

Exclusive Display in Sponsored Search Advertising

Kinshuk Jerath

Assistant Professor of Marketing

Tepper School of Business

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

(412) 268-2215

kinshuk@cmu.edu

Amin Sayedi[†]

Doctoral Student

Tepper School of Business

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

(412) 268-2295

ssayedir@cmu.edu

May 2011

[†]The authors are grateful to Arpita Ghosh, Avi Goldfarb, Ernan Haruvy, Preston McAfee, David Pennock, R. Ravi, Ken Wilbur, Robert Zeithammer and Yi Zhu for useful discussions and suggestions.

Exclusive Display in Sponsored Search Advertising

Abstract

As sponsored search becomes increasingly important as an advertising medium for firms, search engines are exploring more advanced bidding and ranking mechanisms to increase their revenue from sponsored search auctions. For instance, Google, Yahoo! and Bing are investigating auction mechanisms in which each advertiser submits two bids: one bid for the standard display format in which multiple advertisers are displayed, and one bid for being shown exclusively. If the exclusive-placement bid by an advertiser is high enough then only that advertiser is displayed, otherwise multiple advertisers are displayed and ranked based on their multiple-placement bids.

We call such auctions two-dimensional auctions and study two extensions of the Generalized Second Price (*GSP*) mechanism that are currently being evaluated by search engines. We show that allowing advertisers to bid for exclusivity always generates higher revenues for the search engine. Interestingly, even if the final outcome is multiple display, the search engine still extracts higher revenue because of increased competition among advertisers, simply because bidding for exclusivity is allowed. In fact, one of the auctions we consider can extract the full surplus of the bidders as search engine revenue under certain conditions. Under other conditions on the advertisers' valuations for exclusivity as well as the heterogeneity in their valuations for exclusivity, exclusive display auctions can benefit the advertisers also.

Keywords: sponsored search advertising, exclusive display, game theory, position auctions, two-dimensional auctions.

1 Introduction

Online advertising is fast becoming an increasingly important component of any firm's advertising mix. In turn, one of the primary forms of online advertising is sponsored search advertising on popular search engines such as Google, Yahoo! and Bing. In sponsored search advertising, advertisers pay a fee to the search engine to have links to their websites listed as relevant results in response to a keyword search. When a user submits a query on the search engine, she is presented with advertisements (henceforth, ads) that are placed into positions, usually arranged linearly down the side of the page (along with the organic search results which are not sponsored).¹ Sponsored search advertising is the primary source of revenue for search engines; for instance, Google, Yahoo! and Bing earn millions of dollars per day through this channel.²

Being their largest source of revenue, the pricing mechanism for sponsored search advertising is of critical importance to search engines. All the prominent search engines currently run Generalized Second Price (*GSP*) auctions to sell their advertising space. However, this choice of the auction mechanism was not a straightforward one, and the industry went through several phases before the *GSP* auction became the dominant choice. Sponsored search was introduced in 1997 when the search engine company GoTo (renamed Overture in 2001, and acquired by Yahoo in 2003) let advertisers bid to appear among the top search results. GoTo's original sponsored search mechanism was a Generalized First Price auction in which every advertiser submitted a bid and the advertisers were arranged in descending order of bids, with each one paying his bid. This auction was also adopted by both Yahoo! and MSN (now rebranded as Bing). The payment mechanism was also experimented with, and while initially advertisers had to pay every time their ad was shown (pay per impression), this was changed to payment every time their ad was actually clicked (pay per click). The Generalized First Price auction, however, was soon found to be an unstable auction mechanism in which advertisers had the incentive to cyclically bid low and high amounts to game the system (Edelman and Ostrovsky 2007). This motivated the need for a more stable mechanism.

In 2002, Google introduced the Generalized Second Price (*GSP*) auction with the basic rules

¹Throughout the paper, we refer to a user as "she," an advertiser as "he," and the search engine as "it."

²Based on publicly released information available at <http://investor.google.com/earnings.html> for Google, <http://yhoo.client.shareholder.com/results.cfm> for Yahoo!, and <http://www.microsoft.com/investor/EarningsAndFinancials/Earnings/PressReleaseAndWebcast/FY11/Q1/default.aspx> for Bing.

that every advertiser submits his per-click bid (i.e., how much he is willing to pay for every click obtained from a consumer) but has to actually pay only the minimum amount necessary to keep his current position in the list of results (i.e., *GSP* is a “next-price” auction). *GSP*, a much more stable auction mechanism, was gradually adopted by other prominent search engines as well (e.g., Yahoo! and Bing currently also use this mechanism). Search engines also continually conduct their own internal experimentation, based on which they apply slight variations (the exact details of which are often not publicly announced) to the basic *GSP* mechanism. For instance, advertisers are now ranked based not on their submitted bids, but based on their effective bids, where each advertiser’s effective bid is obtained by multiplying his submitted bid with a quality score specific to the advertiser-keyword combination. As another example, search engines have experimented with a pay-per-action payment mechanism in which an advertiser has to pay only if, after clicking on an ad, a user also completes a predetermined action, such as purchasing the product or spending two minutes on the advertiser’s website.

The above discussion shows that the spectacular rise of sponsored search advertising in the last decade has been accompanied by constant effort from the search engines to refine their pricing mechanisms by gradually fixing the deficiencies in them, including developing new auction mechanisms to rank advertisers. Moreover, if one search engine introduces a profitable innovation, other search engines follow suit in a short time. As the industry matures, search engines are looking to further expand their bidding mechanisms by allowing advertisers to be more specific about their utilities and to express a richer set of preferences. For example, Google has been considering “hybrid” advertising auctions which allow advertisers to bid on a per-impression or a per-click basis for the same advertising space. Zhu and Wilbur (2011) show that such auctions can enhance both search engine revenue and the efficiency of advertisers’ allocation to positions.

A recent and very interesting development in this context has been the exploration by search engines of auction mechanisms that allow advertisers to bid for *exclusive display* in response to a user search. In other words, advertisers can bid for their ad to be the only ad displayed, rather than being one of many ads displayed. Exclusive display may be an attractive option for advertisers as an advertiser can create strong brand associations by being the only one displayed in response to certain keywords. For example, if the ad of only the manufacturer Olympus gets displayed in response to the keyword “digital camera,” it can be a significant branding advantage for Olympus

over its competitors such as Canon and Nikon. Moreover, multiple ads shown next to each other may impose negative externalities on each other. For example, if a user who has searched for the keyword “car rental” clicks on the ad of Hertz, chances are that she will also go back to check the ads of some other companies displayed in the sponsored list, such as Avis and Budget, before finalizing the transaction with Hertz. These negative externalities, which can decrease the values of clicks to advertisers, will be smaller if only one ad is shown to the user.

Exclusive display is even more valuable when a particular brand name is the keyword searched, because the brand owner would want to be the only advertiser displayed in response. For example, in high-profile cases, Rosetta Stone and Luis Vuitton sued Google in USA and Europe, respectively (Mullin 2010, Sterling 2010), in an attempt to have laws enacted to prevent bidding on trademarks by competitors. While these companies lost these legal battles, it does reveal the strong incentive of brand owners to be displayed exclusively, which the search engines could profit from. Desai et al. (2010) similarly argue, with experimental support, that when advertisers are listed next to each other, “context effects” influence users’ perceptions of their relative qualities, which in certain cases can hurt the advertisers, especially the high-quality ones. Such effects may motivate advertisers to prefer exclusive display. In summary, exclusive display will not only increase the expected clicks on an ad if the relevant keyword is searched, but may also increase the valuation per click for the advertiser. Therefore, advertisers may be willing to pay a higher price per click for having exclusive appearance for some keywords.

Interestingly, all three of the most popular search engines in the USA, namely Google, Yahoo! and Bing, have explored exclusive display auctions as part of their research efforts in the recent past. Google has considered displaying exclusive ads in response to user queries as part of its “perfect ad” initiative (Metz 2011). Yahoo! has, in fact, advanced even further and, in March 2011, has been granted patents on certain aspects of the idea of exclusive display and on two particular exclusive display auctions that it developed (U.S. Patent 20110071908 and U.S. Patent 20110071909). Our discussions with researchers and executives at Bing indicate that Microsoft is also exploring exclusive display auctions.³ These are strong indicators that exclusive display of ads

³Notably, in the late 1990s and early 2000s, AOL, which had a large share of the online search market at the time, was using exclusive listings. For example, it signed contracts with eBay and Monster which allowed them to be the exclusive providers of certain services when accessed through AOL (Bradley 2001, Hallowell 2002, Rayport 2000). This, however, was different from the sponsored ads arrangement, and the pricing mechanism was not auction based.

may be adopted by popular search engines in the near future. There is, however, also a debate about the value to search engines of switching to exclusive display mechanisms. For instance, Metz (2011) reports that high-level executives at Google were constantly going back and forth between adopting and not adopting exclusive display advertising. Eventually, due to a lack of proper understanding, Google chose to temporarily stay with the status quo of displaying multiple ads. Faced by the same questions, Yahoo! and Bing are conducting independent research on the value of exclusive display auctions.

