

Exclusive Placement in Online Advertising

Amin Sayedi,^a Kinshuk Jerath,^b Marjan Baghaie^c

^a University of Washington, Seattle, Washington 98195; ^b Columbia Business School, Columbia University, New York, New York 10027;

^c Uber Technologies, San Francisco, California 94103

Contact: aminsa@uw.edu (AS); jerath@columbia.edu,  <http://orcid.org/0000-0003-0732-5863> (KJ); marjan@uber.com (MB)

Received: March 16, 2017

Revised: October 11, 2017; February 26, 2018

Accepted: March 1, 2018

Published Online in Articles in Advance:
August 30, 2018

<https://doi.org/10.1287/mksc.2018.1098>

Copyright: © 2018 INFORMS

Abstract. A recent development in online advertising has been the ability of advertisers to have their ads displayed exclusively on (a part of) a web page. We study this phenomenon in the context of both sponsored search advertising and display advertising. Ads are sold through auctions, and when exclusivity is allowed, the seller accepts two bids from advertisers, where one bid is for the standard display format in which multiple advertisers are displayed, and the other bid is for being shown exclusively (therefore they are called two-dimensional, or 2D, auctions). We identify two opposing forces at play in an auction that provides the exclusive placement option—allowing more flexible expression of preferences through bidding for exclusivity increases competition among advertisers, leading to higher bids, which increases the seller’s revenue (between-advertiser competition effect), but it also gives advertisers the incentive to shade their bids for their non-preferred outcomes, which decreases the seller’s revenue (within-advertiser competition effect). Depending on which effect is stronger, the revenue may increase or decrease. We find that the 2D generalized second price (GSP_{2D}) auction, which is an extension of the widely used generalized second price (GSP) auction and on which currently used auctions for exclusive placement are based, may lead to higher or lower revenue under different parametric conditions. Paradoxically, the revenue from allowing exclusive placement decreases as bidders have higher valuations for exclusive placement. We verify several key implications from our analysis of GSP_{2D} using data from Bing for over 100,000 auctions. As a possible solution (applicable to both sponsored search and display advertising), we show that using VCG_{2D}, which is the adaptation of the Vickrey–Clarke–Groves (VCG) auction for the 2D setting, guarantees weakly higher revenue when exclusive display is allowed. This is because it induces truthful bidding, which alleviates the problem of bid shading due to the within-advertiser competition effect.

History: Ganesh Iyer served as the senior editor and Dmitri Kuksov served as associate editor for this article.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mksc.2018.1098>.

Keywords: exclusive display • sponsored search advertising • display advertising • multidimensional auctions • rich ads in search

1. Introduction

Online advertising, which includes display and sponsored search advertising, accounts for over 35% of all advertising expenditure in the United States, and this share is projected to grow further in the future (eMarketer 2016). This constitutes the primary source of revenue for many firms, including Google, Bing, and Facebook, and these and other firms experiment and innovate continually (Manzi 2012) to enhance the advertising options and the associated pricing mechanisms that they provide to advertisers. It is well known, for instance, that search engines arrived at the currently widely used generalized second price (GSP) auction mechanism after trying a number of different pricing mechanisms (Edelman and Ostrovsky 2007, Edelman et al. 2007).

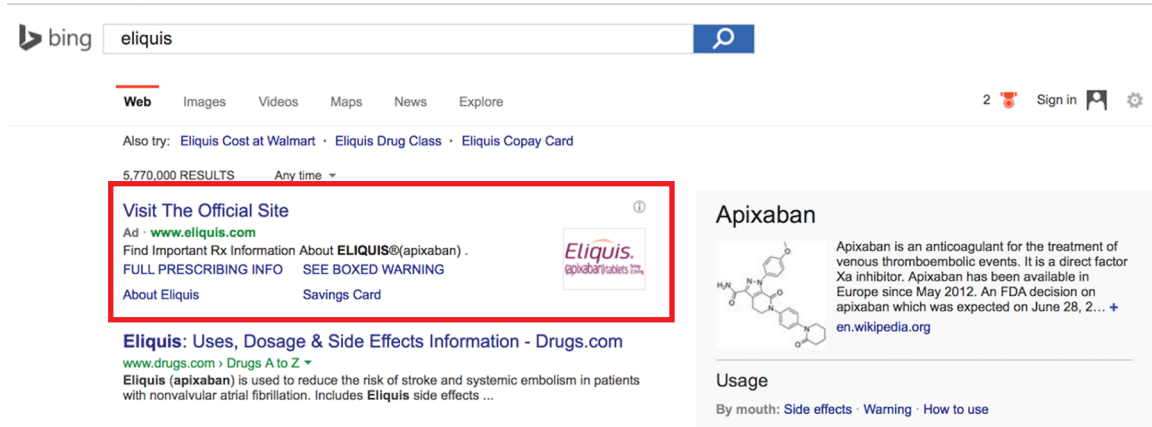
A recent interesting development in this context has been the advent of exclusive placement of ads on a web page in both sponsored search and display advertising.

In the case of sponsored search, an advertiser can bid for its ad to be the only one displayed on the results page, or on the “north” area of the results page (i.e., at the top of a search results page) in response to a keyword search (Figure 1(a)). Similarly, in the case of display advertising, an advertiser can bid for its ad to be displayed exclusively in a panel that might otherwise include ads from multiple advertisers (Figure 1(b)), or to be displayed in multiple or all ad slots on the web page (Figure 1(c)).

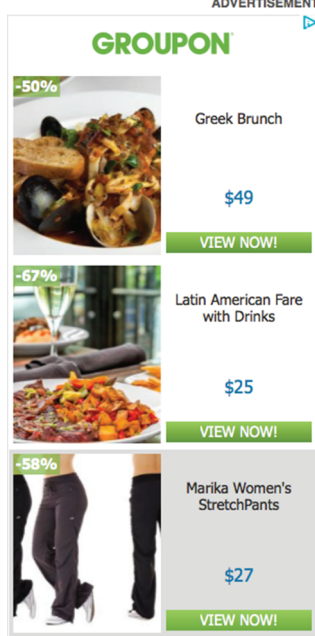
Exclusive placement is an attractive option for advertisers for various reasons. An exclusively displayed ad on a web page can be expected to obtain a larger number of total impressions or clicks compared with an ad displayed along with multiple other ads, because it does not compete with other ads for the viewer’s attention. Jeziorski and Segal (2015) use data from Microsoft’s search engine, Bing, to estimate that an ad

Figure 1. (Color online) Examples of (a) an Exclusive Display Sponsored Search Ad in the North Slot in Response to a Keyword Search on Bing; (b) a Panel Exclusively Displaying Ads for Groupon, When Three Different Ads Could Possibly Have Been Displayed; and (c) a Web Page Exclusively Displaying Ads for grubHub, When Two Different Ads Could Possibly Have Been Displayed


(a) Sponsored search example



(b) Display example



(c) Display example



displayed without competing ads will obtain approximately 50% more clicks, on average. Furthermore, exclusive placement can increase an advertiser's valuation conditional on a click, as it can create strong brand associations on the consumer's side by being the only ad displayed, to prevent negative externalities from other ads. Such effects may motivate advertisers to prefer exclusive placement. This raises the possibility of higher revenues for ad sellers and market makers by allowing exclusive placement because advertisers may bid higher for exclusive ads; however, sellers also run the risk of losing revenue from the advertisers who will not be displayed.

In 2011, Bing and Yahoo! launched the Rich Ads in Search (RAIS) program through which they allowed advertisers to bid for exclusive placement for "north" ads (which account for up to 85% of all clicks on ads on a search results page; Reiley et al. 2010) for their trademarked keywords. However, Bing saw higher revenue in RAIS only for some keywords, and lower revenue for other keywords. In 2016, Bing temporarily suspended the RAIS initiative. Along similar lines, Google also experimented with displaying exclusive ads, that is, only one ad per page, as part of its "perfect ad" initiative (Metz 2008, 2011). After their initial experiments, Google also suspended this initiative.

Nevertheless, our conversations with researchers at Microsoft indicate that the exclusive-display offering is one that advertisers want to use, and Bing and other search engines want to get right. However, there is a lack of clarity on the auction mechanism to use and on the associated impact on revenue. Indeed, efforts are being made to develop mechanisms to address these issues in a robust manner (Simonov et al. 2015).

Interestingly, in parallel, exclusive placement is gaining increasing traction in display advertising, where advertisers are able to place ads in all slots in a panel or multiple ads on the same page. Our conversations with practitioners in the display advertising industry reveal that there is sustained demand from advertisers for exclusive placement ads.¹ However, these ads are typically sold manually, while the display ad market is increasingly adopting the much more efficient programmatic buying and selling. It is expected that if a programmatic auction-based method can be developed to transact exclusive placement ads, then their share will increase significantly.

Overall, the picture that emerges is that while there is promise in exclusive placement advertising, there is a need to obtain clarity on its pros and cons, and develop appropriate mechanisms to transact them. In this paper, we use the tools of game theory to take a first step toward developing an understanding of exclusive ad placement for both sponsored search and display advertising. Using a stylized model, we identify key driving forces at play. This enables us to suggest auction mechanisms that improve on the status quo, while being within the scope of currently used mechanisms.

In the model, we assume that each advertiser can have different per-click valuations for clicks obtained when it is displayed with other advertisers (multiple placement) and clicks obtained when it is the only one displayed (exclusive placement). Ads are sold through auctions, and when exclusivity is allowed, auctioneers accept two bids from advertisers (therefore they are called two-dimensional, or 2D, auctions), where one bid is for the standard display format in which multiple advertisers are displayed and the other bid is for being shown exclusively.

We identify two key opposing forces at play in exclusive placement. First, competition is heightened between advertisers because they can be more expressive in revealing their preferences to the search engine, and they compete not only for positions in the nonexclusive outcome but also compete for the outcome to be exclusive or nonexclusive; we call this the “between-advertiser competition effect,” and this is good for seller revenue. Second, competition between nonexclusive and exclusive outcomes gives an advertiser who prefers exclusive placement the incentive to reduce its bid for the multiple placement outcome because it

wants exclusive placement (with itself as the winner) to be the winning outcome; we call this the “within-advertiser competition effect,” and this is bad for seller revenue. Depending on the rules of the auction and parametric conditions, either force can dominate, and the seller’s revenue may increase or decrease by allowing exclusive placement. However, as mentioned, we also show that extensions of already popular auction mechanisms exist that guarantee weakly higher revenue with exclusive placement.

More specifically, our focal analysis is motivated by the RAIS program of Bing and Yahoo!. We study the GSP_{2D} auction, which is possibly the simplest extension of the widely used GSP mechanism that allows for exclusive placement. This mechanism was developed at, and patented by, Yahoo! (Ghosh et al. 2011a, b) and is believed to be the basis of the exclusive placement auction mechanism used in RAIS. For RAIS, exclusive placement is allowed only on trademarked keywords; that is, it is only applicable when trademarked keywords are searched, and only the advertiser that owns the trademark can bid for exclusive placement. Therefore, we also maintain this assumption for our focal analysis. Our choice is also motivated by the reason that this allows our model to be more relevant for predicting patterns that we then examine in RAIS auction data from Bing. We note that we remove this restriction in subsequent analysis.

