^{\mathcal{N}}$ increases to x_3. As a result, coordinated bids increase in θ and in the degree of substitutability between positions as highlighted by fact 1.

3.4 Link with the LEF criterion and VCG outcome

In this section we show that, if we scale the analysis over the spirit of standard competitive equilibrium analysis by using the LEF criterion in the present framework, then a profitable way to collude does not exist except by achieving the *VCG* outcome. In other words, if deviations are managed by members according to the argument of *local indifference* then the *GSP* auction becomes collusion-proof and the equivalence with the *Vickrey-Clarke-Groves* solution is restored and unique.

What would the effect be if we were to apply the LEF *stability condition* to the collusive game, selecting from the collusive equilibria the one that respects this criterion? In other words, let us consider the LEF criterion as a *stability condition* instead of the expression of a *Walrasian tatônnement* process (as in Börgers et al. (2013)), the latter being necessarily related to a *non-cooperative* analysis. Assume still that, in the case of deviation from the collusive agreement, the defector pays his individual contribution computed *before the targeting auction starts*, where individual contributions are computed unconditionally on deviations and are based on an assortative assignment. In addition, assume still that the *bidding agency* makes its transfer payment after the main auction. The *incentive compatibility* conditions are now given by the following relations:

$$\pi_k^{\mathcal{N}}\left(b_k^{\mathcal{N}}, \mathbf{b}_{-i}^{\mathcal{N}}, x_k\right) - \tilde{\pi}_k\left(\tilde{b}_k, \mathbf{b}_{-i}^{\mathcal{N}}, x_k\right) \quad \geq \quad 0 \tag{10}$$

where $\forall k \in \mathcal{K}$:

$$\tilde{\pi}_k\left(\tilde{b}_k,.\right) = \alpha_{k-1}\left(x_k - b_k^{\mathcal{N}}\right) - \left(\alpha_k b_{k+1} - \alpha_k b_{k+1}^{\mathcal{N}}\right) \tag{11}$$

This condition says that when the player assigned to position k intends to defect from the collusive agreement and contemplates position $k-1$, then he expects to pay the same price as the player assigned to this position. Hence, we give the mechanism a semblance of symmetry in the collusive equilibrium conditions and this results in each member bidding the *non-cooperative symmetric* equilibrium prices.

Proposition 5. *Under the locally envy-free condition and a* uniform-ε *redistribution a cartel with three players can do no better than achieve the* VCG *outcome. That is:* $\boldsymbol{\mu}_{\mathcal{N}}^* = \left(b_1^{\mathcal{N}} > b_2^{\mathcal{N}}; \frac{P_1^v}{\alpha_1}; x_3\right) = \mathbf{b}^{\mathbf{v}}$.

The underlying justification of the choice of the respective *non-cooperative* bids is as follows. In the *non-cooperative* game, if bidders pick, from the set of competitive prices, the one that maximises their individual profits, then the *LEF* criterion results in them implementing the *VCG* outcome. In our present context, the *bidding agency* assumes the role of a *mediator* that manages individual choices. Then, if it were to act as a benevolent *mediator*, under the scope of the competitive equilibrium it would choose the vector of *non-cooperative* prices that maximises the bidders' surplus. The purpose of proposition 5 is to show that the *bidding agency* cannot implement a compensation policy that is compatible with a *first-best* collusive outcome under this stability criterion. It is interesting to note that bidding functions are now completely independent of the *agency*'s compensation policy. Hence, the key idea of this proposition is that the *GSP* auction, together with an efficient collusion, implements, in equilibrium, the *VCG*-equivalent outcome as a result of explicit coordination. Sponsored search markets entail a tri-partite structure in which third parties bid on behalf of different advertisers. As a result, our model offers a consistent justification as to why bidders behave according to the *LEF* refinement when bid delegation is involved.

4 Concluding remarks

Bidding agencies have emerged to help advertisers internalize their costs of knowledge acquisition and to absorb the level of competition they themselves would create by acting non-cooperatively. To the best of our knowledge, this study is the first to offer a specific examination of an explicit bid coordination implemented by a *bidding agency* at a *GSP* auction.

Specifically, the paper proposes a characterisation of the bid profiles that can be implemented in a two-position game involving three players. This characterisation was subsequently increased to include five members and four positions , but it appears a closed-form solution cannot be computed for a higher number of players. Numerical simulations suggest that the same intuition prevails, which is why we opted to focus solely on the case of three advertisers and two positions. Bidders can efficiently collude via the *bidding agency*, but this adversely affects *GSP* auction revenues. This study also proposes a closed-form expression for the optimum fee level that can be implemented and highlights the non-monotonicity of the agency's revenues. It is argued that greedy behaviour on the part of the *bidding agency* is detrimental to the sustainability of an efficient collusion but that it manages to coordinate advertisers at the least efficient *Nash* equilibrium. Our results suggest that an *agency* can deliver a profitable bid delegation service, and that its presence in the sponsored search market increases the advertisers' surplus and contributes to market efficiency. We believe that our model offers a stylized justification of both the prevalence of the tri-partite market structure and of the compensation policies based on a fixed fee that prevail in this market.

The auctioneer is passive, which raises the question as to what it should do and if it it is in its interests to prevent collusion in the auction. We argue that it may not be in the auctioneer's interests to deter collusion if it takes a long-run perspective. A search engine's revenues are based on maintaining a strong market position, which is derived from advertisers' feelings that they are obtaining a high surplus. A one-shot auction might be greatly impaired by delegation but, in the long run, in line with the reasoning of Bulow and Klemperer (1996) and Roughgarden and Sundararajan (2007), the auction should generate near-optimal revenues. Thus, from our perspective it should be in the interests of the auctioneer to maintain this structure and even cooperate with *bidding agencies* as shown by corollary 1 and proposition 4. Finally, the non-cooperative *VCG* solution is uniquely determined through the process of explicit coordination, which offers an intuitive and consistent justification when implemented non-cooperatively. We conclude by commenting on the assumptions made in our analysis and look at possible extensions of our work.

Our analysis rests on four main assumptions. First, we do not consider the possibility of repeated interactions, as we deal solely with the one-shot *GSP* auction game. Second, the *agency* occupies a monopoly position and does not have to face an adverse selection issue when implementing the collusive device. Third, once a potential member rejects the proposed contract then the cartel is not formed. Finally, individual contributions and the lump-sum transfer are computed and redistributed uniformly among members and are independent of individual deviations. In practice, the *GSP* auction is automatically triggered each time a search query is entered by a consumer, and this may favour tacit collusion as bidders have opportunities to compete against the same set of competitors. Our intuition is that the outcome will remain the same as the *bidding agency* can implement the additional threat of triggering non-cooperative bidding once a defection occurs. An interesting consideration here is the possibility of implementing a bid-rotation mechanism as in McAfee and McMillan (1992).

Multiple bidding agencies and incentives for the seller. Even if multiple *agencies* bidding in the name of different advertisers during the same set of keyword auctions appears rarely, an extension would be to integrate a competitive stage that includes different third parties and assess the effect of competition for the auction performance. This would introduce additional strategic scope and might act in favour of an increase in the level of competition with respect to our results. For instance, what would the outcome be of introducing an oligopoly of *agencies* that compete *à la* Bertrand by offering the lowest compatible fee?

Identity-dependent and deviation-dependent side payments. The uniform fee and uniform redistribution of spoils assumptions seem to be good approximations of *agency* practices. They levy a fee, set *ex-ante*, on the revenues generated by each click to each advertisers. We believe that our model, therefore, makes sense in relation to the equal redistribution assumption. Indeed, the model involves taxation on the collusive spoils that does not take into account individual members' identity and importance. As a result, further extensions should consider the case in which the flat-fee becomes part of the individual contract, as a function of the member's importance within the cartel. For instance, considering redistribution as a function of each firm's market power in the product market would be a suitable task to pursue.