Exclusive display of sponsored ads raises several questions. Is it advantageous for the search engine to switch to exclusive display, and why (or why not)? Should firms switch completely to exclusive display or should they adopt a “hybrid” format in which firms can bid for both multiple display and exclusive display, and the final outcome is decided by the search engine based on the submitted bids? What are the forces driving bidding behavior in exclusive display auctions? If firms can extend the current *GSP* mechanism in different ways to allow firms to bid for exclusive display, which of these auction mechanisms will provide the most revenue and under what conditions? Finally, what is the impact on the advertisers and on social welfare? In this paper, we take a step forward towards answering these questions.

Given that *GSP* is the auction mechanism used widely, we consider auction mechanisms that are extensions of *GSP*. We start by assuming that each advertiser can have different per-click valuations for clicks obtained when it is displayed with other advertisers (multiple display) and clicks obtained when it is the only one displayed (exclusive display). We analyze two auction mechanisms that are conceptually simple extensions of *GSP* and were recently patented by Yahoo! as the main candidates for implementation. In these auctions, each advertiser submits a two-dimensional bid—its maximum willingness to pay per click for multiple display, and its maximum willingness to pay per click for exclusive display. In the first auction, NP_{2D} , the next-price rule of *GSP* is extended to two dimensions—every advertiser has to pay the minimum amount necessary to maintain the outcome configuration (multiple or exclusive display) and its position within the configuration. In the second auction, GSP_{2D} , the allocation and pricing rules are defined to be exactly those of *GSP* when multiple ads are displayed. The idea behind defining the GSP_{2D} auction in this manner is that when multiple ads are displayed, advertisers see no difference at all between the new auction and the existing *GSP* system.

First, we develop and analyze a simple game-theoretic model which provides several insights into exclusive-display auctions. We show that allowing the advertisers to bid for exclusivity always generates higher revenue for the search engine. However, our results also make it clear that search engines should not adopt auctions with exclusive placement only; rather, they should adopt hybrid auctions that allow advertisers to bid for multiple as well as exclusive placement. If an advertiser highly values being the only one displayed, then exclusive display can be the outcome and can provide higher revenue for the search engine. More interestingly, we find situations in which search engine revenue increases even when the outcome is multiple display, simply because advertisers could bid for exclusive display also. This is because advertisers compete not only for ranks within the multiple display outcome, but also compete for the outcome itself to be multiple display. In fact, this competition is highest at the point at which the search engine is indifferent between multiple display and exclusive display. At this point, the search engine can extract *all* of the bidders' surplus as its revenue. Under other situations, both search engine revenue and bidders' surplus can simultaneously increase.

We also derive results regarding which auction, NP_{2D} or GSP_{2D} , gives higher search engine revenue and allocative efficiency under different conditions. In general, we find that NP_{2D} is the better auction in terms of revenue as well as allocative efficiency. Moreover, NP_{2D} is revenue monotone in bidders' valuations for clicks but GSP_{2D} is not. These advantages of GSP_{2D} lend support to the idea that the simple "next price" heuristic of GSP is a good heuristic to use while designing extended position auctions.

Next, we run a comprehensive simulation study in which we analyze more realistic, and more complicated, situations. The simulations confirm the results from our simpler analytical model and also provide new insights. For instance, we find that a larger number of bidders significantly increases the revenue advantage to the search engine of using two dimensional auctions. We also find that as heterogeneity among advertisers in their valuations for exclusive display increases, both search engine revenue and bidders' surplus can be significantly higher in two-dimensional auctions as compared to the one-dimensional GSP auction.

Given the keen interest from the industry in exclusive display in sponsored search advertising, but also the lack of good understanding of the same, the insights from our study are very timely and relevant. For instance, Google's dilemma seems to have been that allowing exclusive display might

reduce revenue because of the loss of revenue from the many advertisers who will not be displayed (Metz 2011). We find that allowing advertisers to bid for both multiple and exclusive display can significantly increase search engine revenue. Moreover, we conjecture that multiple display will still be the outcome for most keywords, but search engine revenue will be higher than *GSP*. This is because we can expect that exclusive display will be valued more than multiple display by almost all advertisers for almost all keywords but, for most keywords, may not be valued so much more that one advertiser dominates all others combined.

The rest of this paper is structured as follows. In the next section, we briefly review the literature related to our work. In Section 3, we develop the general framework for our analysis. In Section 4, we define and analyze a simple model within our general framework and obtain insights into the dynamics of exclusive-display auctions. In Section 5, we conduct simulations which numerically confirm the results from our analytical exercise in a more complicated setting and also provide additional insights. In Section 6, we conclude by summarizing our results and laying out directions for future work.

2 Related Literature

Theoretical studies in the Economics and Marketing communities have significantly enhanced our understanding of position auctions used in sponsored search advertising. Edelman and Ostrovsky (2007) studied first-price auctions and established that bidding will be cyclical and unstable in these auctions. Edelman, Ostrovsky and Schwarz (2007) and Varian (2007) showed that bidding is stable in the Generalized Second Price auction (*GSP*), but bids do not truthfully reveal valuations of advertisers for positions. Various other papers that consider different aspects of second-price position auctions include Athey and Ellison (2011), Edelman and Schwarz (2010), Jerath et al. (2011), Katona and Sarvary (2010), Liu et al. (2010) and Desai et al. (2010). Many of the above papers consider both pay-per-impression and pay-per-click payment schemes. Zhu and Wilbur (2011) consider hybrid auctions in which advertisers can choose to bid on a per-impression or a per-click basis, while Agarwal et al. (2009) analyze bidding in a pay-per-action second-price auction. All of the above papers, however, study auctions that only consider displaying multiple advertisers in response to a keyword search.

To our knowledge, only two other papers (in the Computer Science community) analyze position auctions in which advertisers can express their preferences beyond simply turning in bids for a multiple-display 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 he wants to be displayed with. Note that this is a very different auction mechanism from the ones we consider in this paper. Furthermore, the focus of Muthukrishnan (2009) is on developing a fast algorithm to determine the outcome of this auction (which includes deciding how many ads to display, and which advertisers to include and how to rank them), while the revenue and efficiency properties of the auction itself are not analyzed. The paper closest to our work is Ghosh and Sayedi (2010), who analyze the same auctions as we do. However, they derive a very different set of results as their focus is on comparing the properties of the multiple equilibria that the NP_{2D} and GSP_{2D} auctions can attain. In contrast, in this paper, our aim is to understand at an intuitive level how exclusive-display auctions work, and which auction is more beneficial to the search engine and to the advertisers under different conditions. We believe that our results and insights, while being of academic interest, also speak closely to the needs of a managerial audience.

Finally, there is previous work in Economics and Marketing that our paper is related to. There is a small literature on multi-dimensional auctions in which bidders, differentiated on multiple characteristics, submit multi-dimensional bids and a winner is determined (Branco 1997, Che 1993, Mori 2006, Thiel 1988). For example, in an auction for a contract to build an aircraft, bidders quote a price and also specify the components of the aircraft (Branco 1997). Additionally, note that exclusivity contracts are often negotiated between media providers and advertisers for traditional media advertising. For example, Anheuser-Busch and Volkswagen held the rights for advertising exclusively in the beer and automotive categories, respectively, during Super Bowl 2011.⁴ Dukes and Gal-Or (2003) study this market. However, our work is very different from these literature streams. First, the institutional details of our setting introduce several differences (e.g., ranked outcomes, per-click bidding by advertisers, bid-weighting by the auctioneer). 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

⁴<http://beernews.org/2011/01/beer-playing-a-major-role-in-super-bowl-xlv-chatter>;
http://media.vw.com/press_releases/volkswagen-of-america-will-be-the-super-bowls-exclusive-car-advertiser-a-first-for-the-german-automaker

or only one winner.