Interestingly, we find that either the between-advertiser or the within-advertiser competition effect can prevail in GSP_{2D} ; that is, the search engine makes higher or lower revenue depending on parametric conditions. Paradoxically, the revenue is lower when bidders have high valuation for exclusive placement because this is exactly when the within-advertiser competition effect is strong. This is, of course, not a desirable property of the mechanism from the search engine’s point of view. In terms of the bidding strategies of advertisers, we find that because the trademark owner may focus on the exclusive-display outcome and shade its bid for the multiple-display outcome, the other bidders are induced to bid higher, possibly even above their valuations, because they want the multiple placement outcome to be the winning outcome. In addition, advertisers may have lower payoffs in GSP_{2D} than in GSP including the advertiser who bids for exclusivity and is placed exclusively.

We examine a data set on exclusive placement auctions in search advertising from Bing’s RAIS program. Using 39 randomly selected trademarked queries that had exclusive placement allowed, we compared advertiser bids and outcomes before and after exclusive placement was allowed, analyzing a total of over 100,000 auction instances. We found that, as predicted, advertisers who bid for exclusivity reduced their bids

for the nonexclusive outcome, while their competitors who could not bid for exclusivity increased their bids. For a little over 60% of the queries, the revenue increased after exclusivity was allowed, but for the remaining, the revenue decreased. These findings are in line with the insights provided by our theoretical framework.

Next, as a possible solution to the issue of decreasing revenue, we consider the VCG_{2D} auction, which applies the rules of the VCG auction to exclusive placement by allowing an additional exclusive placement bid for the trademark owner. The choice of a VCG-based mechanism is motivated by the fact that, besides having a long history in economics (Vickrey 1961, Clarke 1971, Groves 1973, Krishna 2010), VCG-based auctions are already in use in the industry. For instance, recently, the Russian search engine Yandex has switched to a VCG-style auction from a GSP-style auction (Russian Search Marketing 2015), and major players in the display advertising industry already use the VCG mechanism (Varian and Harris 2014). In the VCG_{2D} auction, the advertisers bid truthfully, that is, they do not shade their multiple-display bids, which eliminates the within-advertiser competition effect. Therefore, we find that the seller's revenue in the VCG_{2D} auction is weakly greater than the revenue in the GSP auction. This implies that if a seller wants to allow for exclusive placement of ads, the VCG_{2D} auction may be a good choice compared to the currently employed GSP auction.²

The rest of this paper is structured as follows. In Section 2, we discuss the literature related to our work. In Section 3, we analyze exclusive placement under the assumption that bidding for exclusivity is restricted to one "most relevant" advertiser (e.g., the trademark owner). We first describe the model and discuss the analysis and the results, and then present empirical analysis that validates our key insights. In Section 4, we analyze the case in which bidding for exclusivity is not restricted to a particular advertiser. In this scenario, we find that all of our key insights continue to hold. In Section 5, we conclude with a discussion.

2. Related Literature

In this section, we briefly discuss the related literature. Theoretical studies in economics and marketing have enhanced our understanding of position auctions used in sponsored search advertising, starting with those by Edelman et al. (2007) and Varian (2007), who showed that bidding is stable but not truthful in the widely used generalized second price auction. Various other papers that consider different aspects of second price position auctions include those by Wilbur and Zhu (2008), Katona and Sarvary (2010), Athey and Ellison (2011), Jerath et al. (2011), Zhu and Wilbur (2011), Amaldoss et al. (2015, 2016), Desai et al. (2014), Gomes

and Sweeney (2014), Sayedi et al. (2014), Lu et al. (2015), and Shin (2015). All of the above papers, however, study auctions that only consider displaying multiple advertisers in response to a keyword search.

There is a nascent literature on expressive auctions, in which advertisers can express their preferences beyond simply turning in bids for a multiple placement outcome. Muthukrishnan (2009) considers a second price auction and allows each advertiser to submit a per-click bid (its maximum willingness to pay) and specify the maximum number of other advertisers it wants to be displayed with; this is a very different auction mechanism from GSP_{2D} . Ghosh and Sayedi (2010) analyze the GSP_{2D} auction as we do; however, their focus is on comparing the worst-case properties of the multiple equilibria that the GSP and GSP_{2D} auctions can attain. By contrast, our aim is to intuitively understand the working of exclusive-display auctions to develop implications for revenue for the seller and bidding strategies for advertisers.

The literature on combinatorial auctions, in which multiple items are for sale and bidders can submit bids for combinations of items, is also related to our work. Crampton et al. (2006) provide a comprehensive survey of advances in combinatorial auctions. Because of the complex auction structure, there has not been much success in optimal mechanism design for combinatorial auctions in general.

Our paper is also related to the literature on multidimensional auctions, in which bids contain multiple attributes (Thiel 1988, Che 1993, Branco 1997, Mori 2006). For example, in an auction for a contract to build an aircraft, bidders quote a price and also specify the components of the aircraft along with the qualities of each component (Branco 1997).

Finally, we note that exclusivity contracts are often negotiated between media providers and advertisers for traditional media advertising (Dukes and Gal-Or 2003). For example, Anheuser-Busch and Volkswagen held the rights for advertising exclusively in the beer and automotive categories, respectively, during the 2011 Super Bowl. Our work is related to, but very different from, the work in this literature stream. First, the institutional details of our setting introduce several differences (e.g., ranked outcomes with position effects, per-click bidding by advertisers, etc.). Second, in our specific case, the auction mechanism allows multiple as well as exclusive winners, and the auctioneer decides after the bidders have submitted their bids whether there will be multiple winners with a rank ordering or only one winner.

3. Restricted Exclusive Placement

This section is motivated by Bing and Yahoo!'s implementation of RAIS. In particular, we assume that only one designated bidder can bid for exclusive placement.

In sponsored search advertising, this is the owner of the trademark that appears in the search term. (Note that the scope of RAIS is limited to search queries that include trademarked keywords.) An important consideration here is that we have obtained data from RAIS auctions at Bing, and results from the model with restricted bidding are more directly relevant to be tested in these data. Furthermore, this assumption facilitates exposition by reducing the number of parameters in the model. Throughout this section, we use terminology from search advertising (e.g., search engine, clicks, and click-through rates (CTRs)). This analysis is relevant to display advertising to the extent that in some cases exclusive placement may be restricted to a bidder that matches closely with the content of the page the ad will be shown on. However, in Section 4, we relax this assumption and analyze a model in which all bidders can bid to be placed exclusively, a scenario more relevant to display advertising. Using that analysis we show the robustness of the results obtained in this section.

We analyze four auction mechanisms: GSP, GSP_{2D}, VCG, and VCG_{2D}. GSP and VCG are currently the most commonly used mechanisms in search and display advertising—GSP is used by Google, Bing, and Yahoo! in their search advertising auctions, while VCG is commonly used by exchange platforms for selling display advertising impressions and is also used by the Russian search engine Yandex for selling search advertising. GSP_{2D} and VCG_{2D} are the extensions of GSP and VCG when allowing for exclusive placement. As mentioned earlier, Bing and Yahoo!’s implementation of RAIS is supposedly based on the GSP_{2D} mechanism (this mechanism has also been patented by Yahoo!).³ VCG_{2D} is a special case of the more general VCG mechanism, which, for expositional clarity, we call VCG_{2D} to emphasize that it allows exclusive placement.

3.1. Model

Page Layout. In the standard, that is, nonexclusive, outcome, there are two slots in the north section of the page, with CTRs given by θ_1 and θ_2 for the top and bottom slots, respectively. There is one slot on the east side of the page with click-through rate θ_3 . In the exclusive outcome, there is one slot on the north and one slot (with the same CTR, θ_3) on the east. Without loss of generality, through normalization, we assume that the CTR of the only north slot in the exclusive outcome is 1. We assume $1 \geq \theta_1 \geq \theta_2 \geq \theta_3 > 0$. Exclusive display in the north side is allowed only for queries that contain trademarked keywords and only for the trademark owner. For instance, if the trademark is Adidas, then the trademark owner of Adidas only can bid for exclusive display in response to any query that contains the keyword Adidas, for example, “Adidas,” “Adidas shoes,” “latest Adidas range,” etc.

The page layout in our model is motivated by Bing and Yahoo!’s implementation of RAIS, where exclusivity applies only to the north section of the page; that is, the east slot is not affected. This makes our model a closer approximation of RAIS (which is desirable, as the data that we use in this paper are from the RAIS implementation). Furthermore, the existence of the east slot also allows us to “keep the losing bidders in the game” when the outcome is exclusive, so that they still bid strategically without us making any assumptions on their behavior. Nonetheless, in Section 4, we analyze a model in which there is only one slot under exclusivity and show the robustness of our insights.

Advertiser Valuations. We assume that there are four advertisers, A, B, C, and D, with nonexclusive valuations $a \geq b \geq c \geq d$, respectively. Advertiser A is the trademark owner of the keyword and can buy exclusive rights on the north section of the search results page through RAIS. We assume that A’s valuation for being shown exclusively is e_a . The exclusive valuations of the other advertisers are irrelevant because, per Bing’s implementation rules, they are not allowed to bid for exclusivity, as they are not trademark holders for the keyword; for this reason, for notational simplicity we denote e_a simply by e . We allow e to be higher or lower than a . Advertiser A submits two bids, a' and e' , for the nonexclusive and the exclusive outcomes, and advertisers B, C, and D submit bids b' , c' , and d' for the nonexclusive outcome, respectively. Note that we have assumed that advertiser A’s valuation for nonexclusive display is the highest among all advertisers.⁴

The model can be readily reinterpreted to appeal to display advertising. In the context of display advertising, θ_i represents the amount of attention that advertising slot i gets when the outcome is nonexclusive. Typically, advertising slots that are larger or are located on top of a page draw more attention than smaller slots or those that are located below the fold. The amount of attention that the exclusive slot gets is normalized to 1. Advertisers’ valuations correspond to advertisers’ willingness to pay for one unit of attention.

Equilibrium Selection. Following Edelman et al. (2007) and Varian (2007), we use the lowest revenue envy-free (LREF) equilibrium as a refinement for equilibrium selection, where the “envy-free” requirement implies that any advertiser does not want to exchange positions with advertisers above and below it.

We assume that the valuations of all advertisers are common knowledge. This is a common assumption in the literature on sponsored search auctions, justified by the argument that these auctions are run continually, and therefore advertisers and the search engine have ample opportunity to learn about each other’s valuation.⁵ We now proceed to the analysis of the model.

3.2. GSP and GSP_{2D} Auctions

In this section, we analyze the GSP and GSP_{2D} auctions in detail and compare their outcomes. We then verify the empirical patterns predicted by our theory in data obtained from Bing.

3.2.1. The GSP Auction. In the GSP auction, positions are allocated in decreasing order of bids, and each advertiser pays the bid of the advertiser directly below it (and the advertiser at the bottom of the ladder pays the highest losing bid). The envy-free condition requires the order of the bids to be in the same order as the valuations, that is, $a' \geq b' \geq c' \geq d'$. Furthermore, it requires the losing advertiser to bid at least its true valuation, and since we choose the LREF equilibrium, it implies that the losing advertiser bids truthfully.⁶ In this case, D is the losing advertiser, which implies that $d' = d$.

Advertiser C's bid should be such that he will not envy advertiser B. This means $\theta_2(c - c') \leq \theta_3(c - d)$, which implies

$$c' \geq \underline{c} \equiv \frac{(\theta_2 - \theta_3)c + \theta_3 d}{\theta_2}.$$

Similarly, the bid of advertiser B has to be such that he will not envy advertiser A. This means that $\theta_1(b - b') \leq \theta_2(b - c')$, which implies

$$b' \geq \underline{b} \equiv \frac{(\theta_1 - \theta_2)b + \theta_2 c'}{\theta_1}.$$

In the LREF equilibrium, advertiser B bids \underline{b} , and advertiser C bids \underline{c} . Advertiser A can bid anything above \underline{b} . The following proposition describes the equilibrium for the GSP auction.