A Appendix

A.1 Proof of proposition 2

We seek to construct a collusive equilibrium bid profile that is compatible with the incentive compatibility constraints defined in (4) and (5) and with the individual rationality constraint defined by the relation (6). We set $\theta = \frac{\alpha_2}{\alpha_1}$.

From player 2's incentive compatibility constraint, we know he will not deviate to position 1 if the following relation is satisfied $\alpha_2\left(x_2 - b_3^{\mathcal{N}}\right) - \omega_2 + \frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} \geq \alpha_1\left(x_2 - b_1^{\mathcal{N}}\right) - \omega_2$. This relation gives the following conditions:

$$
\begin{aligned}
b_3^{\mathcal{N}} &\leq \frac{1}{(4-\varepsilon)\alpha_2}\left((1-\varepsilon)\sum_{i=1}^{2} p_i - (1-\varepsilon)\alpha_1 b_2^{\mathcal{N}} + 3\alpha_1 b_1^{\mathcal{N}} - (\alpha_1 - \alpha_2)3x_1\right) \\
b_1^{\mathcal{N}} &\geq \frac{1}{3\alpha_1}\left((\alpha_1 - \alpha_2)3x_2 - (1-\varepsilon)\sum_{i=1}^{2} p_i - (1-\varepsilon)\alpha_1 b_2^{\mathcal{N}} + (4-\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right)
\end{aligned}
$$

giving the first lower bound for player 1 and an upper bound for player 3. Now from player 1's incentive compatibility constraint, we know that if he does not want to swap his position for position 2, then it should be the case that $\alpha_1\left(x_1 - b_2^{\mathcal{N}}\right) - \omega_1 + \frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} \geq \alpha_2\left(x_1 - b_3^{\mathcal{N}}\right) - \omega_1$. This gives the following relations:

$$
\begin{aligned}
b_2^{\mathcal{N}} &\leq \frac{1}{(4-\varepsilon)\alpha_1}\left((\alpha_1 - \alpha_2)3x_1 + \sum_{i=1}^{2} p_i + (2+\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right) \\
b_3^{\mathcal{N}} &\geq \frac{1}{(2+\varepsilon)\alpha_2}\left((4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} - \sum_{i=1}^{2} p_i - (\alpha_1 - \alpha_2)3x_1\right)
\end{aligned}
$$

the first upper bound for player 2 and the lower bound for player 3. Finally, if we examine player

3's incentive compatibility constraint, we obtain $\frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} \geq \alpha_1\left(x_3 - b_1^{\mathcal{N}}\right)$ and $\frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} \geq \alpha_2\left(x_3 - b_2^{\mathcal{N}}\right)$, giving respectively:

$$b_1^{\mathcal{N}} \geq x_3 - \frac{1}{3\alpha_1}\left((1-\varepsilon)\sum_{i=1}^{2} w_i\right)$$
$$b_2^{\mathcal{N}} \geq \frac{1}{3\alpha_2-(1-\varepsilon)\alpha_1}\left(3\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2} p_i + (1-\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right)$$

the second lower bounds for player 1 and 2. The participation constraint equals $\sum_{i=1}^{2} p_i \geq \sum_{i=1}^{2} \alpha_i b_{i+1}^{\mathcal{N}}$, which implies $b_2^{\mathcal{N}} \leq b_2$ and $b_3^{\mathcal{N}} \leq b_3$. The above inequalities result in the following equilibrium strategy profile $\boldsymbol{\mu}_{\mathcal{N}} = \left(b_i^{\mathcal{N}}\right)_{i=1,2,3}$:

Set of collusive bids 1.

$$\begin{aligned}
&b_1^{\mathcal{N}} \in [\mathcal{A}; \bar{x}] \\
&b_2^{\mathcal{N}} \in [\mathcal{B}; \mathcal{C}] \\
&b_3^{\mathcal{N}} \in \left[\max\left\{0; \frac{1}{(2+\varepsilon)\alpha_2}\left((4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} - (1-\varepsilon)\sum_{i=1}^{2} p_i - (\alpha_1-\alpha_2)3x_1\right)\right\}; \mathcal{D}\right]
\end{aligned} \tag{12}$$

with

$$\begin{aligned}
\mathcal{A} &= \max\left\{x_3 - \frac{(1-\varepsilon)}{3\alpha_1}\sum_{i=1}^{2} w_i; \frac{1}{3\alpha_1}\left((\alpha_1-\alpha_2)3x_2 - (1-\varepsilon)\sum_{i=1}^{2} p_i - (1-\varepsilon)\alpha_1 b_2^{\mathcal{N}} + (4-\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right)\right\} \\
\mathcal{B} &= \max\left\{0; \frac{1}{3\alpha_2-(1-\varepsilon)\alpha_1}\left(3\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2} p_i + (1-\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right)\right\} \\
\mathcal{C} &= \frac{1}{(4-\varepsilon)\alpha_1}\left((\alpha_1-\alpha_2)3x_1 + \sum_{i=1}^{2} p_i + (2+\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right) \\
\mathcal{D} &= \frac{1}{(4-\varepsilon)\alpha_2}\left((1-\varepsilon)\sum_{i=1}^{2} p_i - (1-\varepsilon)\alpha_1 b_2^{\mathcal{N}} + 3\alpha_1 b_1^{\mathcal{N}} - (\alpha_1-\alpha_2)3x_1\right)
\end{aligned}$$

We use the following claim which gives us a tool with which to discriminate among the set of compatible collusive bids:

Claim 2. *A sufficient condition for $b_2^{\mathcal{N}} > 0$ is $\theta = \frac{\alpha_2}{\alpha_1} \geq \frac{(1-\varepsilon)b_2}{3x_3-(1-\varepsilon)b_3}$.*

Proof. Set $b_1^{\mathcal{N}}$ so that neither player 2 nor player 3 has incentives to deviate to the top position.

The incentive compatibility conditions are:

$$
\begin{aligned}
(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)-3\alpha_2 x_3-(1-\varepsilon)\alpha_2 b_3^{\mathcal{N}}+(3\alpha_2-(1-\varepsilon)\alpha_1)b_2^{\mathcal{N}} \geq 0 \quad & IC_3\\
3\alpha_2 x_2+(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)-(4-\varepsilon)\alpha_2 b_3^{\mathcal{N}}+(1-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \geq 0 \quad & IC_2\\
3\alpha_1 x_1+(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)-(1-\varepsilon)\alpha_2 b_3^{\mathcal{N}}+(4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \geq 0 \quad & IC_1\\
(\alpha_1-\alpha_2)3x_1+(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)+(2+\varepsilon)\alpha_2 b_3^{\mathcal{N}}-(4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \geq 0 \quad & IC_1^{'}
\end{aligned}
$$

where IC_3 is player 3's incentives to deviate for position 2, IC_2 is the player 2's incentives to occupy no position, IC_1 is player 1's incentives to occupy no position and $IC_1^{'}$ is player 1's incentives to occupy position 2. It can be observed that IC_3 implies both IC_2 and IC_1. Now, assume that $b_2^{\mathcal{N}}=b_3^{\mathcal{N}}=0$, we have:

$$
\begin{aligned}
(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)-3\alpha_2 x_3 \geq 0 \quad & IC_3\\
(\alpha_1-\alpha_2)3x_1+(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3) \geq 0 \quad & IC_1^{'}
\end{aligned}
$$

Thus, these bids are compatible if the following is true: $(1-\varepsilon)(b_2+\theta b_3)-3\theta x_3 \geq 0$ and $(1-\theta)3x_1+(1-\varepsilon)(b_2+\theta b_3) \geq 0$: that is, if $\theta \leq \frac{(1-\varepsilon)b_2}{3x_3-(1-\varepsilon)b_3}=\eta$ or if $\theta \leq \frac{3x_1+(1-\varepsilon)b_2}{3x_3-(1-\varepsilon)b_3}$, where the first inequality implies the second. Thus, $b_2^{\mathcal{N}}=b_3^{\mathcal{N}}=0$ are compatible bids if $\theta \leq \eta$. □