3 Framework for Analysis

In this section, we describe the general framework that we use in the paper. When a user of the search engine submits a query, she is shown two lists of results—the organic list and the sponsored list. The sponsored list is a ladder of, usually text-only, ads towards the right of the results page. Sometimes, one to four ads are also placed above the organic search results. A position that contains an ad is called a slot, and the search engine basically assigns the ads to the slots. The slots that are placed above the organic search results are more likely to get clicks than those placed on the right and are considered more valuable; similarly, the slots placed at upper positions are more valuable than those placed in lower positions. Therefore, we get a total ordering, and we can model the ad presentation as an array of slots where the earlier positions in the array are more valuable and more likely to get clicks than the later positions.

We assume that there are n advertisers who want to display their ads. In our context, ads can be displayed in one of two formats. In the first format, multiple ads are displayed; specifically, k ads can be displayed, where $k < n$. Slot i is associated with a number $0 < \theta_i \leq 1$ called the click-through rate (CTR) of the slot. The number θ_i indicates the probability of being clicked if multiple ads are displayed and an ad is placed at slot i . According to the above discussion, θ_i s are sorted in descending order along the array of slots.

In the second format, only one ad is displayed exclusively. In this case, we assume that the click-through rate of the only slot shown is $\hat{\theta}$. We assume that $\hat{\theta} \geq \theta_1$, i.e., the only slot shown in the exclusive-display outcome gets at least as many clicks as the first slot in the multiple-display outcome. Since normalizing does not affect our results, we assume $\hat{\theta} = 1$ for simplicity. Note that the total number of clicks on all the sponsored links combined can be higher when multiple links are displayed, i.e., $\sum_{i=1}^k \theta_i$ can be higher than $\hat{\theta}$. In fact, we assume this to be true in our analytical and simulation studies.

Each advertiser i of the n advertisers has a vector of valuations (v_i^N, v_i^E) , where v_i^N is the valuation of a click when displayed with multiple other advertisers and v_i^E is the valuation of a click when displayed alone in response to a keyword search (the superscripts N and E stand for

non-exclusive and exclusive, respectively). As discussed in the introduction, quality perceptions and post-click conversion rates can improve under exclusive display. For these reasons, we make the reasonable assumption that $v_i^E \geq v_i^N$. The search engine runs an auction in which it invites bids from advertisers. We now describe the different auction mechanisms.

3.1 One-Dimensional Auction: *GSP*

Major search engines, such as Google, Yahoo! and Bing, use a Generalized Second Price (*GSP*) auction for allocation and pricing. Consider a keyword for which each advertiser i submits a bid b_i . The search engine sorts the advertisers in descending order of their bids and allocates the first (highest, and most valuable) slot to the highest bidder, the second slot to the second-highest bidder, and so on, until either all slots are allocated, or the bid is lower than the reserve price r . A reserve price of r means that the search engine would rather leave the slot empty than to sell it at a price less than r .⁵ Suppose that the bids submitted by the advertisers are b_1, \dots, b_n , and without loss of generality assume that $b_1 \geq b_2 \geq \dots \geq b_n$. Therefore, as we described, the bidder i with bid b_i gets the i -th slot, as long as $b_i \geq r$ and $i \leq k$, where r is the reserve price and k is the number of slots available.⁶

The bidder who is assigned to the i -th slot has to pay $\max(b_{i+1}, r)$ every time a user clicks on his ad; in other words, the payment rule is pay per click, and every bidder has to pay the minimum amount necessary to keep his position. For example, if bidder i whose bid is b_i changes his bid all the way down to $\max(b_{i+1}, r)$, but not to any anything below that, he would still get the same slot; therefore, the price of bidder i is set to $\max(b_{i+1}, r)$. This characteristic makes the *GSP* a “next-price” auction. The revenue of *GSP* is given by $\sum_{i=1}^k \theta_i \max(b_{i+1}, r)$, and the social welfare is given by $\sum_{i=1}^k \theta_i v_i^N$.⁷ Note that since multiple ads are always displayed, the valuations

⁵We assume that the reserve price is exogenously specified for all auctions, which is the assumption in most papers on position auctions (Edelman, Ostrovsky and Schwarz 2007, Jerath et al. 2011, Katona and Sarvary 2010, Varian 2007). Note that if the number of bidders is greater than the number of slots, the reserve price does not matter (as long as it is not so high that it can preclude some bidders from taking part in the auction).

⁶Search engines often transform a bid b_i to an *effective bid* $\hat{b}_i = \gamma_i \times b_i$, where γ_i is a quality score which depends upon the past performance of the ad (how likely it is to generate clicks), the relevance of the ad to the keyword, the reputation of the advertiser, etc. The search engine only works with the effective bids rather than the original bids. The practice of transforming original bids into effective bids does not play an important role in the analysis presented in our paper and does not change the key insights. Hence, for ease of understanding, we present our results without considering this transformation.

⁷Note that, throughout the paper, social welfare is defined as the sum of the revenue of the search engine and the bidders’ surplus. It does not directly include the impact on the consumers (e.g., consumers have less choice if only one ad is displayed), since we do not model these aspects explicitly.

of advertisers for exclusive display have no relevance in the *GSP* auction.

3.2 Two-Dimensional Auctions: NP_{2D} and GSP_{2D}

The goal of this paper is to understand auctions which allow advertisers to bid for being shown exclusively on the results page. While bidding for exclusivity is not yet implemented at any major search engine, many of them are internally experimenting with such auctions. We analyze the two auction mechanisms, NP_{2D} and GSP_{2D} , recently proposed and patented by Yahoo!. Since the current mechanism used by major search engines, including Yahoo!, is *GSP*, both proposed mechanisms are also extensions of *GSP*.

NP_{2D} and GSP_{2D} are “two-dimensional” auctions, i.e., each advertiser simultaneously submits two bids b^N and b^E where b^N indicates how much they are willing to pay per click if their ad is shown among other ads, and b^E indicates how much they are willing to pay per click if their ad is shown exclusively (as before, the superscripts N and E stand for non-exclusive and exclusive, respectively). Similarly, the outcome of the auction can be either E or N , where E means that only one ad is displayed, and N means that multiple ads are shown to the user. We call the non-exclusive bid of an advertiser, b^N , his N -bid, and the exclusive bid, b^E , his E -bid. If the outcome is N , we assume that k ads are shown and their click-through rates of the k ordered slots are $\theta_1, \dots, \theta_k$. If the outcome is E , the click-through rate of the only slot shown is $\hat{\theta} = 1 \geq \theta_1$.

In the following subsections, we describe the mechanisms of NP_{2D} and GSP_{2D} . We assume that the n bidders submit bids $(b_1^N, b_1^E), \dots, (b_n^N, b_n^E)$, and without loss of generality we assume that $b_1^N \geq b_2^N \geq \dots \geq b_n^N$. Since the ordering of N -bids might be different from that of E -bids, assume that e_1 and e_2 are the indices of the highest and the second highest E -bids, respectively: $b_{e_1}^E \geq b_{e_2}^E \geq b_i^E$ for any $i \notin \{e_1, e_2\}$.

The NP_{2D} Auction

The first extension of *GSP* for the two-dimensional setting, NP_{2D} , is based on the simple “next-price” rule of *GSP*—every winner has to pay the minimum amount necessary to keep his position. In *GSP*, the next price is the bid of the next-highest bidder. In the two-dimensional NP_{2D} , however, maintaining one’s position consists of two things for a winner in outcome N : first, the outcome must remain N and not switch to E ; second, the bid must enable the bidder to maintain

his position amongst the k slots. In a next price auction for our setting, therefore, the payment of a winner in slot i of outcome N is the larger of two terms—the first being the minimum value at which the outcome still remains N , and the second being the bid of the next bidder, b_{i+1}^N , as in *GSP*. The NP_{2D} auction with reserve prices is formally defined below.

Definition 1 (The NP_{2D} auction) Let $\Gamma = \sum_{i=1}^k \theta_i \max(b_i^N, r)$. The mechanism NP_{2D} compares $b_{e_1}^E$ to Γ to decide whether the outcome should be E or N .