Proposition 1 (GSP Equilibrium). *In equilibrium, advertisers A and B are in the north slots, and C is in the east slot. The bids of advertisers B, C, and D are $\underline{b} = ((\theta_1 - \theta_2)b + \theta_2 \underline{c})/\theta_1$, $\underline{c} = ((\theta_2 - \theta_3)c + \theta_3 d)/\theta_2$, and d , respectively. Advertiser A can bid anything above \underline{b} ; that is, the bid of advertiser A is $\underline{b} + \epsilon$, $\epsilon > 0$. Advertisers A, B, and C pay $\theta_1 \underline{b}$, $\theta_2 \underline{c}$, and $\theta_3 d$, respectively. The search engine's revenue is $R_{\text{GSP}} = (\theta_1 - \theta_2)b + 2(\theta_2 - \theta_3)c + 3\theta_3 d$.*

We note that the analysis of GSP under LREF refinement already exists in Edelman et al. (2007). Nonetheless, we present this result for the sake of completeness, and to make it easier for the reader to compare the analysis of GSP to that of GSP_{2D}.

3.2.2. The GSP_{2D} Auction. *Auction Definition.* Assuming that nonexclusive bids are $b_1 \geq b_2 \geq b_3 \geq b_4$ and the exclusive bid is b_e , the search engine chooses the exclusive outcome if and only if $b_e > \theta_1 b_2 + \theta_2 b_3$. If the outcome is exclusive, A pays $\theta_1 b_2 + \theta_2 b_3$ per click, and the east slot is allocated to the bidder with the highest nonexclusive bid among B, C, and D, at the cost per

click of the second highest nonexclusive bid. On the other hand, if the outcome is nonexclusive, payment and allocation will be the same as in GSP, in which positions are allocated in decreasing order of bids, and each advertiser pays the bid of the advertiser directly below it (and the advertiser at the bottom of the ladder pays the highest losing bid).

Note that the rule $b_e > \theta_1 b_2 + \theta_2 b_3$ for choosing the outcome as exclusive or nonexclusive ignores the east slot. This is because only the north section of the search results page is being made exclusive, and it is the rule that Bing uses in its implementation (i.e., Bing chooses the outcome based only on the impact on revenue from the north section). We have also considered and analyzed a rule that incorporates impact on the east slot as well to choose the auction outcome. We obtain qualitatively similar results in that formulation. Specifically, including the east slot in the revenue would make the exclusive outcome even more likely to be chosen—since one of our primary counterintuitive results is that using exclusive auctions could decrease the revenue, the assumption of using only the north section to decide the outcome is a conservative assumption in that regard.

Exclusive Outcome. In GSP_{2D}, if the outcome is exclusive, advertiser A wins the north slot, and advertiser B wins the east slot. In the LREF equilibrium, advertisers C and D bid $c' = c$ and $d' = d$, respectively; that is, losing advertisers bid truthfully. Advertiser B bids $b' = c + \epsilon$. (Any bid above c gives him the same outcome and does not affect any other advertiser's price or allocation.) advertiser A bids $a' = 0$ for the nonexclusive outcome and $e' = e$ (i.e., bids truthfully) for the exclusive outcome.⁷ When advertiser A's value for exclusivity is sufficiently high, equilibrium outcome will be exclusive. In this case, advertiser A pays $\theta_1 c' + \theta_2 d' = \theta_1 c + \theta_2 d$ and advertiser B, who wins the east slot, pays $\theta_3 c' = \theta_3 c$. The search engine's revenue is

$$R_{\text{GSP}_{2D}}^E = \theta_1 c + \theta_2 d + \theta_3 c.$$

This revenue may be higher or lower than in GSP. In particular, the revenue in GSP_{2D} is higher if and only if $R_{\text{GSP}_{2D}}^E > R_{\text{GSP}}$, that is,

$$\theta_1 c + \theta_2 d + \theta_3 c > (\theta_1 - \theta_2)b + 2(\theta_2 - \theta_3)c + 3\theta_3 d.$$

Using this inequality, we can see that the search engine's revenue in GSP_{2D}, compared with that in GSP, increases if θ_1 decreases, b decreases, or θ_3 increases. The effect of θ_2 depends on other parameters: if (and only if) $b + d > 2c$, then an increase in θ_2 also increases the appeal of GSP_{2D}.

Nonexclusive Outcome. In GSP_{2D}, for the outcome to be nonexclusive, we need a few conditions. First,

as discussed in the GSP case, for advertiser B to not envy advertiser A, we need the condition

$$b' \geq \underline{b} \equiv \frac{(\theta_1 - \theta_2)b + \theta_2 c'}{\theta_1},$$

and for advertiser C to not envy advertiser B, we need the condition

$$c' \geq \underline{c} \equiv \frac{(\theta_2 - \theta_3)c + \theta_3 d}{\theta_2}.$$

Furthermore, advertiser A should not benefit from increasing its exclusive bid e' (and dropping its nonexclusive bid a') to change the outcome to exclusive. This means that we have

$$\theta_1(a - b') \geq e - (\theta_1 c' + \theta_2 d).$$

This inequality holds if b' is sufficiently small or c' is sufficiently large. We already have \underline{b} as a lower bound on b' ; to understand when the inequality holds, we need an upper bound on c' . For this, we recognize that advertiser C's bid should be low enough such that advertiser B does not want to move from the second slot to the third slot (and does not envy the advertiser in the third slot). In other words, $\theta_2(b - c') \geq \theta_3(b - d)$, which implies

$$c' \leq \bar{c} \equiv \frac{(\theta_2 - \theta_3)b + \theta_3 d}{\theta_2}.$$

By letting $b' = \underline{b}$ and $c' = \bar{c}$, we get

$$e \leq L_1 \equiv \theta_1 a + \theta_2 d - \frac{\theta_3}{\theta_2}(\theta_1 - \theta_2)(b - d).$$

In other words, if $e > L_1$, a nonexclusive equilibrium does not exist in GSP_{2D} for any values of b' and c' that satisfy other equilibrium conditions (i.e., $b' \geq \underline{b}$ and $c' \leq \bar{c}$). However, an exclusive equilibrium always exists (and this is the one characterized earlier).

If $e \leq L_1$, both exclusive and nonexclusive equilibria could exist (as long as e is still sufficiently large, that is, $e \geq \theta_1 c + \theta_2 d$; otherwise it would always be nonexclusive⁸). In this case, we select the nonexclusive equilibrium because advertiser B prefers the nonexclusive outcome, he is indifferent between all bids greater than c in the exclusive outcome, and by bidding \underline{b} he can allow advertiser C to change the outcome to nonexclusive (if advertiser C bids \bar{c}). In other words, for any value of c' where $\underline{c} \leq c' \leq \bar{c}$, bidding \underline{b} is a weakly dominant strategy for advertiser B. If advertiser B always bids \underline{b} , advertiser C can change the outcome to nonexclusive by bidding \bar{c} . Since both advertisers prefer the nonexclusive outcome, it is reasonable to select the nonexclusive outcome in equilibrium refinement.

Next, we show that when e is sufficiently small, GSP_{2D} and GSP lead to the same equilibrium outcome.

Using Proposition 1, we know that advertiser B bids \underline{b} and advertiser C bids \underline{c} in equilibrium in GSP. If, with those bids, advertiser A prefers the nonexclusive outcome, then the equilibrium of GSP_{2D} will be the same as in GSP. In other words, the equilibrium of GSP_{2D} will be the same as in GSP if

$$\theta_1(a - b') \geq e - (\theta_1 c' + \theta_2 d)$$

for $c' = \underline{c}$ and $b' = \underline{b}$, which simplifies to

$$e \leq L_2 \equiv \theta_1(a - b + c) + \theta_2(b - c + d) - \frac{\theta_3}{\theta_2}(\theta_1 - \theta_2)(c - d).$$

Therefore, for $e < L_2$ (i.e., when e is sufficiently small), introducing exclusivity has no effect on the advertisers' strategies. In this case, the equilibrium is as described in Proposition 1, and the search engine's revenue is $R_{\text{GSP}_{2D}}^{N, e < L_2} = R_{\text{GSP}}$.

When $L_2 < e \leq L_1$, the equilibrium outcome is still nonexclusive. However, equilibrium bids are slightly different. Advertiser C's bid, c' , has to be such that advertiser A cannot benefit from changing the outcome to exclusive. In other words, c' must be large enough such that $\theta_1(a - b') \geq e - (\theta_1 c' + \theta_2 d)$. This simplifies to

$$c' \geq \hat{c} \equiv b + d + \frac{e - \theta_1(a + d)}{\theta_1 - \theta_2}.$$

Since advertiser C increases its bid from \underline{c} to \hat{c} to maintain the nonexclusive outcome, advertiser B also increases its bid (by a smaller amount) to ensure envy-freeness (note that \underline{b} is an increasing function of c'). advertiser B bids

$$\hat{b} \equiv b + \frac{\theta_2(e - \theta_1 a - \theta_2 d)}{\theta_1(\theta_1 - \theta_2)}.$$

Finally, advertiser A's nonexclusive bid can be anything above \hat{b} , and does not affect any of the payments, allocations, or revenue. Since advertiser A prefers the nonexclusive outcome, its exclusive bid will be low enough so that the exclusive outcome does not happen in equilibrium. The search engine's revenue in this case will be $\theta_1 \hat{b} + \theta_2 \hat{c} + \theta_3 d$, which simplifies to

$$R_{\text{GSP}_{2D}}^{N, L_2 < e \leq L_1} = (\theta_1 + \theta_2)b + \frac{2\theta_2(e - \theta_1 a - \theta_2 d)}{\theta_1 - \theta_2} + \theta_3 d.$$

When $L_2 < e \leq L_1$, we have $R_{\text{GSP}_{2D}}^{N, L_2 < e \leq L_1} > R_{\text{GSP}}$, which means that existence of exclusivity increases the search engine's revenue, even though exclusivity does not happen in equilibrium. This is because advertiser C increases its bid to make the exclusive outcome less appealing for advertiser A. Consequently, advertiser B, who now has to pay more because of advertiser C, also increases its bid (by a smaller amount) to increase the payment of advertiser A, so that he will not envy advertiser A. Note that advertiser A ends up paying

more than what he was paying in GSP. In other words, advertiser A's moderately high valuation for exclusivity can hurt him even when the exclusive outcome does not happen in equilibrium (i.e., advertiser A's profit decreases with e when $e \in [L_1, L_2]$).

Based on the above analysis, the following proposition describes the equilibrium for GSP_{2D} .