Since, $b_1^{\mathcal{N}}$ deters any deviation for the top position, it suffices to consider the adjacent deviations. Then, according to the claim, to ensure an efficient assignment we impose $\theta>\eta$ so that $b_2^{\mathcal{N}}>0$ and maintain the bid $b_3^{\mathcal{N}}=0$. Both constraints are,

$$
\begin{aligned}
(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}-3\alpha_2 x_3+(3\alpha_2-(1-\varepsilon)\alpha_1)b_2^{\mathcal{N}} \geq 0 \quad & IC_3\\
(\alpha_1-\alpha_2)3x_1+(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}-(4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \geq 0 \quad & IC_1^{'}
\end{aligned}
$$

as $x_3<x_1$ if player 3's incentive to occupy position 2 is binding, it should also be the case for player 1. Thus, we are left with the constraint IC_3 in equilibrium which implies that $(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}-$

$3\alpha_2 x_3 + (3\alpha_2 - (1-\varepsilon)\alpha_1) b_2^{\mathcal{N}} = 0$. That is,

$$b_2^{\mathcal{N}} = \frac{1}{3\alpha_2 - (1-\varepsilon)\alpha_1}\left(3\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}\right)$$

Notice that, $b_2^{\mathcal{N}} < x_3$ since we can re-write the last expression as,

$$\begin{aligned}
b_2^{\mathcal{N}} &= \frac{1}{3\alpha_2 - (1-\varepsilon)\alpha_1}\left(3\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}\right) \\
&= x_3 - \left(\frac{(1-\varepsilon)}{(1-\varepsilon)\alpha_1 - 3\alpha_2}(\alpha_1(x_3 - b_2) - \alpha_2 b_3)\right) \\
&= x_3 - \lambda(\alpha_2 x_3 - \alpha_2 b_3 - \alpha_2 x_3 + \alpha_1 x_3 - \alpha_1 b_2) \\
&= x_3 - \lambda(\pi_3 + (\alpha_1 - \alpha_2) x_3 - \alpha_1 b_2)
\end{aligned}$$

and that $IC_1^{'}$ is satisfied with such bid,

$$\begin{aligned}
&(\alpha_1 - \alpha_2)3x_1 + (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - (4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \\
&\geq (\alpha_1 - \alpha_2)3x_1 + (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - (4-\varepsilon)\alpha_1 x_3 \\
&\geq (\alpha_1 - \alpha_2)3x_3 + (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - (4-\varepsilon)\alpha_1 x_3 \\
&\geq (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - (3\alpha_2 - (1-\varepsilon)\alpha_1) x_3 \geq 0
\end{aligned}$$

Now, consider the case in which $b_2^{\mathcal{N}} > 0$ and $b_3^{\mathcal{N}} > 0$. This is only possible towards the parameter ε. Assume that there exists a threshold δ so that if $\varepsilon > \delta$ then $b_2^{\mathcal{N}} > 0$ and $b_3^{\mathcal{N}} > 0$. Both constraints equal,

$$\begin{aligned}
(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - 3\alpha_2 x_3 - (1-\varepsilon)\alpha_2 b_3^{\mathcal{N}} + (3\alpha_2 - (1-\varepsilon)\alpha_1) b_2^{\mathcal{N}} \geq 0 \quad & IC_3 \\
(\alpha_1 - \alpha_2)3x_1 + (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} + (2+\varepsilon)\alpha_2 b_3^{\mathcal{N}} - (4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} \geq 0 \quad & IC_1^{'}
\end{aligned}$$

Both constraints are binding at the optimum, otherwise the *bidding agency* can decrease $b_2^{\mathcal{N}}$

and $b_3^{\mathcal{N}}$ accordingly. We obtain,

$$b_2^{\mathcal{N}} = \frac{1}{3\alpha_2 - (1-\varepsilon)\alpha_1}\left(3\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} + (1-\varepsilon)\alpha_2 b_3^{\mathcal{N}}\right)$$
$$b_3^{\mathcal{N}} = \frac{1}{(2+\varepsilon)\alpha_2}\left((4-\varepsilon)\alpha_1 b_2^{\mathcal{N}} - (1-\varepsilon)\sum_{i=1}^{2} p_i - (\alpha_1-\alpha_2)3x_1\right)$$

Plugging the expression of $b_2^{\mathcal{N}}$ in that of $b_3^{\mathcal{N}}$ results in the following relation,

$$b_3^{\mathcal{N}} = \frac{\left(x_1(\alpha_1-\alpha_2)((1-\varepsilon)\alpha_1 - 3\alpha_2) - (1-\varepsilon)(\alpha_1+\alpha_2)\sum_{i=1}^{2}\alpha_i b_{i+1} + (4-\varepsilon)\alpha_1\alpha_2 x_3\right)}{\alpha_2((2+\varepsilon)\alpha_2 - (1-\varepsilon)2\alpha_1)} \tag{13}$$

which we replace in the expression of $b_2^{\mathcal{N}}$ to obtain,

$$b_2^{\mathcal{N}} = \frac{\left((2+\varepsilon)\alpha_2 x_3 - (1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1} - (1-\varepsilon)(\alpha_1-\alpha_2)x_1\right)}{(2+\varepsilon)\alpha_2 - (1-\varepsilon)2\alpha_1} \tag{14}$$

Proof of the monotonicity and the decrease in **b** is straightforward. Denote by $\underline{b}_i^{\mathcal{N}}$ and $\bar{b}_i^{\mathcal{N}}$ the respective collusive bids for $\varepsilon \leq \delta$ and $\varepsilon > \delta$. $\forall \varepsilon \leq \delta$ the derivative of $b_2^{\mathcal{N}}$ with respect to ε is given by $\frac{\partial\left(\underline{b}_2^{\mathcal{N}}\right)}{\partial\varepsilon} = \frac{3\alpha_2(\alpha_1 b_2 + \alpha_2 b_3 - \alpha_1 x_3)}{(3\alpha_2 - (1-\varepsilon)\alpha_1)^2} \geq 0$ and the derivative with respect to $R = (\alpha_1 b_2 + \alpha_2 b_3)$ by $\frac{\partial\left(\underline{b}_2^{\mathcal{N}}\right)}{\partial R} = -\frac{1-\varepsilon}{3\alpha_2 - \alpha_1(1-\varepsilon)}$ which is negative if $3\alpha_2 - \alpha_1(1-\varepsilon) \geq 0$ that is if $\frac{\alpha_2}{\alpha_1} \geq \frac{1-\varepsilon}{3}$ which is satisfied by the restriction that $b_2^{\mathcal{N}} > 0$. Now, $\forall \varepsilon > \delta$ the derivatives are given by:

$$\frac{\partial\left(\bar{b}_2^{\mathcal{N}}\right)}{\partial\varepsilon} = \frac{3\alpha_2(\alpha_1 b_2 + \alpha_2 b_3 + (\alpha_1-\alpha_2)x_1 - 2\alpha_1 x_3)}{(\alpha_2(2+\varepsilon) - 2\alpha_1(1-\varepsilon))^2} \geq 0$$
$$\frac{\partial\left(\bar{b}_3^{\mathcal{N}}\right)}{\partial\varepsilon} = \frac{3(\alpha_1+\alpha_2)(\alpha_1 b_2 + \alpha_2 b_3 + (\alpha_1-\alpha_2)x_1 - 2\alpha_1 x_3)}{(\alpha_2(2+\varepsilon) - 2\alpha_1(1-\varepsilon))^2} \geq 0$$
$$\frac{\partial\left(\bar{b}_2^{\mathcal{N}}\right)}{\partial R} = \frac{-(1-\varepsilon)}{(2+\varepsilon)\alpha_2 - (1-\varepsilon)2\alpha_1} \leq 0$$
$$\frac{\partial\left(\bar{b}_3^{\mathcal{N}}\right)}{\partial R} = \frac{-(1-\varepsilon)(\alpha_1+\alpha_2)}{\alpha_2((2+\varepsilon)\alpha_2 - (1-\varepsilon)2\alpha_1)} \leq 0$$

in which the last two denominators present positive signs if $\theta \geq 2\frac{(2+\varepsilon)}{(1-\varepsilon)}$.