- If $b_{e_1}^E \geq \Gamma$, the outcome is E with per-click payment

$$\max(b_{e_2}^E, \sum_{i=1}^{e_1-1} \theta_i \max(b_i^N, r) + \sum_{i=e_1}^k \theta_i \max(b_{i+1}^N, r)).$$

- If $b_{e_1}^E < \Gamma$, the outcome is N and the bidder winning slot $i \neq e_1$, if his bid is at least r , pays

$$\theta_i p_i = \max(\theta_i \max(b_{i+1}^N, r), b_{e_1}^E - \Gamma + \theta_i \max(b_i^N, r)),$$

while the bidder e_1 , if his bid is at least r , winning slot e_1 pays

$$\theta_{e_1} p_{e_1} = \max(\theta_{e_1} \max(b_{e_1+1}^N, r), b_{e_2}^E - \Gamma + \theta_{e_1} \max(b_{e_1}^N, r)),$$

where p_i is per-click payment of bidder i , and $\theta_i p_i$ is the expected payment of bidder i every time the keyword is searched.

The GSP_{2D} Auction

The second extension of GSP for the two-dimensional setting, GSP_{2D} , is defined to ensure the practical benefit that when multiple ads are displayed, advertisers see no difference at all between the new auction and the existing *GSP* system. In other words, GSP_{2D} is designed to restrict the allocation and pricing to be exactly those of *GSP* whenever multiple ads are shown (the outcome is N). It remains to decide whether to show outcome N or outcome E , as well as the pricing and allocation for outcome E . A formal definition of these rules is as follows.

Definition 2 (The GSP_{2D} auction) The mechanism GSP_{2D} compares $b_{e_1}^E$ to $\sum_{i=2}^{k+1} \theta_{i-1} \max(b_i^N, r)$ to decide whether the outcome should be E or N .

- If $b_{e_1}^E \geq \sum_{i=2}^{k+1} \theta_{i-1} \max(b_i^N, r)$, the outcome is E with winning bidder e_1 , whose payment is $\max(\sum_{i=2}^{k+1} \theta_{i-1} \max(b_i^N, r), b_{e_2}^E)$ per click.
- If $b_{e_1}^E \leq \sum_{i=2}^{k+1} \theta_{i-1} \max(b_i^N, r)$, assign the page to the bidders with k highest N -bids as long as their N -bids are at least r , and charge them according to GSP pricing, i.e. bidder i (if assigned any slot) has to pay $\max(b_{i+1}^N, r)$ per click.

Note that the highest N -bid is ignored when deciding between N and E . Perhaps, the most natural choice would be to compare $b_{e_1}^E$ and $\sum_{i=1}^k \theta_i \max(b_i^N, r)$ for choosing between N and E . However, one can easily see that with such rule, outcome E almost never happens in equilibrium. If advertiser 1 (with highest N -bid) is different from advertiser e_1 (with highest E -bid), advertiser 1 can submit a very large N -bid to make sure that $\sum_{i=1}^k \theta_i \max(b_i^N, r) > b_{e_1}^E$, i.e., the outcome is N . Note that the payment of advertiser 1 is set by GSP in outcome N , i.e., he does not have to pay anything more than b_2^N (the same as before) for submitting such a large bid. This bad equilibrium behavior is fixed by ignoring the highest N -bid.

4 Analytical Study

We build insights into the different auction mechanisms by considering a simplified analytical model which is a special case of the general framework described in the previous section. The insights derived from this simplified analytical exercise are valid in the fully general framework; we choose the simpler model because it allows us to bring to light the key insights through a tractable, closed-form analysis. In Section 5, we report the results of a numerical simulation study using the full model which confirms the results derived in this section.

4.1 Simplified Model

We assume that there are three bidders (i.e., $n = 3$) labeled A, B and C with valuation vectors (a, a) , (b, b) and $(0, c)$, respectively, where $a, b, c > 0$. In other words, the valuation per click of bidders A and B remains the same whether their ads are listed with other ads or listed exclusively. Bidder C, however, only values being listed exclusively.⁸ Without loss of generality, we assume that

⁸We choose this special case because it helps us to obtain the focal insights in a simple way. Other assumptions on valuations are possible (e.g., bidders A and B have valuations (a, a') and (b, b') , respectively, with $a' \geq a$ and $b' \geq b$,

$a \geq b$. We also assume that if the outcome is N , there are two slots available (i.e., $k = 2$). Recall that we have already assumed that the CTR of the only slot in outcome E is 1. To be conservative in our evaluations of the two-dimensional auctions, we assume that the CTR of the first position in N is also 1. This implies that the CTR of the second position in N is p , so that the total expected clicks in N are $1 + p$, which is more than the total expected clicks in E . We are only interested in the case that $b \geq r$ which means that the bidders are both competitive enough and can beat the reserve price; otherwise, the bidder will be completely ignored by the mechanism.

Bidder i submits a two-dimensional bid (b_i^N, b_i^E) where b_i^N is how much he bids for being shown among the others, and b_i^E is how much he bids to appear exclusively. We calculate the equilibrium revenue and efficiency of GSP , GSP_{2D} and NP_{2D} . However, there are multiple Nash equilibria for each of these mechanisms. This makes comparisons among the mechanisms very hard since, in most of the cases, all mechanisms have good and bad equilibria in terms of revenue and efficiency. Hence, we impose the following reasonable equilibrium refinements to reduce the number of equilibria that we obtain.

1. Losers bid at least their true value: Certain Nash equilibria can exist in which the outcome is E and the losing bidders bid $b_i^N < v_i^N$. In such a case, it is unreasonable to expect that the losing bidders bid low to maintain an equilibrium in which they are losing. This refinement helps us to rule out such equilibria.⁹
2. Winners bid the lowest in their best-response set: If two different bids result in the same outcome for a winner, we assume that he chooses the lower one. In other words, as long as a bidder is getting his desired outcome, he prefers to bid the lowest value that guarantees him the same outcome.

Moreover, we assume throughout the paper that the bidders do not play weakly dominated strategies. The above refinements allow us to make easier analytical comparisons among the mechanisms in terms of revenue and efficiency. We also assume that the ties are broken in favor of the

bidder C has valuations (c, c) or (c, c') with $c' \geq c$, etc.) but these provide the same insights. The simulation study covers such cases.

⁹Note that in an equilibrium where losers are bidding below their true values, no loser can single-handedly change the outcome to N ; however, if each loser increases his bid to his true value (note that bidding true value is also a best response for him) the outcome may eventually change to N . Therefore, losing bidders weakly prefer to bid not below their true valuations.

stronger bidder (bidder with higher valuation), and between E and N we assume that the ties are broken in favor of outcome E .

4.2 Analysis

In this section, we derive the equilibrium revenue and social welfare for the different auctions in three lemmas.

We start with GSP . Note that bidders A and B are the only ones interested in a non-exclusive outcome; therefore, one of them wins the first slot and the other one wins the second slot. The winner of the second slot pays the reserve price r per click, no matter what he bids. Furthermore, the winner of the second slot would still win the second slot if he decreases his bid all the way down to the reserve price r . Therefore, by the second equilibrium refinement, we assume that the winner of the second slot bids r ; this makes the per-click price of the first slot also r . As a result, the revenue of GSP is $r + rp$. However, note that the winner of the first slot could be either A or B—as long as the higher bid is high enough so that the winner of the second slot does not want to deviate to the first slot, we have an equilibrium. This could happen with either A or B being the highest bidder. This gives us the following lemma.

Lemma 1 *The revenue of GSP is $r + rp$, and its social welfare is $a + bp$ or $b + ap$.*

Note that there are multiple equilibria for GSP , but all of them lead to the same revenue $r + rp$. In case of welfare, we will show that the welfare of two-dimensional auctions is generally better than the one-dimensional auction. To favor the one-dimensional auction in this comparison, we pick the equilibrium of GSP which has the higher welfare, $a + bp$.

Next, we analyze NP_{2D} . First, consider the case where the exclusive bid of C is high enough to win the E outcome. In this case, C has to pay at least $a + bp$ per-click because A and B are losing and hence are bidding at least their true value. Moreover, he does not have to pay more than $a + bp$, because if he bids $a + bp$, the other two bidders cannot change the outcome to N without incurring negative utility.

Similarly, if $c < a + bp$, bidders A and B can force the N outcome; in this case, if c is small, specifically $c \leq r + rp$, the revenue will be $r + rp$ as in the GSP case. However, if $c > r + rp$, because of the NP_{2D} payment scheme, bidders A and B must pay highly enough to keep the outcome N ,

which means that their total payment would be c . This gives us the following lemma. A detailed proof is presented in the appendix.