Proposition 2 (GSP_{2D} Equilibrium). *The equilibrium of GSP_{2D} is the following:*

- For $e \leq L_2 \equiv \theta_1(a - b + c) + \theta_2(b - c + d) - (\theta_3/\theta_2) \cdot (\theta_1 - \theta_2)(c - d)$, advertiser A does not bid for exclusivity, and the rest of the outcome is exactly the same as the outcome of GSP described in Proposition 1.
- For $L_2 < e \leq L_1 \equiv \theta_1 a + \theta_2 d - (\theta_3/\theta_2)(\theta_1 - \theta_2)(b - d)$, the bids of advertisers B, C, and D are $\hat{b} = b + (\theta_2(e - \theta_1 a - \theta_2 d))/(\theta_1(\theta_1 - \theta_2))$, $\hat{c} = b + d + e - (\theta_1(a + d))/(\theta_1 - \theta_2)$, and \hat{d} , respectively. The bid of advertiser A is $\hat{b} + \epsilon$, $\epsilon > 0$, for the nonexclusive outcome and low enough for the exclusive outcome such that the exclusive outcome does not happen (we assume this to be 0). The outcome is nonexclusive, with A and B in the north slots and C in the east slot. Advertisers A, B, and C pay $\theta_1 \hat{b}$, $\theta_2 \hat{c}$, and $\theta_3 \hat{d}$, respectively. The search engine's revenue is $R_{\text{GSP}_{2D}}^N = (\theta_1 + \theta_2)\hat{b} + (2\theta_2(e - \theta_1 a - \theta_2 d))/(\theta_1 - \theta_2 + \theta_3 d)$.
- For $e > L_1$, the bids of advertiser A are e and 0 for the exclusive and nonexclusive outcomes, respectively. The bids of advertisers B, C, and D are $c + \epsilon$, c , and d , respectively, where $\epsilon > 0$. The outcome is exclusive, with A in the north slot and B in the east slot. Advertiser A pays $\theta_1 c + \theta_2 d$ and B pays $\theta_3 c$. The search engine's revenue is $R_{\text{GSP}_{2D}}^E = \theta_1 c + \theta_2 d + \theta_3 c$.

The revenue of GSP_{2D} can decrease in e ; specifically, there is a discrete downward jump in revenue at $e = L_1$.

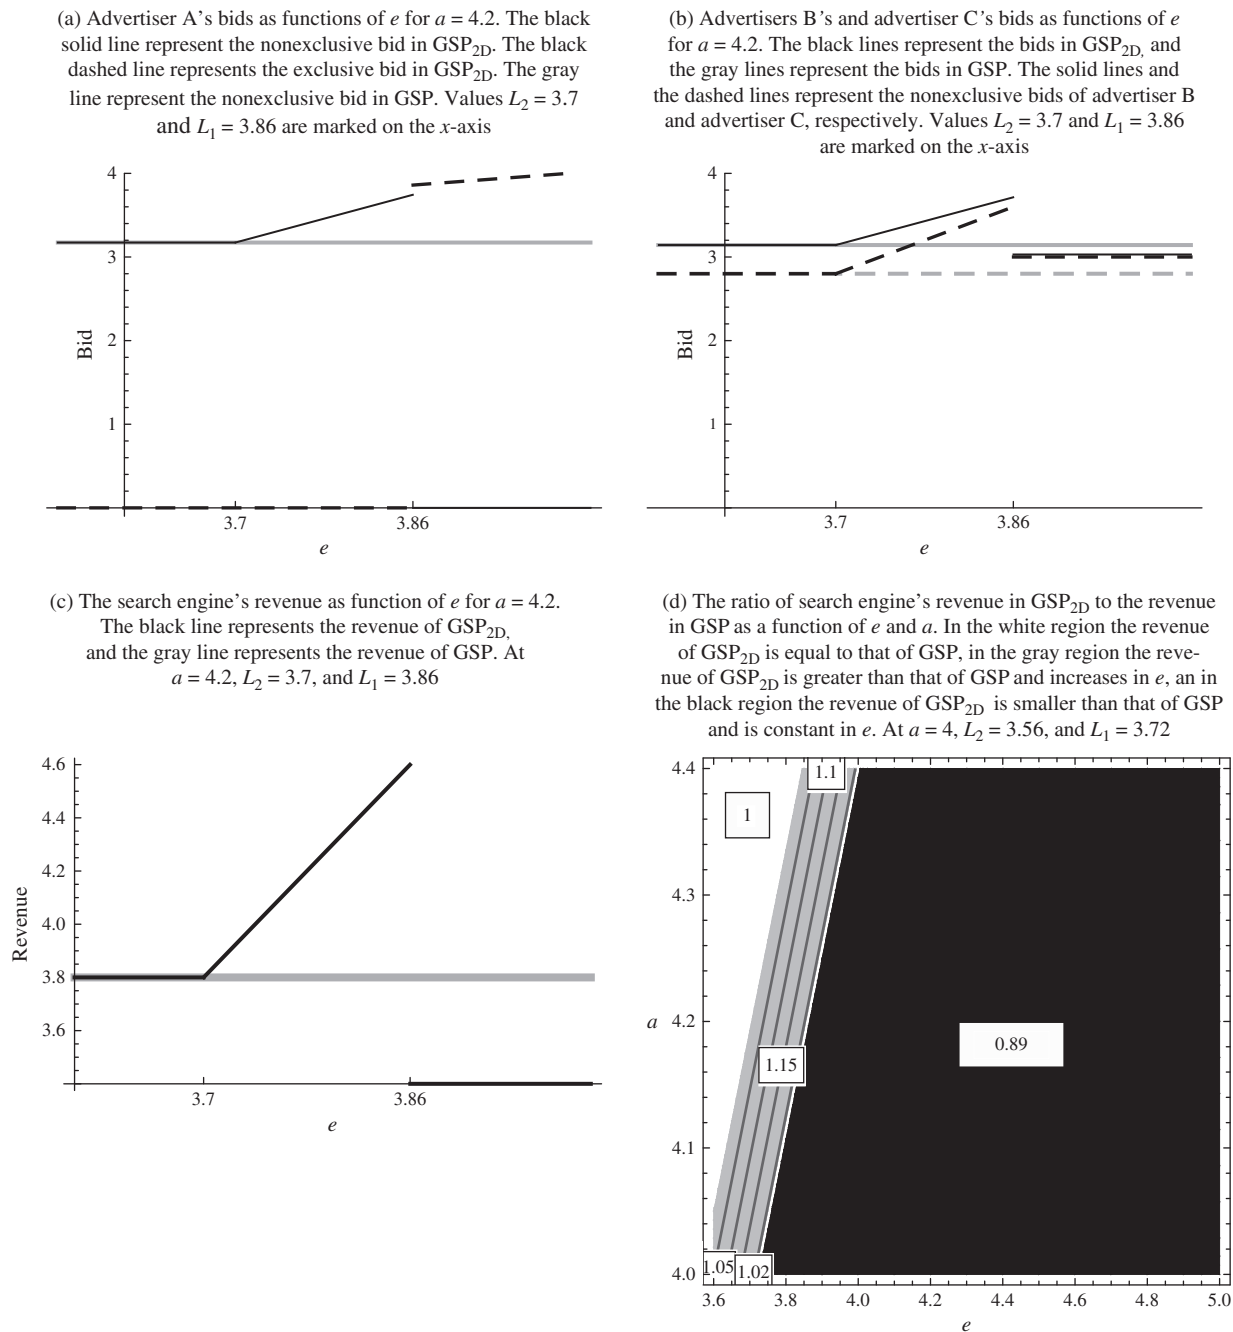
In the above equilibrium, we observe that advertisers have the incentive to bid high because they compete not only for positions in the nonexclusive outcome but also for the outcome to be exclusive or nonexclusive; we call this the between-advertiser competition effect. However, there is a countervailing force to this—Advertiser A has the incentive to bid low for the exclusive outcome when it bids high for the nonexclusive outcome and vice versa; that is, there is a downward pressure on bids because of the competition between the two display formats; we call this the within-advertiser competition effect. In fact, this effect can be strong enough that a counterintuitive implication that we obtain is that the revenue of GSP_{2D} can decrease with increasing exclusive-display valuation because advertiser A significantly reduces its nonexclusive bid.

Proposition 3 (Search Engine Revenue Comparison). *If $e \leq L_2$, then GSP_{2D} has the same equilibrium as in GSP. If $L_2 < e \leq L_1$, then the outcome of GSP_{2D} is nonexclusive, however, the search engine has a higher revenue in GSP_{2D} ,*

and the revenue is an increasing function of e . If $e > L_1$, then the outcome of GSP_{2D} is exclusive. In this case, the search engine's revenue in GSP_{2D} might be higher or lower than that of GSP, depending on θ_i 's and advertisers' valuations.

Panels (a) and (b) in Figure 2 show the advertisers' bids (for clarity, we show the bids of advertiser A in Figure 2(a), and those of advertisers B and C in Figure 2(b); the axes for both plots are identical), and panel 2(c) shows the search engine's revenues for GSP and GSP_{2D} as functions of e . First, note that advertiser A bids 0 for exclusivity for $e \leq L_1$ and bids truthfully for exclusivity for $e > L_1$. Next, we turn to the nonexclusive bids. If the value of e is sufficiently small ($e \leq L_2$), the nonexclusive bids (and the placement and search engine's revenue) in GSP_{2D} are the same as those in GSP. For a medium value of e ($L_2 < e \leq L_1$), the nonexclusive bids in GSP_{2D} increase with e . Note that in this region the placement is multiple display, and the search engine's revenue is higher in GSP_{2D} even though the placement is the same as in GSP. This is because advertisers B and C want to keep the outcome nonexclusive and, therefore, they bid higher than in GSP. Note that advertiser C (with valuation assumed to be 3 for the purposes of the figure) bids more than its valuation in a part of this region. When e is sufficiently large ($e > L_1$), the placement is exclusive, and the nonexclusive bids remain constant in e . Note that the nonexclusive bid of advertiser A in this region is 0, and the search engine's revenue is lower in GSP_{2D} than in GSP. Finally, note that the search engine's revenue is nonmonotonic in e in GSP_{2D} (specifically, there is a discrete downward jump in revenue at $e = L_1$).

Figure 2(d) illustrates how the search engine's revenue compares in GSP_{2D} and GSP with respect to e and a , that is, the exclusive and nonexclusive valuations, respectively, of advertiser A. It plots the ratio of the search engine's revenue in GSP_{2D} to that in GSP. The white region is where the two mechanisms have the same revenue (ratio = 1) and corresponds to $e \leq L_2$. The gray region in the middle is where the revenue of GSP_{2D} is higher than the revenue of GSP (ratio > 1), and corresponds to $L_2 < e \leq L_1$. Finally, the black region (ratio = 0.89) is where the revenue of GSP is higher than the revenue of GSP_{2D} , and corresponds to $e > L_1$. For a fixed value of a (say 4.2), first, for small values of e , the two mechanisms provide the same revenue; then, for medium values of e , GSP_{2D} provides higher revenue than GSP, and finally, for large values of e , GSP_{2D} provides lower revenue than GSP. In fact, for the parameter values used in Figure 2, for all values of e such that $e \geq a$ (i.e., advertiser A values exclusivity more than nonexclusivity, which is reasonable to expect), the revenue of GSP_{2D} is lower than that of GSP. Intuitively, one would expect GSP_{2D} to outperform GSP for high values of e in terms of revenue; that is, if there is an advertiser who values exclusivity highly, then the auction that allows exclusive display should provide

Figure 2. Comparison of Advertisers' Bids and Search Engine's Revenue in GSP and GSP_{2D}

Note. The values of the other parameters are $b = 4$, $c = 3$, $d = 2$, $\theta_1 = 0.7$, $\theta_2 = 0.5$, and $\theta_3 = 0.1$. For these values, $L_2 = 0.76 + 0.7a$ and $L_1 = 0.92 + 0.7a$ are functions of a .

high revenue. We can see, however, that this is not the case for GSP_{2D} (for $e > L_1$). The reason is that, because of the within-advertiser competition effect, advertiser A lowers its nonexclusive bid very significantly (in this case, to 0) to make sure that the exclusive outcome is selected in this region.⁹

Proposition 4 (Advertiser Payoff Comparison). When advertiser A's valuation for exclusivity is sufficiently low, $e \leq L_2$, all advertisers' profits are the same in GSP_{2D} as in

GSP. When advertiser A's valuation for exclusivity is moderately high ($L_2 < e \leq L_1$), the equilibrium outcome is nonexclusive. However, advertisers A and B have lower payoff in GSP_{2D} than in GSP, and their payoffs are decreasing functions of e . Advertiser C has the same payoffs in GSP_{2D} and GSP in this case. When a advertiser A's valuation for exclusivity is sufficiently high ($e > L_1$), advertisers B and C have lower payoffs in GSP_{2D} than in GSP. Advertiser A may have higher or lower payoff, depending on the values of other parameters.