A.2 Proof of proposition 3

To find the value of the threshold δ simply find the value of ε that solves,

$$\frac{\left((2+\varepsilon)\alpha_2 x_3-(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}-(1-\varepsilon)(\alpha_1-\alpha_2)x_1\right)}{(2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1}-\frac{\left(3\alpha_2 x_3-(1-\varepsilon)\sum_{i=1}^{2}\alpha_i b_{i+1}\right)}{3\alpha_2-(1-\varepsilon)\alpha_1}=0$$

$$and$$

$$\frac{\left(x_1(\alpha_1-\alpha_2)((1-\varepsilon)\alpha_1-3\alpha_2)-(1-\varepsilon)(\alpha_1+\alpha_2)\sum_{i=1}^{2}\alpha_i b_{i+1}+(4-\varepsilon)\alpha_1\alpha_2 x_3\right)}{\alpha_2((2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1)}=0$$

this gives,

$$\varepsilon^*=\delta=\frac{\left(1+\frac{\alpha_2}{\alpha_1}\right)\sum_{i=1}^{2}\alpha_i b_{i+1}-\left(1-\frac{\alpha_2}{\alpha_1}\right)(\alpha_1-3\alpha_2)x_1-4\alpha_2 x_3}{\left(1+\frac{\alpha_2}{\alpha_1}\right)\sum_{i=1}^{2}\alpha_i b_{i+1}-\left(1-\frac{\alpha_2}{\alpha_1}\right)\alpha_1 x_1-\alpha_2 x_3}$$

$\forall\varepsilon\in[0,\delta^*]$ the equilibrium strategy profile is characterised by the sets (9). Thus the *bidding agency*'s profit is equal to the following quantity:

$$\Gamma'_{\mathcal{N}}=\varepsilon\left\{\alpha_1\left(b_2-b_2^{\mathcal{N}}\right)+\alpha_2\left(b_3-b_3^{\mathcal{N}}\right)\right\}=\varepsilon\left\{R-\alpha_1\left(\frac{3\alpha_2 x_3-(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)}{3\alpha_2-(1-\varepsilon)\alpha_1}\right)\right\}$$

and $\forall\varepsilon\in(\delta^*,1)$, by the equilibrium bids (13) and (14), it takes the following quantity:

$$\begin{aligned}\Gamma''_{\mathcal{N}}=\varepsilon\Bigg\{R-\alpha_1&\left(\frac{(2+\varepsilon)\alpha_2 x_3-(1-\varepsilon)(\alpha_1 b_2+\alpha_2 b_3)-(1-\varepsilon)(\alpha_1-\alpha_2)x_1}{(2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1}\right)\\ -\alpha_2&\left(\frac{x_1(\alpha_1-\alpha_2)((1-\varepsilon)\alpha_1-3\alpha_2)-(1-\varepsilon)(\alpha_1+\alpha_2)(\alpha_1 b_2+\alpha_2 b_3)+(4-\varepsilon)\alpha_1\alpha_2 x_3}{\alpha_2((2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1)}\right)\Bigg\}\end{aligned}$$

The first and second derivative of $\Gamma'_{\mathcal{N}}$ are respectively equal to:

$$\frac{\partial}{\partial\varepsilon}\left(\Gamma'_{\mathcal{N}}\right)=3\alpha_2(3\alpha_2-\alpha_1)\frac{(\alpha_1 b_2+\alpha_2 b_3-\alpha_1 x_3)}{(3\alpha_2-(1-\varepsilon)\alpha_1)^2}$$

which is positive if $3\alpha_2\geq\alpha_1$ and

$$\frac{\partial^2}{\partial\varepsilon^2}\left(\Gamma'_{\mathcal{N}}\right)=-6\alpha_1\alpha_2(3\alpha_2-\alpha_1)\frac{(\alpha_1 b_2+\alpha_2 b_3-\alpha_1 x_3)}{(3\alpha_2-(1-\varepsilon)\alpha_1)^3}$$

which presents a negative sign whenever $3\alpha_2\geq\alpha_1$. Thus, in the domain $[0,\delta^*)$ the *bidding agency*'s profit function is concave. The function is non-differentiable in $\varepsilon = \delta^*$ and the first and second derivatives of $\Gamma''_{\mathcal{N}}$ are, respectively, given by:

$$\frac{\partial}{\partial \varepsilon}\left(\Gamma''_{\mathcal{N}}\right) = -6\alpha_2(\alpha_1-\alpha_2)\frac{(\alpha_1 b_2+\alpha_2 b_3+x_1(\alpha_1-\alpha_2)-2\alpha_1 x_3)}{((2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1)^2}$$

which presents a negative sign and

$$\frac{\partial^2}{\partial \varepsilon^2}\left(\Gamma''_{\mathcal{N}}\right) = 12\alpha_2(\alpha_1-\alpha_2)(2\alpha_1+\alpha_2)\frac{(\alpha_1 b_2+\alpha_2 b_3+x_1(\alpha_1-\alpha_2)-2\alpha_1 x_3)}{((2+\varepsilon)\alpha_2-(1-\varepsilon)2\alpha_1)^2}$$

which presents a positive sign. Thus, the profit function is a convex function over the domain $(\delta^*, 1)$.

A.3 Proof of claim 1

From proposition 3, we know that the *bidding agency* has no incentives to set $\varepsilon > \delta$; thus, it is sufficient to focus on the bid functions of relation (9). Take the equilibrium bid profile $\boldsymbol{\mu}_{\mathcal{N}}$ of proposition 2 and rearrange the expressions so that the bids become a function of θ. From proposition 3, we know that the *agency* has no incentives to set $\varepsilon > \delta$; thus, it is sufficient to focus on the bid functions of relation (9). Consider the situation in which $\alpha_1 \mapsto \infty$ so that $\theta = 0$, from claim 2 we have that $b_2^{\mathcal{N}} = b_3^{\mathcal{N}} = 0$. Hence, the auctioneer's revenues would fall to zero. If now the auctioneer sets $\alpha_1 = \alpha_2$ so that $\theta = 1$, rearrange relation (9) to obtain $b_2^{\mathcal{N}}(\theta) = \frac{3\theta x_3-(1-\varepsilon)(b_2+\theta b_3)}{3\theta-(1-\varepsilon)}$. From proposition 3 it is optimal for the *bidding agency* to set $\varepsilon = \delta$; thus, we obtain $b_2^{\mathcal{N}} = \frac{\alpha_2 x_3+\alpha_2(b_2-b_3)-2\alpha_1 b_2}{2\alpha_2(x_3-\alpha_1 b_2+\alpha_2 b_3)}$, which equals $\frac{1}{2}x_3$ if the outside bid profile is assumed to be $\mathbf{b}^l = \left(b_2^l = x_3, b_3^l = x_1 - \frac{\alpha_1}{\alpha_2}(x_1-x_3)\right)$ and x_3 if $\mathbf{b}^l = \left(b_2^l = x_3, b_3^l = 0\right)$ is considered. As a result, the auctioneer's revenue cannot be higher than in the corresponding equilibrium profile $\mathbf{b}^l$.