Lemma 2 *The revenue of NP_{2D} is $\min(\max(c, r+rp), a+bp)$ and its social welfare is $\max(c, a+bp)$.*

Finally, we analyze GSP_{2D} . First, consider the case where the exclusive bid of C is high enough to win the E outcome. In this case, C must be bidding at least $a + rp$ because otherwise A and B could change the outcome to N (for example, by both bidding a). Also, if C is bidding $a + rp$, the other bidders cannot change the outcome to N without incurring negative utility. According to GSP_{2D} , bidder C's payment in this case is $\max(b + rp, a)$, and the social welfare is c .

If C cannot win the E outcome, he will be a loser. The outcome however could still be N or E . If the outcome is E , A would be the winner with payment $\max(c, b, r + rp)$. If the outcome is N , the situation would be the same as in GSP except that bidders A and B must make sure that C cannot change the outcome to E . This pushes the bids of A and B up and guarantees the revenue of the outcome to be at least c . Therefore, when the outcome is N , the revenue is $\max(c, r + rp)$. Clearly, if A is the exclusive winner, the social welfare is a , otherwise, as in GSP , the social welfare is $a + bp$. This gives us the following lemma. A detailed proof is presented in the appendix.

Lemma 3 *If $c \geq a + rp$, then the revenue of GSP_{2D} is $\max(b + rp, a)$ and its social welfare is c . If $c < a + rp$, the revenue of GSP_{2D} is either $\max(c, b, r + rp)$ or $\max(c, r + rp)$ and its social welfare is either a or $a + bp$, respectively.*

Note that there are multiple equilibria in some cases for GSP_{2D} . In particular, if c is small, there are two possible equilibria: (i) the outcome is N , and A wins the first slot and B wins the second slot, (ii) A wins E . The revenue of the first equilibrium is $\max(c, r + rp)$, while the revenue of the second equilibrium is $\max(b, c, r + rp)$.

4.3 Results and Insights

A representative comparison among the revenues of the auctions as calculated in the previous section is shown in Figure 1(a). The figure shows that the revenue of either NP_{2D} or GSP_{2D} always dominates the revenue of GSP . Therefore, it is always better for the search engine to


use a two-dimensional auction allowing exclusive display instead of the currently prevailing one-dimensional GSP auction. However, which one of these two auctions performs better? Figure 1(a) shows that, generally speaking, NP_{2D} has higher revenue; however, the revenue of GSP_{2D} *might* be better than the revenue of NP_{2D} for small values of c . (Note that $c < a + rp$ is the region in which GSP_{2D} has multiple equilibria. The dashed green line corresponds to the higher-revenue equilibrium; the revenue of the other equilibrium of GSP_{2D} is not necessarily more than the revenue of NP_{2D} .) We state this below as a proposition.

Proposition 1 *If $b > r + rp$, then if c is below a threshold value (specifically, $c < b$) then GSP_{2D} can provide highest revenue for the search engine, and if c is above this threshold value then NP_{2D} provides highest revenue for the search engine. If $b < r + rp$ then NP_{2D} always provides highest revenue for the search engine. Either GSP_{2D} or NP_{2D} always provides higher revenue than GSP .*


There are two primary reasons why exclusive-display auctions provide higher revenue for the search engine. First, NP_{2D} and GSP_{2D} allow advertisers to express their valuations for exclusive display while GSP ignores these valuations. Clearly, if the valuation for exclusivity is large enough even for one advertiser, allowing this advertiser to express this valuation through his bid will switch the outcome to E from N and increase search engine revenue. This is the case when the value of c is large.

However, there is a second, and more interesting, phenomenon at play which leads to higher revenue in auctions allowing exclusive display. Note that even when the outcome of NP_{2D} is N (which happens if $c < a + bp$, which is the shaded region in Figure 1(a)), we can still see higher revenue in NP_{2D} than in GSP . In other words, even if exclusive display is not the equilibrium outcome, auctions allowing bidders to bid for exclusive display can provide higher revenue to the search engine than the GSP auction. In fact, we can see that the N outcome in GSP_{2D} can also generate higher revenue for the search engine than GSP , even though in this case, by the definition of GSP_{2D} , the bidders will be ranked in GSP_{2D} exactly as they would be ranked in GSP .


To understand the intuition behind this, consider the case in which there is a bidder who values the exclusive-display outcome more than the multiple-display outcome, but this valuation for exclusivity is not too high (which is true for medium values of c). In this case, the other bidders who value the N outcome close to the E outcome have to bid higher than they do in GSP to



(a) Search Engine Revenue



(b) Social Welfare



(c) Fraction of Social Welfare as Search Engine Revenue

Figure 1: Revenues, social welfare and fraction of social welfare extracted as revenue as functions of c with parameters $p = 0.6$, $a = 3$, $b = 2$ and $r = 1$. GSP_{2Dh} denotes the higher-revenue equilibrium of GSP_{2D} , defined only for $c < a + rp$. The shaded region in panel (a) shows the values of c below which the outcome of NP_{2D} is N .

actually keep the outcome as N (i.e., to prevent the outcome from becoming E , in which case they will not be displayed and be worse off), which leads to higher revenue. Said in another way, two-dimensional auctions give more degrees of freedom to the bidders which also increases the competition among them. This increased competition, in equilibrium, leads to higher revenue for the search engine. We state this below as a proposition.

Proposition 2 *Even if the outcome of a two-dimensional auction (NP_{2D} or GSP_{2D}) allowing for exclusive display of an ad is such that multiple ads are actually displayed, it can generate higher revenue for the search engine than the GSP auction due to greater competition among the bidders.*

Note that an interesting property of the revenues of NP_{2D} and GSP_{2D} auctions is that NP_{2D} is revenue monotone in valuation for clicks while GSP_{2D} is not. In other words, if the valuation of a bidder for clicks increases, the revenue of NP_{2D} will not decrease, while this is not true for GSP_{2D} . In fact, from Figure 1(a) it is clear that the revenue of GSP_{2D} drops down at some point as c increases.

Another interesting and important metric to study for these auctions is the social welfare, which is the sum of the search engine profits and the bidders' surplus. To study this, we define a new two-dimensional auction called VCG_{2D} . VCG_{2D} is an extension of the one-dimensional social-welfare-maximizing VCG auction, and maximizes social welfare in the two-dimensional setting because the outcome is chosen as E or N and slots are subsequently allocated to the advertisers based on who values them more. We define and analyze VCG_{2D} in the appendix, and find that its social welfare is given by $\max(c, a + bp)$. Interestingly, this is exactly the social welfare of the NP_{2D} auction. This implies that, for our simplified model, NP_{2D} is a social-welfare-maximizing auction. The other auctions, GSP and GSP_{2D} , also achieve the maximum social welfare for some regions of the parameter space, but there are other regions of the parameter space where they do not. This is shown with the help of Figure 1(b).

For small and medium c , compared to GSP , both NP_{2D} and GSP_{2D} can increase search engine revenue without increasing social welfare, which implies that two-dimensional auctions help the search engine to extract more from the bidders. Furthermore, for large c , the social welfare increases, and both search engine revenue and bidders' surplus increase simultaneously. These patterns are clear if we inspect Figures 1(a) and 1(b) simultaneously. We state the interesting

results in the following proposition.

Proposition 3 *Two-dimensional auctions can simultaneously increase both search engine revenue and bidders' surplus.*

Next, we investigate at how good the two-dimensional setting performs in extracting as much revenue as *possible*. For this, we use the fact the the revenue of a mechanism in equilibrium can never be more than the social welfare, otherwise, some bidders would incur negative utility which is against the definition of equilibrium and individual rationality. Therefore, the ratio of search engine revenue to the maximum possible social welfare (obtained from the VCG_{2D} auction) is a measure of how good an auction is at extracting as much revenue as possible.

Figure 1(c) plots the fraction of the maximum possible revenue extracted by the various auctions considered. It shows that two-dimensional auctions not only extract more revenue than GSP in absolute value, but also the fraction of the total possible revenue extracted is larger than the fraction extracted by GSP .¹⁰ By looking at the curve of the fraction of possible revenue extracted by NP_{2D} , we see that the curve increases up to the value 1 (which is the point $c = a + pb$), which implies that at this point NP_{2D} extracts *all* of the social welfare as profit for the search engine. Subsequently, as c increases, this fraction starts to decrease. The interpretation is as follows. As c increases from a small value, the competition between bidder C who wants outcome E and bidders A and B who want outcome N becomes more intense, and the increased competition makes the revenue larger. The point at which the search engine can extract all of the social welfare as its revenue is the point at which its revenue from both N and E is the same and it is indifferent between choosing N or E . However, as c increases further so that it becomes large enough such that the optimal outcome for the search engine is E , the extracted revenue remains constant with increasing c . This is because, due to the next-price characteristic of the auctions, the revenue is determined by the bids of the other advertisers who do not value exclusivity highly, while the maximum possible revenue, which is now c , keeps increasing. Therefore, the ratio of the extracted revenue to the maximum possible revenue becomes smaller. We state this result as a proposition.