Proposition 4 shows that, since GSP_{2D} increases the between-advertiser competition, it could lead to higher nonexclusive bids, which in turn could lower the advertisers' equilibrium payoffs.

3.2.3. Empirical Support. In this section, we provide empirical support for a number of sharp predictions regarding the GSP and GSP_{2D} auctions from our theoretical model. We use a data set from RAIS auctions for search ads at Bing. We note that our aim in this section is to provide directional empirical validation of the predictions of our model, and we do not claim to provide conclusive evidence or conduct scientific hypothesis testing.

To conduct our analysis, we collected data on 100 randomly chosen queries at Bing that contain trademarked keywords that have at least 100 instances of exclusive display on July 15, 2015, but were not enrolled in RAIS on July 15, 2014. By comparing outcomes on these two days, we can assess differences in behavior due to the availability of the exclusive-display option under RAIS. We note that July 15 was a randomly chosen date, and we compared data for this date on two days one year apart to minimize seasonality effects. From our initial set of 100 queries, we removed those that were synonyms, those that had no competition, and those for which advertiser sets were very different in 2014 and 2015. We also removed the queries for which the advertiser ID of the trademark owner changed from 2014 to 2015.¹⁰ After this trimming, we were left with 39 queries, spanning 114,052 auction instances, on which we conducted our analysis. These 39 queries are from a wide range of industries including retail, finance, travel, and education. Using these data, we are able to find support for the following predictions from our model.

First, our model suggests that the advertisers that bid for being displayed exclusively in response to queries containing their trademarked keywords have an incentive to decrease their bids for the multiple placement outcome. In accordance with this, we find that, averaged across all auctions for the 39 keywords, advertisers that submit exclusive-display RAIS bids in 2015 decrease their multiple placement bids by 46% compared to 2014. This highlights the insight that advertisers that own the trademark, who could be expected to be among the highest bidders in a multiple placement format, bid for being displayed exclusively but at the same time might reduce their bids for being displayed with multiple others. This can reduce the search engine's revenue.

Second, our model suggests that the advertisers who cannot or do not want to bid for exclusive display will increase their bids for the multiple placement outcome because they want the multiple placement outcome to be the winning outcome. In accordance with this, we find that, averaged across all auctions for the 39 keywords, advertisers that do not bid for exclusive display

in 2015 increase their multiple placement bids by 33% compared to 2014. This is good for the search engine's revenue.

Taken together, the first two predictions above do not give a clear picture of whether the search engine's revenue will increase or decrease on using GSP_{2D} for the RAIS system. Indeed, we find that in our data, for 24 out of 39 queries, the revenue increased, and for 15 queries, the revenue decreased.

3.3. VCG and VCG_{2D} Auctions

In this section, we analyze the VCG and VCG_{2D} auctions. We compare their outcomes to each other and to the GSP-based auctions.

3.3.1. The VCG Auction. The Vickrey–Clarke–Groves auction is a widely studied and applied auction that generalizes the basic idea of a second-price auction (also known as a Vickrey auction). This auction has attractive theoretical properties such as inducing truthful bidding by bidders and maximizing social welfare (which are not properties of the GSP and GSP_{2D} mechanisms). Recently, the Russian search engine Yandex has adopted a VCG-style auction (Russian Search Marketing 2015). In this section, we use the VCG mechanism's allocation and payment rule for the position auction with exclusive display allowed only for the trademark owner. Intuitively, the VCG mechanism uses the welfare-maximizing allocation and charges each bidder the "harm" that the bidder's presence causes to the other bidders.

Since truthful bidding is a weakly dominant strategy in the VCG auction, we have $a' = a$, $b' = b$, $c' = c$, and $d' = d$; therefore, advertisers A, B, and C get the first, second, and third slots, respectively. The payments of the advertisers are calculated in Section A.1 in the appendix. The search engine's revenue is the sum of the payments of advertisers A, B, and C, which is given by $R_{VCG} = (\theta_1 - \theta_2)b + 2(\theta_2 - \theta_3)c + 3\theta_3d$. We note that all advertisers' payments and allocations in VCG are the same as those in GSP.¹¹

3.3.2. The VCG_{2D} Auction. The VCG_{2D} mechanism uses the same general allocation and payment rules as those in the VCG auction, but it allows advertiser A to bid for exclusivity. Although the VCG_{2D} auction is a special case of the VCG mechanism, for expositional clarity, we call it the VCG_{2D} auction to emphasize that it allows exclusive display. Given the desirable properties of the VCG auction, this analysis is of interest, though we note that, to our knowledge, the VCG_{2D} auction is not being used by any search engine at the moment.

We provide the analysis of the auction in Section A.2 in the appendix. Here, we provide the results and insights from this analysis. Let $K_3 = \theta_1a + (\theta_2 - \theta_3)c + \theta_3d$, $K_2 = \theta_1a + (\theta_2 - \theta_3)b + \theta_3d$, and $K_1 = \theta_1a + (\theta_2 - \theta_3)b + \theta_3c$. Then, we obtain the following proposition.

Proposition 5 (VCG_{2D} Equilibrium). *The equilibrium of the VCG_{2D} auction is the following:*

- If $e \leq K_3$, the outcome is nonexclusive. The equilibrium payment and allocation of VCG_{2D} are exactly the same as those in GSP for all advertisers.
- If $K_3 < e \leq K_2$, the outcome is nonexclusive. Advertiser B has to pay a higher price for the same allocation as in GSP; furthermore, its payment is an increasing function of e . The payments of advertisers A and C will be the same as those in GSP.
- If $K_2 < e \leq K_1$, the outcome is nonexclusive. Advertisers B and C have to pay a higher price for the same allocation as in GSP; furthermore, their payments are increasing functions of e . The payment of advertiser A remains the same as that in GSP.
- Finally, if $e > K_1$, the outcome is exclusive. The revenue of VCG_{2D} weakly increases in e .

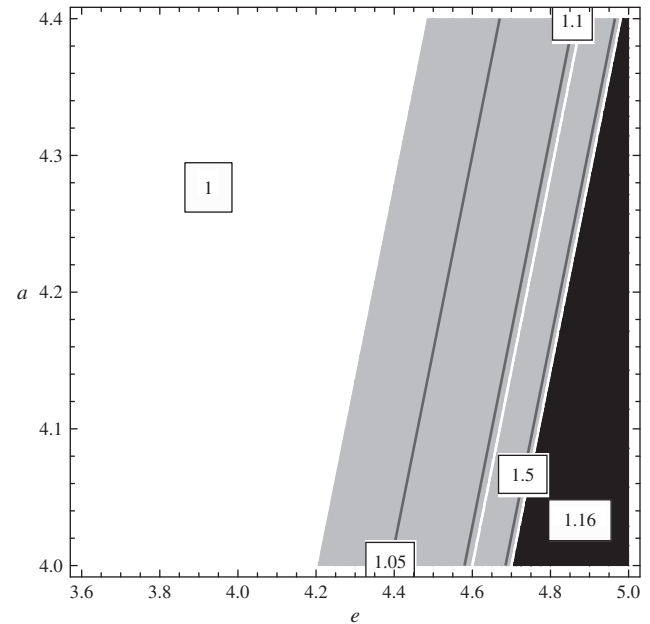
As in GSP_{2D}, when the valuation for exclusive placement of advertiser A is small enough, that is, $e \leq K_1$, the outcome is nonexclusive display, and when this valuation is large enough, that is, $e > K_1$, the outcome is exclusive display. However, unlike GSP_{2D}, the revenue of the search engine is (weakly) increasing in e . As long as the outcome is nonexclusive, it is clear that the revenue does not go down in e as the payments of advertisers B and C could only increase with e , while the payment of advertiser A is unchanged. When the outcome is exclusive, the revenue is constant in e .

Proposition 6 (Revenue of VCG_{2D} vs. GSP). *The search engine's revenue in VCG_{2D} is always greater than or equal to that of GSP; that is, introducing exclusivity with VCG_{2D} will (weakly) increase revenue.*

An important takeaway from Proposition 6 is that the search engine's revenue is always higher in VCG_{2D} than in GSP, which cannot be guaranteed for the GSP_{2D} auction, as our previous analysis shows. This is shown in Figure 3, which plots the ratio of revenue of VCG_{2D} to the revenue of GSP with respect to e and a (for specific values of the other parameters). Note that for the parameter values used in Figure 3, for values of e such that $e \geq a$ (i.e., advertiser A values exclusivity more than nonexclusivity, which is reasonable to expect), the revenue of VCG_{2D} is higher than that of GSP, except in a small region where they are equal. In other words, compared to the status quo of using GSP, incorporating exclusivity with VCG_{2D} does not have a downside for the search engine in terms of revenue.¹²

Finally, we know from Edelman et al. (2007) that under the LREF refinement, VCG and GSP have the same equilibrium outcome. Therefore, Proposition 6 implies that incorporating exclusivity also (weakly) increases the revenue of a search engine that is currently using VCG (e.g., Yandex).

Figure 3. The Ratio of the Search Engine's Revenue in VCG_{2D} to the Revenue in GSP as a Function of e and a



Notes. In the white region, the revenue of VCG_{2D} is equal to that of GSP; in the gray region, the revenue of VCG_{2D} is greater than that of GSP and increases in e ; and in the black region, the revenue of VCG_{2D} is greater than that of GSP and is constant in e . The values of the other parameters are $b = 4$, $c = 3$, $d = 2$, $\theta_1 = 0.7$, $\theta_2 = 0.5$, and $\theta_3 = 0.1$ (these are the same values as for Figure 2). For these values, $K_1 = 1.9 + 0.7a$, $K_2 = 1.8 + 0.7a$, and $K_3 = 1.4 + 0.7a$ are functions of a , and at $a = 4$, $K_1 = 4.7$, $K_2 = 4.6$, and $K_3 = 4.2$. The range of e on the x -axis is chosen to be the same as that in Figure 2(c).

Corollary 1 (Revenue of VCG_{2D} vs. VCG). *The search engine's revenue in VCG_{2D} is always greater than or equal to that of VCG; that is, introducing exclusivity with VCG_{2D} will (weakly) increase revenue.*

Remark on Equilibrium Refinement. The LREF equilibrium refinement has been widely used in the previous literature on search advertising, and has been shown to be the unique outcome of simple dynamic strategies (Cary et al. 2008, Decarolis et al. 2017). Furthermore, the equivalence between the LREF equilibrium of GSP and the outcome of VCG allows us to use VCG as a benchmark for our analysis in 1D and 2D settings. However, we note that the two main forces of our model, namely, the within-advertiser competition effect and the between-advertiser competition effect, continue to hold under many other equilibrium refinements. For example, if we use the highest-revenue envy-free equilibrium, advertiser A continues to lower its bid for its nonpreferred outcome in GSP_{2D}, which in turn could lead to a lower search engine revenue in GSP_{2D} than in GSP.