A.4 Proof of proposition 5

In order for player 2 to be indifferent to winning the second position at price $b_3^{\mathcal{N}}$ or winning top position at price $b_2^{\mathcal{N}}$, the following relation should be satisfied:

$$\alpha_2\left(x_2 - b_3^{\mathcal{N}}\right) - \omega_2 + \frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} = \alpha_1\left(x_2 - b_2^{\mathcal{N}}\right) - \omega_2$$

For player three to be indifferent to being assigned to the second position or not being assigned at all under the collusive agreement it should be that $\frac{1-\varepsilon}{3}\Pi_{\mathcal{N}} = \alpha_2\left(x_3 - b_3^{\mathcal{N}}\right)$. Rearranging both relations, we obtain the following pair of equations:

$$\alpha_1 b_2^{\mathcal{N}} = \frac{3}{2+\varepsilon}\left(x_2(\alpha_1-\alpha_2)\right) + \frac{4-\varepsilon}{2+\varepsilon}\alpha_2 b_3^{\mathcal{N}} - \frac{1-\varepsilon}{2+\varepsilon}\left(\sum_{i=1}^{2} P_i^v\right) \tag{15}$$

$$\alpha_2 b_3^{\mathcal{N}} = \frac{3}{2+\varepsilon}\alpha_2 x_3 + \frac{1-\varepsilon}{2+\varepsilon}\left(\alpha_1 b_2^{\mathcal{N}} - \sum_{i=1}^{2} P_i^v\right) \tag{16}$$

Recall that $P_i^v = \sum_{k=i+1}^{m+1} x_k(\alpha_{k-1}-\alpha_k)$. Rearranging terms, plugging (16) into (15) and using equation (8), we obtain $\alpha_1 b_2^{\mathcal{N}} = x_2(\alpha_1-\alpha_2) + \alpha_2 x_3 = P_1^v$ and $\alpha_2 b_3^{\mathcal{N}} = \alpha_2 x_3 = P_2^v$. The solution corresponds for any $\varepsilon \in [0,1]$ to the same equilibrium bids and payments of the *VCG*-equivalent equilibrium bids profile of equation (8) and thus results in the same outcome as the *symmetric non-cooperative* equilibrium.

B Experimentations

We run 1000 instance of the one-shot *GSP* game, using the General Algebraic Modeling System (GAMS), in which, following Cary et al. (2008), each valuation is drawn from a distribution $G(x) \sim \mathcal{N}(500,200)$ setting the value of X_1, X_2 and X_3 respectively to $x_1 = 592.7$, $x_2 = 565.535$ and $x_3 = 437.331$. On average, we observe a *ctr* of 0.23% on higher positions which allows us to reasonably set $\lambda = 0.23$ giving $E(\alpha) = 4.3$ (Synodiance 2013 synodiance.ctr.study2013) thus $H(\alpha) \sim \mathcal{E}_{(0.23)}$. The exponential law generates numbers lying between 0 and 1 which can be directly interpreted as click probability or click rates. For each instance, we computed the collusive bids resulting from

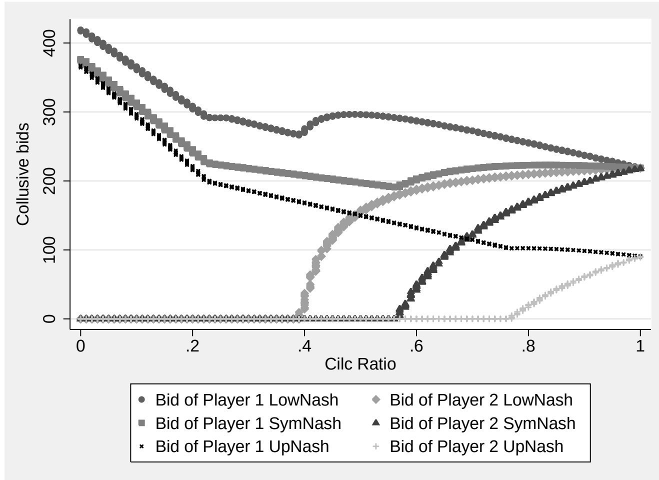

Figure 6: Equilibrium collusive bids when $\varepsilon = 0$ against θ respectively with outside bid profile $\mathbf{b}^l, \mathbf{b}^v, \mathbf{b}^u$ when $H(\alpha) \sim \mathcal{E}_{(0,23)}$.

the *agency*'s maximisation problem for each $\varepsilon \in [0,1]$. The outside option $\pi_i(b_i, \mathbf{b}_{-i}, x_i)$ is *ex-ante* determined according to $\mathbf{b}^l, \mathbf{d}^d, \mathbf{b}^v, \mathbf{b}^u$. The bid vector $\mathbf{b}^l$ involves each bidder playing at the lower bound of equation (7), $\mathbf{b}^d$ is simply $\mathbf{b}^l$ with the restriction that bidder 3 plays truthfully, $\mathbf{b}^v$ are the bids sustaining the *VCG-equivalent* outcome defined by the bid vector in equation (8) and $\mathbf{b}^u$ corresponds to the upper bound of equation (7).

B.1 Bids, profits and revenues

Figure 6 depicts the positive relation existing between the ratio of clicks, we denoted by θ, when $\varepsilon = 0$ and the cooperative bids of players 1 and 2. In such a situation, the rivalry from the lowest advertiser is muted and the collusive outcome is maximised. This figure shows the underlying idea of proposition 2 that the degree of difference in clicks between both positions has to be low enough to induce an efficient allocation. That is, to avoid a random assignment between the lowest-valuing member and the second-highest valuing one.

Figure 7 expresses the underlying idea of corollary 1. That is, all individual surpluses are lower-bounded by the individual surpluses produced in the game sustained by the corresponding *non-cooperative* price vector. We represented the average difference between collusive surpluses when contributions are computed according to each outside option price, i.e., $p_k^g = b_{k+1}^g$ with $g = l,d,v,u$, and: (i) the *non-cooperative* surplus produced by bid vector $\mathbf{b}^l$ and (ii) the surplus that

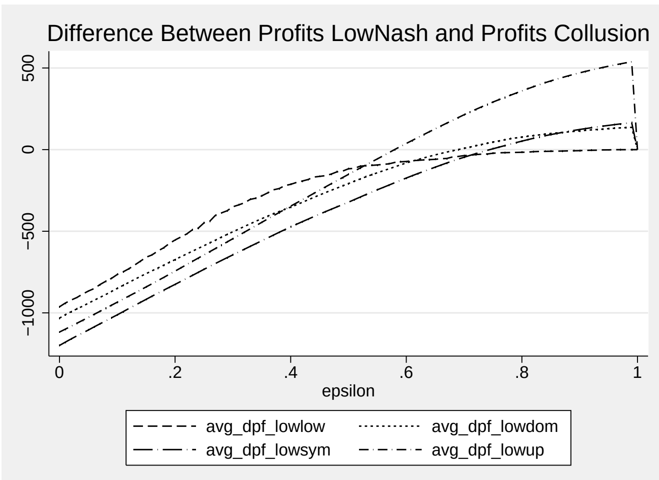

(a) Difference in profits with profits in *Low*-Nash

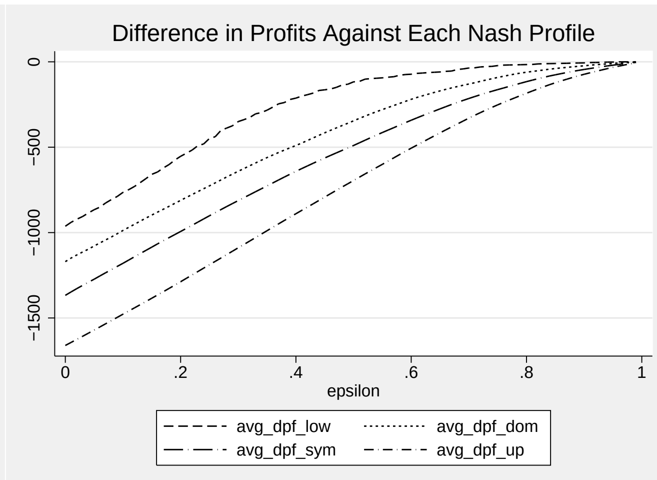

(b) Difference in profits with each $\mathbf{b}^g$

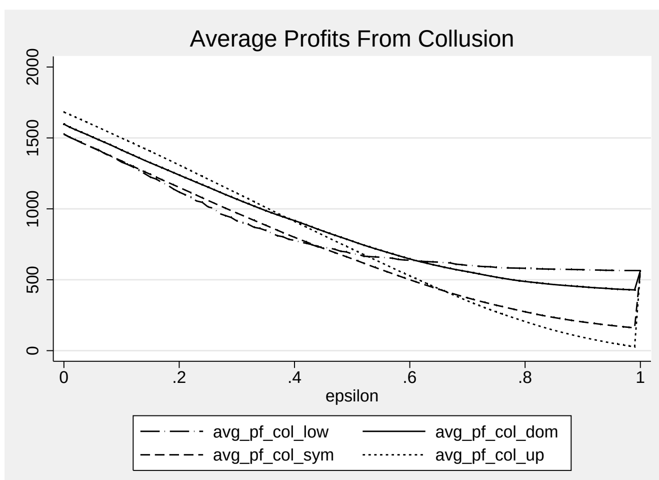

(c) Joint profits from collusion

Figure 7: Average difference in surplus with the bids sustaining $\mathbf{b}^l$, in payoffs against each outside option and joint profits when the *agency* set a fee $\varepsilon \geq 0$.