¹⁰Note that two-dimensional auctions extract a larger fraction of the total possible revenue even when, to calculate the fraction for GSP , we use the social welfare of the one-dimensional setting as the denominator. This comparison normalizes the effect that only two-dimensional auctions (but not GSP) benefit from increasing c . Still, we see that NP_{2D} and GSP_{2D} perform better.

Proposition 4 *The search engine can use two-dimensional auctions to extract a larger fraction of the social welfare as its profit (as compared to the one-dimensional auction GSP). Furthermore, under the condition $c = a + bp$, when the search engine is indifferent between the outcomes N and E , NP_{2D} extracts all of the social welfare as profit for the search engine.*

From the results above, we see that if $c > a + bp$, then as it becomes larger, the search engine only makes the profit $a + bp$ and therefore also extracts a progressively smaller fraction of the social welfare. (In the next section, we will see that having multiple bidders with large per-click valuations for exclusive placement mitigates this problem to a large extent.) However, a solution to this problem is to move to a new mechanism which only allows outcome E . This mechanism invites only one-dimensional bids for outcome E and dominates NP_{2D} in terms of revenue for large values of c by setting a high-enough reserve price. Therefore, interestingly, as c becomes very large, it is better for the search engine to use a one-dimensional exclusive-only bidding mechanism.

Overall, we can summarize the insights from our analytical study as follow. First, two-dimensional auctions that allow advertisers to bid for exclusive display in the sponsored search section in response to a user query have the potential to increase search engine profits. This may happen even if the equilibrium outcome is to actually display multiple ads. Generally speaking, the NP_{2D} auction provides highest profit. Second, two-dimensional auctions can increase social welfare. Furthermore, how social welfare is split between the search engine and the advertisers depends on the value attached to exclusive display by advertisers. In fact, the NP_{2D} auction can appropriate all of the welfare for the search engine and leave the bidders with zero surplus, which the search engine is unable to achieve using the one-dimensional GSP auction. Under other conditions, the surplus of the bidders can also increase. Finally, the good revenue and efficiency properties of NP_{2D} suggest that the simple “next price” heuristic of GSP is a good heuristic for designing two-dimensional position auctions as well.

5 Simulation Study

In this section, we conduct simulations for exclusive-display auctions that reflect more realistic situations (as compared to our simplified analytical model). The simulations support the results


from our analytical study and also add new insights.

5.1 Design of the Study

We use the general model defined in Section 3 and assume that there are $n = 10$ advertisers and $k = 5$ slots in the non-exclusive outcome. Previous research has found that click-through rates decrease exponentially with descending position (Feng et al. 2007). In accordance with this, we assume that $\theta_i = p^{i-1}\theta_1$ for $i \in \{2, 3, \dots, k\}$, where $0 < p < 1$. We define $\theta_i = 0$ for $i > k$. We assume that the CTR of the first slot in the non-exclusive outcome is $\theta_1 = p, 0 < p \leq 1$. This implies that the CTR of slot i in the non-exclusive outcome is $p^i, i \in \{1, 2, 3, 4, 5\}$, which in turn implies that the total expected clicks in the non-exclusive outcome are $\sum_{i=1}^5 p^i$. We use three values of p in our simulations, specifically, 0.55, 0.7 and 0.85, which implies that the total number of clicks in the non-exclusive outcome are 1.16, 1.94 and 3.15, respectively. Recalling that the CTR of the only slot in the exclusive outcome is normalized to $\hat{\theta} = 1$, we note that, for the chosen values of p , the total number of clicks in the non-exclusive outcome, compared to the exclusive outcome, are about the same, about double and about triple, respectively.

The valuation of each advertiser j for the non-exclusive outcome, v_j^N , is drawn uniformly randomly between $[0, 1]$. The valuation of each advertiser for the exclusive outcome is set to $v_j^E = h_j v_j^N$, where $h_j \geq 1$ is a ‘‘heterogeneity multiplier’’ and is set in the following manner. For some given parameter $a > 0$, a value h_j is drawn independently randomly for each advertiser with probability distribution defined by $\Pr(h_j > 1+x) = (2-ax)^2/4, 0 \leq x \leq 2/a$. This implies that $h_j \in [1, 1+2/a]$ and the probability density function for h_j is a right-angled triangle with height a at 1 and height 0 at $1 + 2/a$. We show this in Figure 2. Given this formulation, a smaller value of a implies that the value of h_j can be drawn from a wider domain. In other words, a smaller value of a implies that the heterogeneity among the advertisers’ valuations for the exclusive slot is greater and the exclusive slot is also valued more on average. We use 39 different values of a between 3 and 0.01 (shown in Table 1), which implies that the domains of h_j vary between $[1, 1.67]$ (small heterogeneity) and $[1, 201]$ (large heterogeneity), respectively. We set the reserve price of this auction to $r = 0$.

Given the three values of p , and the 39 values of a , we get 117 different scenarios. For each scenario, we simulate the outcomes of the *GSP*, *NP_{2D}* and *GSP_{2D}* auctions 100 times. Before proceeding further, we note that the main differences between the analytical study and the simula-



(a) Uniform distribution for sampling v_j^N (b) Triangular distribution for sampling h_j

Figure 2: A uniform distribution to sample v_j^N and a triangular distribution to sample h_j , where $v_j^E = h_j v_j^N$. Smaller values of a imply larger mean and variance (heterogeneity) in the exclusive-placement valuations.

tion study are that there are more bidders and each bidder has randomly-drawn valuations for the non-exclusive and exclusive outcomes from the same distributions.

5.2 Computing the Equilibrium

We adopt and extend the procedure in Cary et al. (2008) to find the equilibrium in a given auction. This procedure is as follows. We start with all bids, exclusive and non-exclusive, being zero. At each step, an advertiser is selected uniformly at random. This advertiser updates his strategy by playing best response to the other advertisers' current bids. This best response is found by a brute force technique—the advertiser tries different bids that will determine the configuration and put him in the different slots available in the configuration, and chooses the bid that maximizes his payoff. If he does not win any slot when playing his best response strategy, we make him bid at least his true value. (The reason for this is that we want to avoid those equilibria in which some bidders are losing but still bidding below their true valuation to maintain the equilibrium. Note that, in this situation, bidding his true valuation belongs to the set of best response strategies of the losing advertiser.) We continue repeating the above step until we reach an equilibrium, in which case no advertiser wants to update his strategy any further and all losers are bidding at least their true value.¹¹

¹¹Note that this procedure only gives us equilibria in pure strategies. We do not consider equilibria in mixed strategies, a choice which most papers on sponsored search auctions make (Athey and Ellison 2011, Edelman and Ostrovsky 2007, Edelman, Ostrovsky and Schwarz 2007, Varian 2007).

5.3 Results and Insights

We summarize the results for the simulations with $p = 0.7$ in Table 1 and Figure 3.¹² In Table 1, the column titled “ h ” shows the domain of h for the corresponding value of a in the leftmost column. For NP_{2D} and GSP_{2D} , the column titled “E%” shows what percent of time the auction outcome is the exclusive outcome, the column titled “Rev” shows the revenue from the auction, the column titled “Welf” shows the social welfare from the auction and the column titled “Frac” shows the fraction of revenue to social welfare, each quantity averaged over 100 runs. (The “Frac” column for the GSP auction is calculated by dividing the revenue of GSP by the social welfare of NP_{2D} .)

The results illustrate that two-dimensional auctions—especially NP_{2D} —give greater revenue to the search engine than GSP . Moreover, this greater revenue is obtained even if the outcome of the two-dimensional auctions is non-exclusive display, for the reasons explained in Section 4. The social welfare is higher than GSP , and the fraction of social welfare extracted as revenue is also higher. These results get stronger as the value of a decreases (going down the rows in Table 1).