To close this section, we consider the question of what a search engine that is currently using the GSP mechanism with multiple placement should do if it wants to introduce exclusive placement. Employing

the GSP_{2D} auction (or a variant thereof) entails minimal changes in the nature and rules of the auction (i.e., it stays as a second-price auction), but the search engine runs the risk of reduced revenue, as we saw in the case of Bing. Surprisingly, revenue can be lower when the exclusive bidder has a high value for exclusivity, which is exactly the opposite of what the auctioneer would expect and desire. Therefore, if using GSP_{2D}, the search engine should introduce exclusive placement only if the within-advertiser competition effect is small and the between-advertiser competition effect dominates; however, this may not be easy to determine a priori, and it is cumbersome from a practical point of view to have different auction mechanisms under different conditions. On the other hand, the search engine could use the VCG_{2D} auction, which would guarantee that it obtains (weakly) higher revenue. While this entails a substantial change in the nature and rules of the auction, the Russian search engine Yandex switched to a VCG-style auction from a GSP-style auction in 2015 (Russian Search Marketing 2015), and Varian and Harris (2014) state that search engines such as Google have also contemplated switching to a VCG-style auction from the current GSP auction.

4. Unrestricted Exclusive Placement

In the previous section, we assumed that only one advertiser, the trademark owner, can bid for exclusivity. In this section, we relax this assumption by allowing all advertisers to bid for exclusivity. This is particularly relevant for display advertising settings (see Figure 1, (b) and (c)) in which the page is not associated with any specific keyword, and therefore there is no trademark owner (or another preferred or designated bidder who alone can bid for exclusivity). We establish the robustness of our results from Section 3 and derive new insights as well. The organization of this section is similar to that of Section 3—first, we specify the model; then we analyze GSP-based auctions; and then we analyze the VCG-based auctions and compare outcomes across auctions. The details of the analysis of this section are available in the online appendix.

4.1. Model

Since allowing to bid for exclusivity for all advertisers adds many cases to the analysis, we simplify the model of Section 3 by assuming that there are two slots and three advertisers.

Page layout. In the standard, that is, nonexclusive, outcome, there are two slots in the page with click-through rates θ_1 and θ_2 . In the exclusive outcome, there is one slot on the page. Without loss of generality, through normalization, we assume that the CTR of the only slot in the exclusive outcome is 1. We assume $1 \geq \theta_1 \geq \theta_2 > 0$. In the display advertising context, the click-through rates represent the amount of attention each slot draws.

Advertiser valuations. We assume that there are three advertisers, A, B, and C. The per-click valuations of advertisers A, B, and C are a , b , and c for the nonexclusive outcome and e_a , e_b , and e_c for the exclusive outcome, respectively. Without loss of generality, we assume that $a \geq b \geq c$. We also assume that an advertiser's valuation for the exclusive outcome is weakly higher than that for the nonexclusive outcome, that is, $e_a \geq a$, $e_b \geq b$, and $e_c \geq c$.¹³ We do not make any assumption on the order of e_a , e_b , and e_c .

Equilibrium selection. As in Section 3, we use the LREF equilibrium as a refinement for equilibrium selection, where the envy-free requirement implies that any advertiser does not want to exchange positions with advertisers above and below it.

We now proceed to the analysis of the model.

4.2. GSP and GSP_{2D} Auctions

4.2.1. The GSP Auction. The GSP auction does not allow for the exclusive outcome. The analysis is similar to the previous analysis and is presented in the online appendix. The search engine's equilibrium revenue is given by $R_{\text{GSP}} = \theta_1 b' + \theta_2 c' = (\theta_1 - \theta_2)b + 2\theta_2 c$.

4.2.2. The GSP_{2D} Auction. *Auction Definition.* Assuming that the nonexclusive bids are $b_1 \geq b_2 \geq b_3$ and the exclusive bids are b_{e_1} , b_{e_2} , and b_{e_3} , the search engine chooses the exclusive outcome if and only if $\max(b_{e_1}, b_{e_2}, b_{e_3}) > \theta_1 b_2 + \theta_2 b_3$. If the outcome is exclusive, the advertiser with the highest exclusive bid pays (per click) the maximum of $\theta_1 b_2 + \theta_2 b_3$ and the second-highest exclusive bid. On the other hand, if the outcome is nonexclusive, payment and allocation will be the same as in GSP, in which positions are allocated in decreasing order of bids and each advertiser pays the bid of the advertiser directly below it (and the advertiser at the bottom of the ladder pays the highest losing bid).

We need to consider three cases: $e_a \geq \max(e_b, e_c)$, $e_b \geq \max(e_a, e_c)$, and $e_c \geq \max(e_a, e_b)$. We present the complete analysis of these three cases in the online appendix. Here we discuss the case when $e_a \geq \max(e_b, e_c)$, which, given that $a \geq \max(b, c)$, is arguably the most reasonable case. The following proposition summarizes the LREF equilibrium of GSP_{2D} when $e_a \geq \max(e_b, e_c)$.

Proposition 7 (GSP_{2D} Equilibrium When $e_a \geq \max(e_b, e_c)$). Let $z = \max(e_b - \theta_2(b - c), e_c)$.

- If $z \leq \theta_1 \bar{b} + \theta_2 c$ and $e_a \leq \theta_1 a + \theta_2 c$, the equilibrium outcome is nonexclusive, the nonexclusive bids are the same as those in GSP, and the exclusive bids are $e'_a = e'_b = \theta_1 \bar{b} + \theta_2 c$ and $e'_c = e_c$.

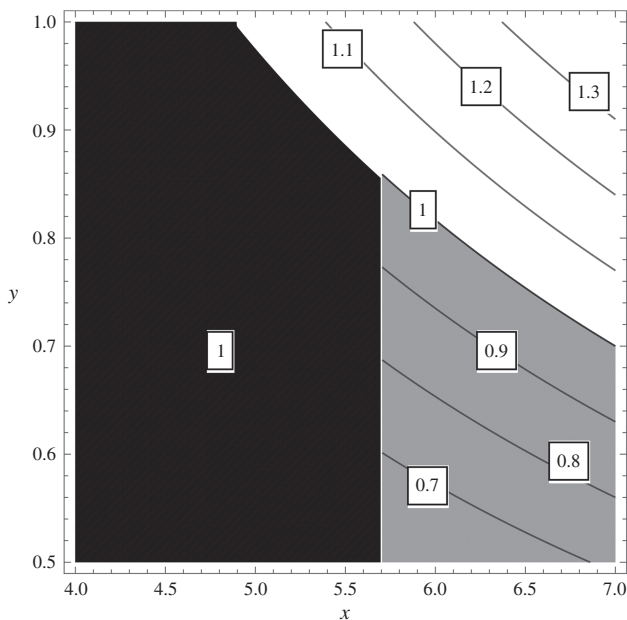
- If $z \in (\theta_1 \bar{b} + \theta_2 c, \theta_1 \bar{b} + \theta_2 c]$ and $e_a \leq \theta_1 a + \theta_2 c$, the outcome is nonexclusive; the bids are $a' = a$, $b' = \bar{b}$, $c' = c$, $e'_a = e'_b = \theta_1 \bar{b} + \theta_2 c$, and $e'_c = e_c$; and (even though the outcome is nonexclusive) the search engine's revenue, $\theta_1 \bar{b} + \theta_2 c$, is greater than that of GSP.

• If $z > \theta_1 \bar{b} + \theta_2 c$ or $e_a > \theta_1 a + \theta_2 c$, the equilibrium outcome is exclusive; the bids are $e'_a = e_a$, $e'_b = e_b$, $e'_c = e_c$, $a' = 0$, $b' = \bar{b}$, and $c' = c$; and in this case, the search engine's revenue, $\max(e_b, e_c)$, could be higher or lower than that of GSP.

Proposition 7 shows the robustness of our result from Proposition 3. In particular, the first case in Proposition 7 corresponds to when $e \leq L_2$ in Proposition 3; in both settings, allowing exclusivity has no effect on the equilibrium outcome of GSP. The second case in Proposition 7 corresponds to when $L_2 < e \leq L_1$ in Proposition 3; in both settings, allowing exclusivity increases the nonexclusive bid of Advertiser B, and even though the outcome of GSP_{2D} is nonexclusive, the revenue in GSP_{2D} is greater than that of GSP. Finally, the third case in Proposition 7 corresponds to when $L_1 < e$ in Proposition 3; in both settings, the outcome is exclusive, and the search engine's revenue in GSP_{2D} can be higher or lower than that in GSP depending on the other parameters in the model.

Figure 4 compares the search engine's revenue of GSP_{2D} to that of GSP. In this figure, we assume that $e_a = x$ and $e_b = e_c = yx$ to reduce the number of parameters in the model. The graph shows that when the value of advertiser A for exclusive outcome is sufficiently low, the equilibrium outcome is nonexclusive, and the revenue of GSP and GSP_{2D} are the same. When the value of advertiser A for exclusivity is sufficiently

Figure 4. The Ratio of Search Engine's Revenue in GSP_{2D} to the Revenue in GSP as a Function of x and y , Where $e_a = x$ and $e_b = e_c = yx$



Notes. In the black region, the revenue of GSP_{2D} is equal to that of GSP; in the gray region, the revenue of GSP is greater than that of GSP_{2D} ; and in the white region, the revenue of GSP_{2D} is greater than that of GSP. The values of the other parameters are $a = 4$, $b = 3.5$, $c = 3$, $\theta_1 = 0.9$, and $\theta_2 = 0.7$.

large, the search engine's revenue in GSP_{2D} is larger than in GSP only if advertisers B and C also have a sufficiently large valuation for exclusivity; otherwise, allowing exclusivity lowers the search engine's revenue.

4.3. VCG and VCG_{2D} Auctions

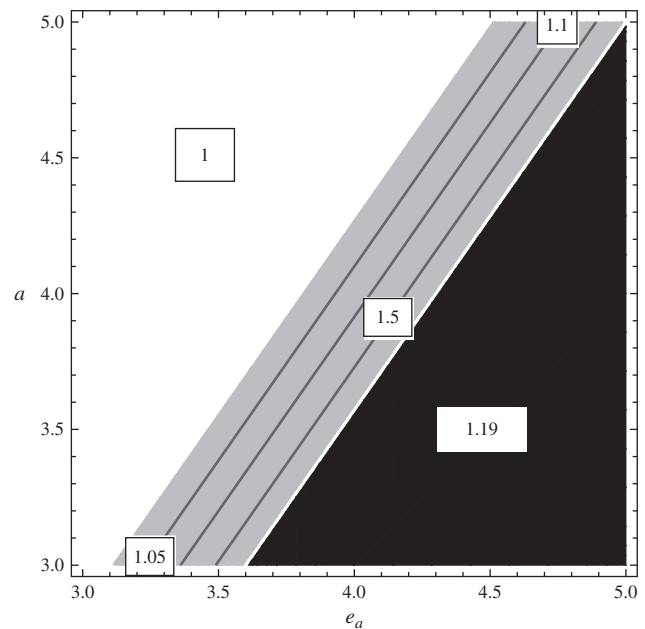
4.3.1. The VCG Auction. In the VCG auction, bids are invited only for the nonexclusive outcome. Advertisers bid truthfully; that is, Advertisers A, B, and C bid a , b , and c , respectively. The revenue from the auction is $(\theta_1 - \theta_2)b + 2\theta_2 c$. More details are provided in the online appendix.

4.3.2. The VCG_{2D} Auction. VCG_{2D} is a special case of the general VCG mechanism, and the advertisers bid truthfully. VCG_{2D} chooses the allocation that maximizes the advertisers' joint valuation, and charges each advertiser the harm that its presence causes to other advertisers. We obtain the following proposition (the analysis is provided in the online appendix).