individuals would have in each corresponding *non-cooperative* equilibrium. Let $\psi_{(.)}$ be the average difference for each case. We computed $\psi_{(i)} = \pi_k\left(p_k^l\right) - \pi_k^{\mathcal{N}}\left(p_k^g\right)$ and $\psi_{(ii)} = \pi_k\left(p_k^g\right) - \pi_k^{\mathcal{N}}\left(p_k^g\right)$. We can observe that $\psi_{(ii)} \mapsto 0$ as $\varepsilon \mapsto 1^-$ which makes sense as both individual contributions tend to zero and each individual rationality constraint is binding as the collusive bids increase with ε. Although for high enough values of ε, the bid profile computes according to the outside option $\mathbf{b}^l$ strictly dominates all other collusive profiles from the bidders' viewpoint. Note the disruption at $\varepsilon = 1$. Joint profits jump to the joint surplus achieved by $\mathbf{b}^l$ (Figures 7a and 7c).

Figure 8 depicts the idea of proposition 4 that the revenue of the seller is upper-bounded by the revenue produced under price vector $\mathbf{b}^l$. If the collusive mechanism is computed according to the outside option $\mathbf{b}^l$, $\mathbf{b}^d$, $\mathbf{b}^v$ or $\mathbf{b}^u$, we can observe in Figure 8d, that as $\varepsilon \mapsto 1$ collusive prices converge to the level sustaining the *non-cooperative* outcome $\mathbf{b}^l$. As a result, revenues for the seller also converge to the same revenue level (around 3372) as shown in Figure 8c.

B.2 Comparison between the benchmark and the collusive outcome and thresholds for ε

Table 1 describes the collusive outcome implemented by a *bidding agency* that acts as a benevolent agent, i.e. $\varepsilon = 0$ and the *non-cooperative* outcome corresponding to each Nash extremum we consider. Conversely, Tables 3 and 4 describe the solution implemented by the mechanism when the *agency* imposes a flat-fee $\varepsilon > 0$ on collusive gains and the corresponding *non-cooperative* outcomes, again for each Nash extremum. We represent each individual profit, contribution, seller's revenue, joint profit (cartel surplus) and individual bid.

Description of the non-cooperative equilibria and of the collusive mechanism with a benevolent *bidding agency* from Table 1: If each player behaves according to the lower bound of relations (7), i.e., $\mathbf{b}^l$, then equilibrium bids involve bid shading from each player with respect to their valuations ($b_1 = 453.53 < 592, b_2 = 437.3 < 565.5$ and $b_3 = 170.1 < 437.3$), the final allocation is efficient as player 1 obtains the first position and player 2 the second (however it is not *locally-envy free*) and produces a total revenue of $R_u = 3450.8$ for the seller. Notice that

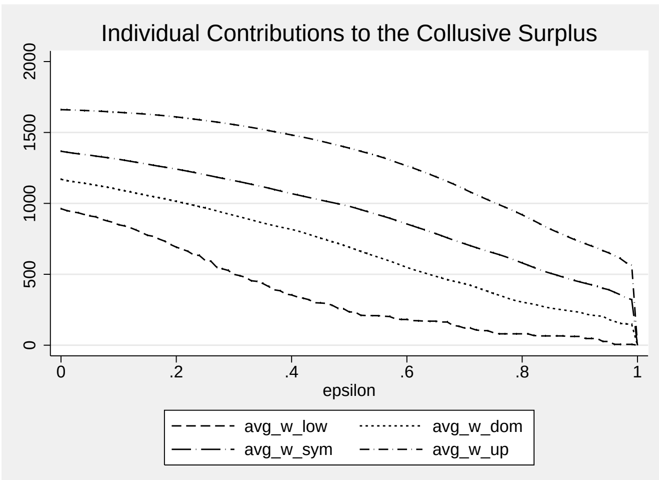

(a) Individual Contributions

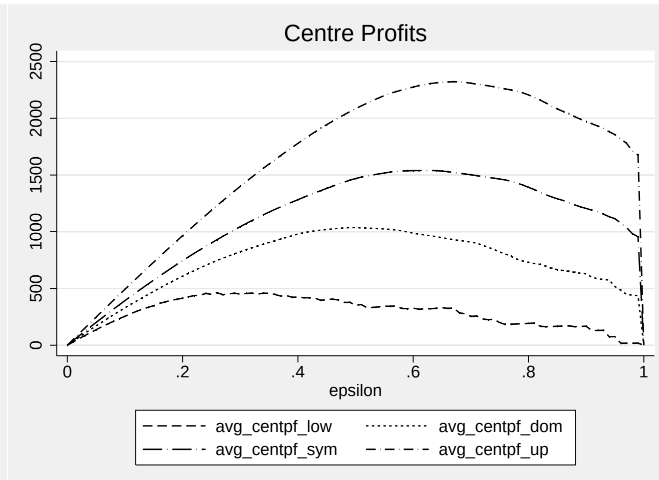

(b) Revenues of the *agency*

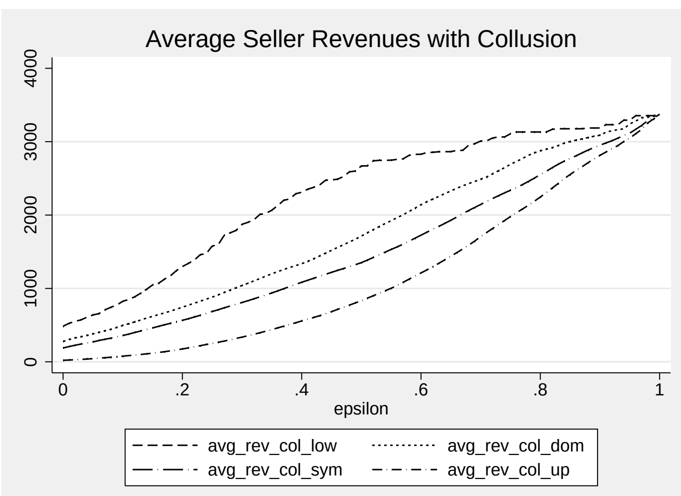

(c) Seller's revenues

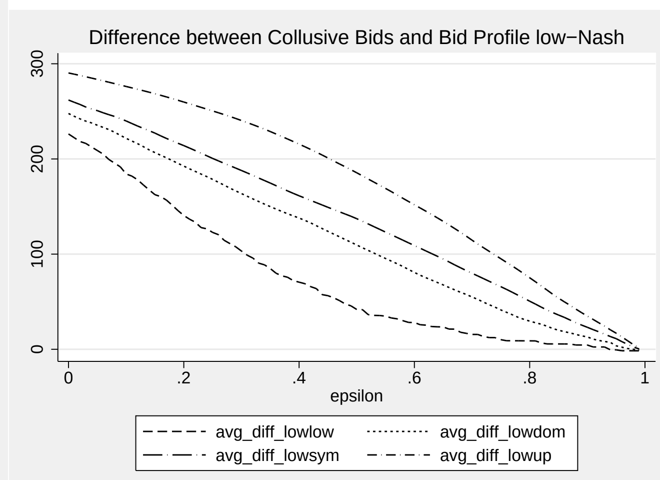

(d) Difference in bids

Figure 8: Average seller's revenues, individual contributions, revenues of the bidding agency according to outside options $\mathbf{b}^l, \mathbf{b}^d, \mathbf{b}^v$ and $\mathbf{b}^u$ and difference between bid profile $\mathbf{b}^l$ and each collusive bid profile.