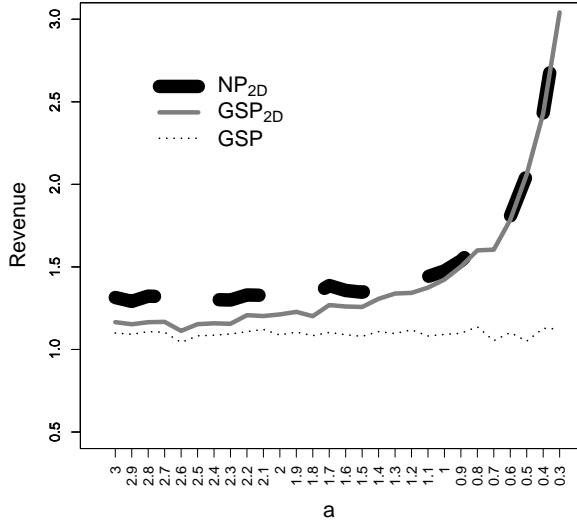
An interesting observation is that, for large a , the exclusive outcome is obtained more often for GSP_{2D} . The reason is that, the manner in which GSP_{2D} is defined, it ignores the highest non-exclusive bid. For large a , this is actually a disadvantage because it unnecessarily prefers the exclusive outcome, even when the non-exclusive outcome is the better one for the search engine. This disadvantage of GSP_{2D} is reflected in a significantly lower revenue than NP_{2D} for large values of a . In fact, this disadvantage can lead to a lower social surplus of GSP_{2D} than even GSP for large a , as observed in Figure 3(b). GSP , on the other hand, ignores the exclusive bids completely. NP_{2D} does not suffer from either problem, and always provides the highest revenue. As a decreases, the valuations for exclusivity (and therefore the bids for exclusivity) are high enough that it does not hurt to ignore the highest non-exclusive bid; therefore, GSP_{2D} has better revenue properties for small a .

Another interesting point to note here is that our analytical study showed that GSP_{2D} has two equilibria when valuations for the exclusive outcome are small, and one of these equilibria has higher revenue than NP_{2D} (Proposition 1). Our simulation study, however, consistently shows higher revenue for NP_{2D} on average. Upon deeper investigation, we find that the higher-revenue


¹²The results from simulations with $p = 0.55$ and $p = 0.85$ are qualitatively similar and we omit them from the paper. They can be obtained from the authors upon request.

a	h	NP_{2D}				GSP_{2D}				GSP		
		E%	Rev	Welf	Frac	E%	Rev	Welf	Frac	Rev	Welf	Frac
3	[1, 1.67]	3%	1.31	1.52	86%	17%	1.17	1.49	76%	1.1	1.52	72%
2.9	[1, 1.69]	4%	1.29	1.51	85%	26%	1.15	1.46	76%	1.09	1.51	72%
2.8	[1, 1.71]	8%	1.32	1.52	86%	31%	1.17	1.48	76%	1.11	1.53	72%
2.7	[1, 1.74]	5%	1.32	1.53	86%	24%	1.17	1.49	76%	1.1	1.53	72%
2.6	[1, 1.77]	9%	1.27	1.49	84%	33%	1.11	1.45	74%	1.04	1.49	69%
2.5	[1, 1.8]	5%	1.31	1.52	86%	32%	1.15	1.46	75%	1.08	1.52	71%
2.4	[1, 1.83]	10%	1.3	1.51	86%	40%	1.16	1.46	76%	1.09	1.51	71%
2.3	[1, 1.87]	17%	1.3	1.51	85%	37%	1.16	1.47	76%	1.09	1.5	72%
2.2	[1, 1.91]	13%	1.33	1.55	85%	48%	1.21	1.49	77%	1.11	1.54	71%
2.1	[1, 1.95]	21%	1.33	1.54	85%	45%	1.2	1.5	77%	1.12	1.53	72%
2	[1, 2]	14%	1.34	1.54	86%	46%	1.21	1.49	78%	1.09	1.53	70%
1.9	[1, 2.05]	28%	1.33	1.53	86%	57%	1.23	1.49	80%	1.11	1.52	72%
1.8	[1, 2.11]	25%	1.33	1.52	87%	61%	1.2	1.46	78%	1.08	1.51	70%
1.7	[1, 2.18]	41%	1.39	1.57	88%	68%	1.27	1.54	80%	1.1	1.55	70%
1.6	[1, 2.25]	37%	1.36	1.56	86%	68%	1.26	1.52	80%	1.09	1.52	69%
1.5	[1, 2.33]	36%	1.35	1.55	87%	68%	1.26	1.5	81%	1.08	1.51	69%
1.4	[1, 2.43]	47%	1.38	1.62	85%	73%	1.31	1.59	80%	1.11	1.54	68%
1.3	[1, 2.54]	58%	1.41	1.65	86%	81%	1.34	1.62	81%	1.1	1.52	67%
1.2	[1, 2.67]	62%	1.42	1.66	86%	80%	1.34	1.64	81%	1.12	1.53	67%
1.1	[1, 2.82]	67%	1.44	1.72	84%	88%	1.37	1.69	80%	1.08	1.53	64%
1	[1, 3]	74%	1.48	1.74	86%	93%	1.42	1.71	82%	1.09	1.52	64%
0.9	[1, 3.22]	80%	1.54	1.9	82%	95%	1.5	1.88	80%	1.1	1.54	59%
0.8	[1, 3.5]	89%	1.63	2.07	80%	97%	1.6	2.06	78%	1.14	1.55	56%
0.7	[1, 3.86]	96%	1.64	2.15	78%	100%	1.6	2.15	77%	1.05	1.48	51%
0.6	[1, 4.33]	94%	1.81	2.3	80%	99%	1.79	2.29	79%	1.1	1.52	50%
0.5	[1, 5]	96%	2.06	2.63	80%	100%	2.06	2.64	80%	1.05	1.49	42%
0.4	[1, 6]	99%	2.43	3.22	78%	99%	2.43	3.22	78%	1.13	1.56	38%
0.3	[1, 7.67]	100%	3.04	4.26	73%	100%	3.04	4.26	73%	1.12	1.55	28%
0.2	[1, 11]	100%	3.85	5.49	72%	100%	3.85	5.49	72%	1.09	1.52	21%
0.1	[1, 21]	100%	7.9	10.63	76%	100%	7.9	10.63	76%	1.1	1.5	11%
0.09	[1, 23.22]	100%	8.33	11.75	73%	100%	8.33	11.75	73%	1.1	1.5	10%
0.08	[1, 26]	100%	8.85	12.36	73%	100%	8.85	12.36	73%	1.09	1.51	10%
0.07	[1, 29.57]	100%	10.25	14.29	74%	100%	10.25	14.29	74%	1.07	1.5	8%
0.06	[1, 34.33]	100%	11.24	15.7	73%	100%	11.24	15.7	73%	1.09	1.54	8%
0.05	[1, 41]	100%	12.54	18.46	71%	100%	12.54	18.46	71%	1.06	1.49	7%
0.04	[1, 51]	100%	16.5	24.42	70%	100%	16.5	24.42	70%	1.08	1.51	5%
0.03	[1, 67.67]	100%	23.5	34.4	69%	100%	23.5	34.4	69%	1.08	1.53	3%
0.02	[1, 101]	100%	33.85	48.64	72%	100%	33.85	48.64	72%	1.06	1.51	2%
0.01	[1, 201]	100%	68.86	97.51	71%	100%	68.86	97.51	71%	1.11	1.53	1%


Table 1: Outcomes of auctions for $p = 0.7$.



(a) Search Engine Revenue



(b) Social Welfare



(c) Fraction of Social Welfare as Search Engine Revenue

Figure 3: Revenues, social welfare and fraction of social welfare extracted as revenue as functions of a with parameter $p = 0.7$. The x -axis is not to scale.

equilibrium of GSP_{2D} indeed occurs in the simulation, but does not occur sufficiently often to lead to higher average revenue than NP_{2D} . This indicates that, in general, NP_{2D} might be the better choice of auction in terms of revenue.

Note that in the simulation setup there is more competition among bidders (than in the simplified analytical setup in Section 4). This is because in the model in Section 4, two bidders (A and B) have the same per-click valuation for multiple and exclusive display and only bidder C values exclusive display more than multiple display. In the simulation, each bidder has the pair of per-click valuations drawn from the same distributions (which implies that all bidders value exclusive display more than multiple display, albeit to different extents) and there are more bidders.