Proposition 8 (Revenue of VCG_{2D} vs. VCG). *The revenue of VCG_{2D} is always greater than or equal to that of VCG; that is, introducing exclusivity with VCG_{2D} will (weakly) increase revenue.*

This proposition shows that the platform would obtain weakly higher revenue on allowing exclusive display with VCG_{2D} compared to the status quo of VCG with multiple display. This is shown in Figure 5, which

Figure 5. The Ratio of Search Engine's Revenue in VCG_{2D} to the Revenue in VCG as a Function of e_a and a



Notes. In the white region, the revenue of VCG_{2D} is equal to that of VCG; in the gray region, the revenue of VCG_{2D} is greater than that of VCG and increases in e ; and in the black region, the revenue of VCG_{2D} is greater than that of VCG and is constant in e . The values of the other parameters are $b = e_b = 3$, $c = e_c = 2$, $\theta_1 = 0.7$, and $\theta_2 = 0.5$.

plots the ratio of revenue of VCG_{2D} to the revenue of VCG with respect to e_a and a (for specific values of the other parameters). In other words, incorporating exclusivity does not have a downside for the publisher in terms of revenue. As discussed earlier, since the revenue of VCG and GSP are the same under LREF refinement, Proposition 8 also implies that the revenue of VCG_{2D} is greater than or equal to that of GSP. Next, we have the following proposition.

Proposition 9 (Variation of Revenue with Valuations). *In VCG, the auction revenue increases in the advertisers' nonexclusive placement valuations (i.e., a , b , and c). In VCG_{2D} , the auction revenue may decrease in the advertisers' nonexclusive placement valuations (i.e., a , b , and c); however, it (weakly) increases in the advertisers' exclusive placement valuations e_a , e_b , and e_c .*

This proposition shows that for the VCG_{2D} auction, the revenue increases in the exclusive-display valuations of the advertisers. However, a counterintuitive property that it states is that its revenue may decrease in the nonexclusive valuations of the advertisers (which does not happen in the VCG auction).

5. Conclusions and Discussion

We study exclusive placement in online advertising wherein advertisers can bid to be the only one placed on (a part of) a web page. We build a game theory model for the same, and our analysis sheds light on different forces at play when exclusive placement is allowed, and how these forces can impact sellers' revenues compared to the currently popular multiple placement format. We analyze currently used exclusive placement auction formats and provide empirical evidence to support our key insights. Our modeling helps us to understand the key drawbacks of currently used mechanisms, and we suggest realistic auction mechanisms that alleviate these drawbacks.

Our focal analysis studies the GSP_{2D} auction (which is an extension of the GSP auction and is supposedly employed by Bing and Yahoo! in their RAIS program) that enables bidders to submit separate bids for exclusive and nonexclusive placements. Our stylized model shows that an advertiser who prefers the exclusive placement bids high for that outcome, and in response, other advertisers who do not prefer or are not allowed to bid for exclusivity increase their nonexclusive bids (between-advertiser competition effect); however, the advertiser who prefers exclusivity also lowers its nonexclusive bid to induce the search engine to select the exclusive outcome (within-advertiser competition effect). Depending on the advertisers' valuations, and the click-through rates of the advertising slots, the seller's revenue may increase or decrease in the GSP_{2D} auction compared to the GSP auction; counterintuitively, revenue from allowing exclusive display

may decrease when exclusive-display valuations are high rather than low. A data set from Bing's RAIS program provides empirical support for several of our theoretical findings.

We then analyze the VCG_{2D} auction, which applies VCG-type rules to exclusive display, and show that under VCG_{2D} the search engine's revenue weakly increases compared to the status quo of GSP. The reason that VCG_{2D} performs well is that it induces advertisers to bid truthfully and, therefore, it alleviates the issue of the shading of bids due to the within-advertiser competition effect (which reduces revenue), while preserving the between-advertiser competition effect (which increases revenue). In short, using VCG-type rules for allowing exclusive display, appropriately adapted to the situation at hand, improves on the status quo both in search advertising (i.e., multiple display only with GSP) and in display advertising (i.e., multiple display only with VCG).

Our conversations with practitioners indicate that there is a robust demand by advertisers for exclusively placed ads. However, publishers and other market makers are unclear about the best way to sell them—in search advertising, their attempts have led to mixed results for revenue (as the RAIS experience of Bing shows), and in display advertising, no simple and efficient programmatic way exists of selling these ads. Our research can provide guidance to sellers regarding how to implement exclusive placement in online advertising. In fact, a possible solution is to use the same VCG_{2D} auction in both search and display advertising. Some search engines have already started implementing the VCG auction (e.g., Yandex), and the VCG auction is widely used in display advertising—the VCG_{2D} auction is a simple extension of this and guarantees that revenues will be (weakly) higher compared to current practices. Our analysis also provides guidance to advertisers on optimal bidding strategies in the different auctions that we consider (for VCG-based auctions, it is simply truthful bidding).

To our knowledge, our work is one of the first to model exclusive display in sponsored search advertising, and there are numerous avenues for future research. First, we analyze the GSP_{2D} auction that has been patented by Yahoo! and the VCG_{2D} auction. However, other auction mechanisms can also be used for exclusive display. We expect the basic forces that we identify to be at play in other mechanisms as well, and future research can explore this. A related and challenging research question is of deriving the optimal auction mechanism for exclusive display.

Second, to keep the model simple, we make the assumption that advertisers' valuations are independent. Explicitly modeling the effect of one advertiser on another advertiser is an interesting direction for future work. For example, a luxury car manufacturer such as

Lexus may want to be listed exclusively if the competitive advertiser is another luxury car manufacturer such as Acura, but may care less about being listed next to a lower-quality manufacturer such as Kia. In other words, in the spirit of Jerath et al. (2011), the competitive environment of a firm may significantly influence its valuation for exclusive display and therefore its bidding strategy. Desai et al. (2014) study such context effects in a one-dimensional multiple-display auction. Future work can explicitly model these phenomena with the exclusive-display option also available to advertisers.

Finally, allowing each bidder to submit bids for exclusive placement is simply one way to make the currently prevailing auction format more expressive. However, there may be various other formats in which advertisers can reveal their preferences in more detail (e.g., Muthukrishnan 2009). Future research can work toward a general theory of expressive bidding in online advertising auctions.

Acknowledgments

The authors are grateful to Kevin Cabral from AppNexus for detailed discussions regarding practices in the display advertising industry, and to Matt Backus, Arpita Ghosh, Avi Goldfarb, Ernan Haruvy, Preston McAfee, David Pennock, Justin Rao, R. Ravi, Miklos Sarvary, Andrey Simonov, Ken Wilbur, Robert Zeithammer, Yi Zhu, and seminar participants at the 2012 Marketing Science Conference, the 2013 Columbia–NYU–Wharton–Yale Four School Conference, the 2013 UTD Bass Conference, the 2016 Choice Symposium, IDC Herzliya, Koç University, Sabancı University, the University of Alberta, and the University of Texas at Austin for comments on this paper. An earlier draft of this paper was circulated under the name “Exclusive Display in Sponsored Search Advertising.” This work was done when Marjan Baghaie was working at Microsoft. Opinions expressed are those of the authors and do not necessarily represent those of Microsoft.

Appendix

A.1. VCG with Restricted Exclusive Placement

In this section, we calculate advertisers’ payments in the VCG auction with restricted exclusive placement. First note that truthful bidding is a weakly dominant strategy in the VCG auction;¹⁴ we have $a' = a$, $b' = b$, $c' = c$, and $d' = d$. In the presence of advertiser A, the sum of the valuations of the other advertisers is $\theta_2b + \theta_3c$. If Advertiser A did not exist, the sum of the valuations of the other advertisers would be $\theta_1b + \theta_2c + \theta_3d$. Therefore, the payment of advertiser A, when the outcome is nonexclusive, is given by

$$\theta_1b + \theta_2c + \theta_3d - (\theta_2b + \theta_3c) = (\theta_1 - \theta_2)b + (\theta_2 - \theta_3)c + \theta_3d.$$

In the presence of advertiser B, when the outcome is nonexclusive, the sum of the valuations of other advertisers is $\theta_1a + \theta_3c$. If Advertiser B did not exist, the sum of the valuations of the other advertisers would be

$$\theta_1a + \theta_2c + \theta_3d.$$

Therefore, the payment of advertiser B is given by

$$\theta_1a + \theta_2c + \theta_3d - (\theta_1a + \theta_3c) = (\theta_2 - \theta_3)c + \theta_3d.$$

Similarly, we can calculate the payment of advertiser C to be θ_3d .

A.2. VCG_{2D} with Restricted Exclusive Placement

Since truthful bidding is a weakly dominant strategy in the VCG_{2D} auction, we have $a' = a$, $b' = b$, $c' = c$, $d' = d$, and $e' = e$. VCG_{2D} chooses the allocation that maximizes the sum of advertisers’ valuations. Therefore, the outcome is exclusive if and only if the sum of advertisers’ valuations in the exclusive outcome, $e + \theta_3b$, is greater than the sum of advertisers’ valuations in the nonexclusive outcome, $\theta_1a + \theta_2b + \theta_3c$, that is,

$$e + \theta_3b > \theta_1a + \theta_2b + \theta_3c.$$

To calculate the payment of advertiser A, when the outcome is exclusive, we have to calculate how much existence of advertiser A hurts the other advertisers. If advertiser A did not exist, the sum of (other) advertisers’ valuations would be $\theta_1b + \theta_2c + \theta_3d$. In the presence of advertiser A, however, the sum of other advertisers’ valuations is θ_3b . The difference determines the payment of advertiser A, that is, the payment of advertiser A when the outcome is exclusive is

$$(\theta_1b + \theta_2c + \theta_3d) - \theta_3b = (\theta_1 - \theta_3)b + \theta_2c + \theta_3d.$$

Similarly, to calculate the payment of advertiser B, when the outcome is exclusive, we calculate how much its presence hurts the other advertisers. If advertiser B did not exist, the sum of other advertisers’ valuations would be $e + \theta_3c$. In the presence of advertiser B, however, the sum of other advertisers’ valuations is e . Therefore, the payment of advertiser B, when the outcome is exclusive, is given by

$$e + \theta_3c - e = \theta_3c.$$

The search engine’s total revenue is the sum of the payments of advertisers A and B, that is,

$$R_{VCG_{2D}}^E = (\theta_1 - \theta_3)b + (\theta_2 + \theta_3)c + \theta_3d.$$

It is easy to see that $R_{VCG_{2D}}^E \geq R_{GSP}$.

Next, we consider the nonexclusive outcome and use the same method as before to calculate advertisers’ payments. In the presence of advertiser A, the sum of the valuations of the other advertisers is $\theta_2b + \theta_3c$. If advertiser A did not exist, the sum of the valuations of the other advertisers would be $\theta_1b + \theta_2c + \theta_3d$. Therefore, the payment of advertiser A, when the outcome is nonexclusive, is given by

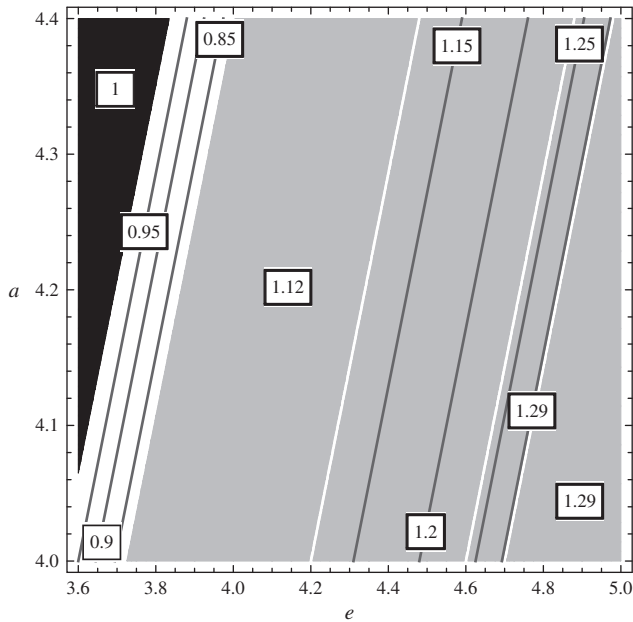
$$\theta_1b + \theta_2c + \theta_3d - (\theta_2b + \theta_3c) = (\theta_1 - \theta_2)b + (\theta_2 - \theta_3)c + \theta_3d.$$

In the presence of advertiser B, when the outcome is nonexclusive, the sum of the valuations of other advertisers is $\theta_1a + \theta_3c$. If advertiser B did not exist, the sum of the valuations of the other advertisers would be

$$\max(\theta_1a + \theta_2c + \theta_3d, e + \theta_3c),$$

where the first term is when the outcome in the absence of advertiser B is still nonexclusive, and the second term is when the outcome in the absence of advertiser B becomes exclusive.