	Benevolent *agency*									
	Non-Cooperative Play					Cartel				
	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Mean	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Mean
Surplus										
Player 1	1024.32	1024.32	466.81	60.98	644.11	2039.84	2231.134	1878.852	1756.32	1976.53
Player 2	700.97	287.65	287.65	0	319.06	1716	1494.49	1699.67	1695.34	1651.52
Player 3	.	.	.	.	.	1015.54	1206.84	1412.02	1695.34	1332.43
Contributions										
Player 1	.	.	.	.	.	2478.76	2639.27	3254.81	3817.13	3047.5
Player 2	.	.	.	.	.	567.86	981.24	981.24	1268.9	949.8
Player 3	.	.	.	.	.	0	0	0	0	
Average Contributions	.	.	.	.	.	1015.542	1206.84	1412.02	1332.43	
Seller Revenue	3450.82	3864.2	4421.7	5115.17	4212.9	404.2	243.7	185.64	29.15	215.67
Total Surplus	5176.15	5176.15	5176.15	5176.15	5176.15	5176.15	5176.15	5176.15	5176.15	5176.15
Cartel Surplus	.	.	.	.	.	4771.95	4932.46	4990.5	5147	4960.48
Bids										
Player 1	453.53	514.83	592.71	592.71	538.44	304.38	274.31	244.81	194.88	254.6
Player 2	437.3	437.3	514.83	581.96	492.86	78.89	49.9	38.98	7.04	43.73
Player 3	170.14	437.3	437.3	565.54	402.59	0	0	0	0	0
Average Bids	353.67	463.16	514.96	580.07		127.76	108.1	94.6	67.31	

Table 1: Equilibrium outcomes with a benevolent *agency* for $H(\alpha) \sim \mathcal{E}_{(0.23)}$ and $G(x) \sim \mathcal{N}_{(500,200)}$

Threshold δ^*

	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$
δ^*	0.143	0.241	0.318	0.516

Table 2: Threshold δ^* according to $\mathbf{b}^l, \mathbf{b}^d, \mathbf{b}^v, \mathbf{b}^u$

this outcome strictly dominates the other three in terms of bidders' surplus. It transpires that if we consider the *bidding agency* as endorsing the role of a mediating device , when setting $\varepsilon = 1$, which is an equivalent situation as collusive profits strictly equate *non-cooperative* profits, the *agency* implements a Nash outcome that Pareto-dominates any other *non-cooperative* outcome from the bidders' viewpoint.

If the *agency* were to set $\varepsilon = 0$, i.e., no surplus extraction from the collusion, then internal rivalry is reduced to the lowest level compatible with an efficient allocation and the mechanism implements the *first-best* outcome as defined in definition 2. Only the highest and the second-highest valuing member are active at the main auction whereas the lowest-valuing member is inactive. For instance, take the solution implemented according to $\mathbf{b}^u$. Bidder 1 bids $b_1^{\mathcal{N}} \simeq 195$, bidder 2 bids $b_2^{\mathcal{N}} \simeq 7.04$ and bidder 3 does not bid. Although it does not affect allocative efficiency, the active collusion destroys the seller's revenues as they fall from $R^u = 5115.17$ to $R_u^{\mathcal{N}} = 29.15$ if bidders behave according to $\mathbf{b}^u$, which would also correspond to the best solution in terms of the bidders' welfare, i.e., $BS_u^{\mathcal{N}} = 5147$.

Description of Tables 3 and 4: The computed values of the optimal fee threshold δ^* are presented for each outside option in Table 2. We take the overall average value of the 1000 instances run. We compute in Table 4 the collusive solution when the *agency* set $\varepsilon > 0$ for which data were arbitrarily divided in three parties. Results for the optimum situation $\varepsilon^* = \delta^*$ are depicted in the middle class of data, the left-hand side of the table corresponds to a flat-fee $\varepsilon' = \frac{\delta^*}{2}$ and the right hand side to a flat-fee $\varepsilon'' = 1.5\delta^*$. In contrast to the benevolent situation, the *agency*, by setting the flat-fee to its optimal level, implements an outcome that, from the bidders' viewpoint, strictly Pareto-dominates all other collusive solutions, i.e., the solution computed according to $\mathbf{b}^l$ for which

	Corresponding NC outcomes				
	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Average
Surplus					
Player 1	790.28	790.28	633.53	104.9	579.77
Player 2	685.22	495.27	495.27	0	418.94
Player 3	.	.	.	.	.
Seller Revenue	3723.88	3913.85	4070.6	5094.43	4200.69
Total Surplus	5199.3	5199.3	5199.3	5199.3	5199.3
Bids					
Player 1	431.12	466.65	592.71	592.71	520.79
Player 2	437.3	437.3	466.65	571.75	478.25
Player 3	385.58	437.3	437.3	565.54	456.45
Average Bids	418	447.1	498.9	576.67	

Table 3: Equilibrium *NC* outcomes corresponding to the non-neutral *agency* case for $H(\alpha) \sim \mathcal{E}_{(0.23)}$ and $G(x) \sim \mathcal{N}_{(500,200)}$

individual surpluses are $\pi_1^{\mathcal{N}} \simeq 1448.38$, $\pi_2^{\mathcal{N}} \simeq 1343.4$ and $\pi_3^{\mathcal{N}} \simeq 658$.

It can be observed that the *bidding agency*'s profits are maximum when $\varepsilon^* = \delta^*$ and strictly increase whenever the outside option is assumed to be that of high equilibrium prices, in contrast to the bidders' individual surpluses. Once $\varepsilon > \delta^*$, these profits mechanically fall as collusive bids strictly increase causing contributions to fall. This corresponds to the underlying idea of proposition 3. For instance, the solution implemented according to the outside option $\mathbf{b}^l$ gives $\Gamma_{\mathcal{N}} \simeq 201.31$, an average contribution $\bar{\omega} \simeq 831.77$ when $\varepsilon' = 0.5 * \delta^*$. This gives $\Gamma_{\mathcal{N}} \simeq 366.8$, an average contribution $\bar{\omega} \simeq 780.37$ when $\varepsilon' = \delta^*$. This gives $\Gamma_{\mathcal{N}} \simeq 0$, the average contribution $\bar{\omega} \simeq 0$ when $\varepsilon' = 1.5 * \delta^*$. The seller strictly benefits from an increase in ε and, note, its revenues are quite close to their corresponding *non-cooperative* level (around $R^{\mathcal{N}} \simeq 3723$). Once the fee level crosses the threshold δ^*, collusive bids strictly converge to the lowest Nash equilibrium in price and the lowest-valuing member competition can no longer be constricted.