Figure A.1. Comparison of GSP_{2D} and VCG_{2D} Revenues with Restricted Exclusive Placement



Notes. The figure shows the ratio of the search engine's revenue in VCG_{2D} to the revenue in GSP_{2D} as a function of e and a . In the black region, the revenue of VCG_{2D} is equal to that of GSP_{2D} . In the white region, the revenue of VCG_{2D} is lower than that of GSP_{2D} and the ratio decreases in e . In the gray region, the revenue of VCG_{2D} is greater than that of GSP_{2D} and the ratio increases in e . The values of the other parameters are $b = 4$, $c = 3$, $d = 2$, $\theta_1 = 0.7$, $\theta_2 = 0.5$, and $\theta_3 = 0.1$. (These are the same values as for Figure 2(d) and the other similar figures in this paper. The choice of axis ranges is also the same.) Note that if we consider the region $e \geq a$ in the figure (which is a reasonable condition to assume), then the revenue of VCG_{2D} is always larger than the revenue of GSP_{2D} .

Using the VCG payment rule, the payment of advertiser B is given by

$$\begin{aligned} & \max(\theta_1 a + \theta_2 c + \theta_3 d, e + \theta_3 c) - (\theta_1 a + \theta_3 c) \\ & = \max((\theta_2 - \theta_3)c + \theta_3 d, e - \theta_1 a). \end{aligned}$$

Similarly, we can calculate the payment of advertiser C to be $\max(\theta_3 d, e - \theta_1 a - (\theta_2 - \theta_3)b)$. Together, we can see that when the outcome is nonexclusive, advertiser A's payment and allocation in VCG are the same as in GSP. Advertisers B and C also get the same allocation as in GSP; however, their payments might be higher than in GSP. In fact, the first term in the $\max(\cdot)$ expression for advertisers B and C is what they would pay in GSP. Overall, the revenue of the search engine will be weakly higher in VCG_{2D} compared with GSP.

Endnotes

¹These are often classified as "share of voice" ads, such as page takeovers (wherein an ad is displayed on the whole web page for a few seconds or the user has to click to advance to the content) and roadblock ads (wherein an advertiser runs its ads in multiple placements on the same web page; see Figures 1(b) and 1(c)). We especially thank Kevin Cabral from AppNexus for detailed discussions regarding practices in the display advertising industry.

²Note that we are not aiming to develop optimal auction mechanisms; rather, we are analyzing auction mechanisms used in the industry, or their variants. Having said that, we note that the VCG-based auctions are welfare maximizing.

³For legal reasons, we cannot reveal the details of the exact mechanisms being used by Yahoo! and Bing.

⁴This assumption is not critical to our insights and is made to facilitate exposition as it reduces the number of cases that need to be considered. This is a reasonable assumption that has been pointed out in empirical papers on branded keywords (e.g., Simonov et al. 2015) and is supported by the bids in the data that we obtain from Bing.

⁵We have also developed and analyzed a model in which others' valuations are known to advertisers only up to a distribution. This model is cumbersome to solve for both the GSP and GSP_{2D} auctions; however, we obtain all of the insights that are obtained through the current model. Details of this analysis are available on request.

⁶They could possibly bid above their valuation if we do not choose the LREF equilibrium.

⁷Advertiser FMA's nonexclusive bid could actually be any amount less than or equal to d . However, since it would not affect any of the payments or the allocation, the outcome would be the same as when it is 0.

⁸We assume that the parameter values satisfy the condition $\theta_1 c + \theta_2 d < L_1$.

⁹In general, depending on the values of b , c , d , and the θ_i 's, the search engine's revenue in GSP_{2D} in this region might be higher or lower than in GSP; the values in Figure 2 have been chosen to illustrate that it can be lower.

¹⁰Such a change is usually due to a change in the advertiser's ad agency, and could imply a change in the advertiser's overall strategy.

¹¹The fact that the LREF equilibrium of GSP has the same payment and allocation as in VCG is not new, and is discussed in Edelman et al. (2007) and Varian (2007). We present the results here for the sake of completeness.

¹²Note that the revenue of GSP_{2D} may exceed the revenue of VCG_{2D} under some conditions. However, comparing Figures 2(d) and 3 shows that this is the case only for a narrow range of parameter values, and this actually does not happen for large values of e (this can readily be shown analytically), which is exactly where using VCG_{2D} has a large advantage. We illustrate this with the help of Figure A.1 in the appendix.

¹³Note that nonexclusivity does not mean that the advertiser is guaranteed to be shown with other advertisers; therefore, it is reasonable to assume that an advertiser's valuation with no guarantees is less than or equal to when it is guaranteed to be shown exclusively.

¹⁴For a proof of truthfulness of VCG, see Krishna (2010).

References

- Amaldoss W, Desai P, Shin W (2015) Keyword search advertising and first-page bid estimates: A strategic analysis. *Management Sci.* 61(3):507–519.
- Amaldoss W, Jerath K, Sayedi A (2016) Keyword management costs and "broad match" in sponsored search advertising. *Marketing Sci.* 35(2):259–274.
- Athey S, Ellison G (2011) Position auctions with consumer search. *Quart. J. Econom.* 126(3):1213–1270.
- Branco F (1997) The design of multidimensional auctions. *RAND J. Econom.* 28(1):63–81.
- Cary M, Das A, Edelman B, Giotis I, Heimerl K, Karlin A, Mathieu C, Schwarz M (2008) On best-response bidding in GSP auctions. Working paper, Harvard University, Boston.

- Che YK (1993) Design competition through multidimensional auctions. *RAND J. Econom.* 24(4):668–680.
- Clarke EH (1971) Multipart pricing of public goods. *Public Choice* 11(1):17–33.
- Crampton P, Shoham Y, Steinberg R (2006) *Combinatorial Auctions* (MIT Press, Cambridge, MA).
- Decarolis F, Goldmanis M, Penta A (2017) Marketing agencies and collusive bidding in online ad auctions. Working Paper, National Bureau of Economic Research, Cambridge, MA.
- Desai P, Shin W, Staelin R (2014) The company that you keep: When to buy a competitor's keyword. *Marketing Sci.* 33(4):485–508.
- Dukes A, Gal-Or E (2003) Negotiations and exclusivity contracts for advertising. *Marketing Sci.* 22(2):222–245.
- Edelman B, Ostrovsky M (2007) Strategic bidder behavior in sponsored search auctions. *Decision Support Systems* 43(1):192–198.
- Edelman B, Ostrovsky M, Schwarz M (2007) Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Amer. Econom. Rev.* 97(1):242–259.
- eMarketer (2016) Digital ad spending to surpass TV next year. (March 8), <http://bit.ly/1TH8wBZ>.
- Ghosh A, Sayedi A (2010) Expressive auctions for externalities in online advertising. Rappa M, Jones P, Freire J, Chakrabarti S, eds. *Proc. Nineteenth Internat. World Wide Web Conf. (ACM, New York)*, 371–380.
- Ghosh A, McAfee P, Sayedi A (2011a) Expressive bidding in online advertising auctions. U.S. Patent 20110071908, filed September 23, 2009, issued March 24, 2011.
- Ghosh A, McAfee P, Sayedi A (2011b) Expressive bidding online advertising auction mechanisms. U.S. Patent 20110071909, filed September 23, 2009, issued March 24, 2011.
- Gomes R, Sweeney K (2014) Bayes–Nash equilibria of the generalized second-price auction. *Games Econom. Behav.* 86:421–437.
- Groves T (1973) Incentives in teams. *Econometrica* 41(4):617–631.
- Jerath K, Ma L, Park Y-H, Srinivasan K (2011) A “position paradox” in sponsored search auctions. *Marketing Sci.* 30(4):612–627.
- Jeziorski P, Segal I (2015) What makes them click: Empirical analysis of consumer demand for search advertising. *Amer. Econom. J. Microeconomics* 7(3):24–53.
- Katona Z, Sarvary M (2010) The race for sponsored links: Bidding patterns for search advertising. *Marketing Sci.* 29(2):199–215.
- Krishna V (2010) *Auction Theory*, 2nd ed. (Academic Press, Burlington, MA).
- Lu S, Zhu Y, Dukes A (2015) Position auctions with budget-constraints: Implications for advertisers and publishers. *Marketing Sci.* 34(6):897–905.
- Manzi J (2012) *Uncontrolled: The Surprising Payoff of Trial-and-Error for Business, Politics, and Society* (Basic Books, New York).
- Metz C (2008) Google remodels top secret money machine. *The Register* (September 2), http://www.theregister.co.uk/2008/09/02/changes_for_adwords.
- Metz C (2011) Google exec announces departure as Page takes reins. *The Register* (April 5), http://www.theregister.co.uk/2011/04/05/senior_jonathan_rosenberg_to_leave_google.
- Mori PA (2006) Design of reasonable multidimensional auctions under regulator ignorance. *Proc. Internat. Public Procurement Conf.*, 781–811.
- Muthukrishnan S (2009) Bidding on configurations in Internet ad auctions. Ngo HQ, eds. *Comput. Combinatorics, Lecture Notes Comput. Sci.*, Vol. 5609 (Springer, Berlin Heidelberg), 1–6.
- Reiley DH, Li SM, Lewis RA (2010) Northern Exposure: A field experiment measuring externalities between search advertisements. Parkes DC, Deliarocas C, Tennenholtz M, eds. *Proc. 11th ACM Conf. Electronic Commerce (EC-2010)* (Cambridge, MA), 297–304.
- Russian Search Marketing (2015) Yandex.Direct auction change. (July 30), <http://russiansearchmarketing.com/yandex-direct-auction-change>.
- Sayedi A, Jerath K, Srinivasan K (2014) Competitive poaching in sponsored search advertising and strategic impact on traditional advertising. *Marketing Sci.* 33(4):586–608.
- Shin W (2015) Keyword search advertising and limited budgets. *Marketing Sci.* 34(6):882–896.
- Simonov AD, Li E, Rao JM (2015) Search advertising auction with dominant advertiser. Working paper, Microsoft Research.
- Thiel SE (1988) Multidimensional auctions. *Econom. Lett.* 28(1):37–40.
- Varian HR (2007) Position auctions. *Internat. J. Indust. Organ.* 25(6):1163–1178.
- Varian HR, Harris C (2014) The VCG auction in theory and practice. *Amer. Econom. Rev.* 104(5):442–445.
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1):8–37.
- Wilbur KC, Zhu Y (2008) Click fraud. *Marketing Sci.* 28(2):293–308.
- Zhu Y, Wilbur KC (2011) Hybrid advertising auctions. *Marketing Sci.* 30(2):249–273.