Active *agency*

Fee	$\varepsilon' = 0.5\delta^*$					$\varepsilon = \delta^*$					$\varepsilon'' = 1.5\delta^*$				
	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Mean	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Mean	$\mathbf{b}^l$	$\mathbf{b}^d$	$\mathbf{b}^v$	$\mathbf{b}^u$	Mean
Threshold	0.085	0.115	0.136	0.236	.	0.169	0.23	0.271	0.472	.	0.254	0.345	0.407	0.709	.
Surplus															
Player 1	1554.94	1604.94	1487.98	1203.46	1462.83	1448.38	1448.38	1291.63	763.1	1237.87	790.272	984.68	867.175	324.98	741.77
Player 2	1449.94	1309.97	1349.76	1098.47	1302.03	1343.4	1153.41	1153.41	658.10	1077.07	685.28	689.71	728.85	219.98	580.98
Player 3	764.66	814.67	854.46	1098.47	883.06	658.1	658.1	658.1	658.1	658.1	0	194.41	233.65	219.98	162.01
Agency Payoff	201.31	314.73	410.47	1025.63	488.03	366.8	556.77	713.52	1737.35	843.61	0	287.91	458.4	1455.03	550.33
Seller Revenue	1228.58	1155.12	1096.76	773.41	1063.47	1382.78	1382.78	1382.78	1382.78	1382.78	3723.9	3042.72	2911.25	2979.45	3164.33
Total Surplus	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43	5199.43
Cartel Surplus	3769.54	3729.58	3692.2	3400.4	3647.93	3449.86	3259.88	3103.14	2079.31	2973.05	1475.55	1868.8	1829.78	764.95	1484.8
Contributions															
Player 1	995.7	1069.15	1284.28	2136.15	1371.34	841.5	841.5	998.25	1526.78	1052	0	243.38	446.24	947.5	409.23
Player 2	1499.6	1689.58	1689.58	2184.9	1765.9	1499.6	1689.58	1689.58	2184.9	1765.9	0	627.76	713.1	1167.49	627.1
Player 3	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
Average	831.77	919.58	991.28	1440.34	.	780.37	843.69	895.94	1237.22	.	0	290.38	386.45	705	.
Bids															
Player 1	283.71	274.37	266.93	218.38	260.85	303.73	303.73	303.73	303.73	303.73	437.3	397.39	389.06	390.81	403.65
Player 2	240.38	226.73	215.9	152.52	208.88	269.17	269.17	269.17	269.17	269.17	437.3	388.1	378.43	382.65	396.63
Player 3	.	.	.	.	.	.	.	.	.	.	385.58	274.14	253.07	265.43	294.56
Average Bids	174.7	167.03	160.94	123.63	.	190.97	190.97	190.97	190.97	190.97	420.1	353.21	340.18	346.3	.

Table 4: Collusive outcomes when the *agency* sets $\varepsilon = \{0.5\delta^*, \delta^*, 1.5\delta^*\}$ for $H(\alpha) \sim \mathcal{E}_{(0.23)}$ and $G(x) \sim \mathcal{N}_{(500,200)}$

References

Aggarwal, G., Muthukrishnan, S. and Feldman, J.: 2006, Bidding to the top: Vcg and equilibria of position-based auctions, *CoRR* .

Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R. and Tennenholtz, M.: 2010, Position auctions with budgets: Existence and uniqueness, *The BE Journal of Theoretical Economics* **10**(1).

Ashlagi, I., Monderer, D. and Tennenholtz, M.: 2009, Mediators in position auctions, *Games and Economic Behavior* **67**, 2–21.

Ashlagi, I., Monderer, D. and Tennenholtz, M.: 2011, Simultaneous ad auctions, *Mathematics of Operations Research* **36**(1), 1–13.

Athey, S. and Ellison, G.: 2011, Position auction with consumer search, *Quarterly Journal of Economics* **126**(3), 1213–70.

Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M. and Weretka, M.: 2014, Demand reduction and inefficiency in multi-unit auctions, *Review of Economic Studies* **81**(4), 1366–1400.

Balseiro, S. R. and Candogan, O.: 2017, Optimal contracts for intermediaries in online advertising, *Operations Research* **65**(4), 878–896.

Bernheim, B. D. and Whinston, M. D.: 1985, Common marketing agency as a device for facilitating collusion, *The RAND Journal of Economics* pp. 269–281.

Bernheim, B. D. and Whinston, M. D.: 1986a, Common agency, *Econometrica* pp. 923–942.

Bernheim, B. D. and Whinston, M. D.: 1986b, Menu auctions, resource allocation, and economic influence, *Quarterly Journal of Economics* **101**(1), 1–31.

Börgers, T., Cox, I., Persendorfer, M. and Petricek, V.: 2013, Equilibrium bids in sponsored search auctions: Theory and evidence, *American Economic Journal: Microeconomics* **5**(4), 163–187.

Brusco, S. and Lopomo, G.: 2002, Collusion via signalling in simultaneous ascending bid auctions with heterogenous objects, with and without complementarities, *The Review of Economic Studies* **69**(2), 407–463.

Bulow, J. and Klemperer, P.: 1996, Auctions versus negotiations, *The American Economic Review* **86**(1), 180–194.

Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A. R., Mathieu, C. and Schwarz, M.: 2008, On best-response bidding in gsp auctions, *NBER Working Paper W13788*, National Bureau of Economic Research.

Decarolis, F., Goldmanis, M. and Penta, A.: 2017, Marketing agencies and collusive bidding in online ad auctions, *NBER Working Paper No. 23962* .

Decarolis, F. and Rovigatti, G.: 2017, Online auctions and digital marketing agencies, *NET Institute Working Paper No. 17-08* .

Edelman, B., Ostrovsky, M. and Schwarz, M.: 2007, Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords, *American Economic Review* **97**(1), 242–259.

Edelman, B. and Schwartz, M.: 2010, Optimal auction design and equilibrium selection in sponsored search auctions, *American Economic Review* **100**, 597–602.

Engelbrecht-Wiggans, R. and Kahn, C. M.: 1998, Multi-unit auctions with uniform prices, *Economic Theory* **12**(2), 227–258.

Feldman, J., Mirrokni, V., Muthukrishnan, S. and Pai, M. M.: 2010, Auctions with intermediaries, *Proceedings of the 11th ACM conference on Electronic commerce*, ACM, pp. 23–32.

Feldman, M., Meir, R. and Tennenholtz, M.: 2011, Revenue enhancement in ad auctions, *Internet and Network Economics*, Springer, pp. 391–398.

Gomes, R. and Sweeney, K.: 2014, Bayes-nash equilibria of the genealized second-price auction, *Games and Economic Behavior* **86**, 421–437.

Graham, D. and Marshall, R.: 1987, Collusive bidder behavior at single-object second-price and english auctions, *Journal of Political Economy* **95**(6), 1217–1239.

Lopomo, G., Marx, L. M. and Sun, P.: 2010, Bidder collusion at first-price auctions, *Review of Economic Design* .

Lucier, B., Paes Leme, R. and Tardos, E.: 2012, On revenue in the generalized second price auction, *Proceedings of the 21st international conference on World Wide Web*, ACM, pp. 361–370.

Mailath, G. and Zemsky, P.: 1991, Collusion in second price auctions with heterogenous bidders, *Games and Economic Behavior* **3**(4), 467–486.

Marshall, R. C. and Marx, L. M.: 2007, Bidder collusion, *Journal of Economic Theory* **133**(1), 374–402.

McAfee, R. P.: 2011, The design of advertising exchanges, *Review of Industrial Organization* **39**(3), 169–185.

McAfee, R. P. and McMillan, J.: 1992, Bidding rings, *American Economic Review* **82**(3), 579–599.

Milgrom, P. R.: 2000, Putting auction theory to work: The simultaneous ascending auction, *Journal of Political Economy* **108**(2), 245–272.

Monderer, D. and Tennenholtz, M.: 2009, Strong mediated equilibrium, *Artificial Intelligence* **173**(1), 180–195.

Muthukrishnan, S.: 2009, Ad exchanges: Research issues, *Internet and network economics* pp. 1–12.

Roughgarden, T. and Sundararajan, M.: 2007, Is efficiency expensive.

Varian, H.: 2007, Position auctions, *International Journal of Industrial Organization* **25**, 1163–1178.

Fondation Paris-Dauphine

Chaire Gouvernance et Régulation
Fondation Paris-Dauphine
Place du Maréchal de Lattre de Tassigny - 75016 Paris (France)
http://chairgovreg.fondation-dauphine.fr