This greater competition manifests itself in two main differences in the outcomes as shown in Figures 1 and 3. To facilitate these comparisons, note that, loosely speaking, a smaller value of a has the same effect on the valuation for the exclusive outcome as a larger value of c in Section 4. First, note from the revenue plots that as a decreases the revenues increase from the two-dimensional auctions in Figure 3(a), which does not happen as c increases in Figure 1(a). This is because, in the simulation, there are many bidders with high valuations for the exclusive outcome, not just one. This implies that the winning bid does not just have to beat the non-exclusive bids, but also other comparable exclusive bids. Therefore, the exclusive bids are higher and so is the revenue. Second, note from the plots showing the fraction extracted that the peak in Figure 3(c) is less pronounced than the peak in Figure 1(c). The peak occurs at the point at which the search engine is indifferent between choosing outcome N or E . With many bidders in the simulated auctions, the fraction of social welfare extracted as revenue is quite large for all values of a . Therefore, the peak is not very distinct. However, as a decreases, the fraction of social welfare extracted as revenue from GSP falls sharply because it completely ignores exclusive-placement valuations.

Overall, observations from the simulation study parallel those from the analytical study. In fact, allowing for many bidders makes an even stronger case for two-dimensional auctions because the revenue and the fraction of social welfare extracted as revenue are both consistently significantly higher than in GSP .

6 Conclusions

Most search engines run auctions to price the ranked list of ads presented to a user in response to a keyword search. In the last decade, the type of auction used has evolved from a Generalized First Price auction to a Generalized Second Price (*GSP*) auction with numerous small adjustments, and search engines continue to explore new auction mechanisms that can improve revenue. Recently, the three largest search engines, namely Google, Yahoo! and Bing, have been experimenting with the idea of allowing an advertiser to bid to display his ad exclusively rather than in a list of multiple ads. Advertisers can be expected to value exclusive display more because they can create strong branding effects by being displayed exclusively with a keyword, and avoid negative externalities that may impact post-click actions of consumers due to the presence of other ads in a ranked list.

In this paper, we study two auctions, NP_{2D} and GSP_{2D} , recently patented by Yahoo! as the primary candidates for implementation as exclusive-display auctions. In these auctions, each advertiser submits a two-dimensional bid, one for a multiple-display format and another for an exclusive-display format. We find that these exclusive-display auctions always generate higher revenue than the *GSP* auction. This is easy to see if some advertiser values exclusive display very highly—he will submit a high exclusive-display bid and the search engine will choose the higher-revenue exclusive display outcome. Interestingly, however, search engine revenue can increase even if the final outcome is multiple display. This is because advertisers are competing not only for ranks within the multiple-display outcome, but are also competing for the outcome itself. Therefore, if an advertiser values exclusive display more than multiple display, other advertisers who want the multiple display outcome (because they themselves cannot bid highly for being displayed exclusively, and will get nothing if someone else gets displayed exclusively) will have to bid higher than in *GSP* to make the multiple-display outcome more profitable for the search engine. We also run numerical simulations which show that the revenue advantage of two-dimensional auctions over *GSP* is especially high with a large number of advertisers.

We conjecture that if search engines adopt two-dimensional exclusive display auctions then, for a large majority of keywords, the outcome will still be multiple display while the search engine will make higher profits than in *GSP*. This is because, for most keywords, advertisers can be expected to value exclusive display more than multiple display which will lead to the competitive

effect explained above, but one advertiser may not value exclusive display so much that his bid by itself can dominate the total potential revenue from multiple advertisers combined. For other keywords, for instance, brand names, one advertiser may indeed value exclusive display higher enough than others that the outcome can very well be exclusive display. Such insights can help to resolve the dilemma that search engines face regarding whether or not to adopt exclusive-display auctions (Metz 2011).

We also find that two-dimensional auctions increase social welfare. Furthermore, they help a search engine to extract a larger percentage of the maximum possible revenue that can be extracted. In other words, the ratio of search engine revenue to social welfare is higher when a search engine employs two-dimensional auctions. Under other conditions, bidders' surplus can increase simultaneously with search engine revenue.

Among the NP_{2D} and GSP_{2D} auctions, we find that NP_{2D} , which is a next-price auction, has better revenue and allocative efficiency properties. This suggests that the simple next-price heuristic of the one-dimensional GSP is a good heuristic for designing two-dimensional position auctions as well.

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 auctions that have been proposed as candidates for implementation at Yahoo!. One of these auctions, NP_{2D} , has good revenue properties and can even extract the full social welfare as search engine revenue under some conditions. It is not clear, however, what is the optimal mechanism for an exclusive-display auction, and future research can explore this.

Second, we make the assumption that every advertiser values exclusive display at least as much as multiple display, which is a very reasonable assumption. However, these valuations are assumed to be independent across competitors to keep the model simple. 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. (2010) 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 multiple and exclusive display is simply one way to make the currently-prevailing auction format more expressive. However, there may be numerous other formats in which advertisers can reveal their preferences in more detail (e.g., Muthukrishnan (2009) discussed earlier). Future research can work towards a general theory of “expressive ad auctions.” This theory should also consider practical limitations such as ease of bidding by advertisers and the real-time calculation and implementation of auction outcomes by the search engine.

Appendix

Proofs of the Lemmas

Lemma A1 *The revenue of NP_{2D} is $\min(\max(c, r + rp), a + bp)$ and its social welfare is $\max(c, a + bp)$.*

Proof: First, consider the case where the exclusive-bid of C is high enough to win E outcome. In this case, C has to pay at least $a + bp$ per-click because A and B are losing and hence are bidding at least their true value; moreover, he does not have to pay more than $a + bp$, because if he bids $a + bp$, the other two bidders cannot change the outcome to N without incurring negative utility. In other words, suppose that C bids $a + bp$ and that the outcome is N with x and y being the non-exclusive bids ($x \geq y$.) For the outcome to be N , we must have $x + yp \geq a + bp$; this means that either A is bidding more than his true value or B is bidding more than his true value. Moreover, if the bidder who is bidding more than his value decreases his bid to the true value, he will become a loser. By definition of NP_{2D} , this bidder must actually be paying more than his true value which means that he is getting negative utility. Therefore, if C bids $a + bp$, he wins the E outcome. The social welfare is c in this case.

However, if $c < a + bp$, bidders A and B can force the N outcome by bidding their true value. In this case, if c is small, specifically $c \leq r + rp$, the revenue will be $r + rp$ as in the GSP case. However, if $c > r + rp$, because of NP_{2D} payment scheme, bidders A and B must pay highly enough to keep the outcome N , which means that their total payment would be c . In other words, if A and B are bidding x and y (with $x \geq y$), for N to be the outcome we must have $x + yp \geq c$, and since $x + yp = c$ is enough for the outcome to be N , by our equilibrium refinement, we assume $x + yp = c$. But in this case, according to NP_{2D} pricing, the bidder of x must pay exactly $\max(y, c - x - yp + x) = \max(y, c - yp) = x$ per-click, and the bidder of y must pay exactly $\max(rp, c - x - yp + yp)/p = \max(rp, c - x)/p = y$ per-click, meaning that the revenue will be $x + yp = c$. The social welfare in this case is $a + bp$. \square

Lemma A2 *If $c \geq a + rp$, then the revenue of GSP_{2D} is $\max(b + rp, a)$ and its social welfare is c . If $c < a + rp$, the revenue of GSP_{2D} is either $\max(c, b, r + rp)$ or $\max(c, r + rp)$ and its social welfare is a or $a + bp$, respectively.*

Proof: First, consider the case where the exclusive-bid of C is high enough to win E outcome. In this case, C must be bidding at least $a + rp$ because otherwise A and B could change the outcome to N : if A is a loser, he is bidding at least a , therefore, if B also bids a , the outcome switches to N without any of A or B being worse off. Also, if C is bidding $a + rp$, the other bidders cannot change the outcome to N without incurring negative utility: to change the outcome to N , the lower N -bid should be more than a which means that the per-click price of the winner of the first slot will be more than a which contradicts the outcome being an equilibrium. Therefore, according to GSP_{2D} , bidder C's payment in this case is $\max(b + rp, a)$ and the social welfare is c .

If C cannot win the E outcome, he will be a loser. The outcome however could still be N or E . If the outcome is E , A would be the winner with payment $\max(c, b, r + rp)$. If the outcome is N , the situation would be the same as in GSP except that bidders A and B must make sure that C cannot change the outcome to E . This means that, assuming that x is the lower N -bid among A and B, we must have $x + rp \geq c$; but at the same time, we know that $x + rp$ is the revenue of GSP_{2D} when the outcome is N . Therefore, when the outcome is N , the revenue is $\max(c, r + rp)$. Clearly, if A is the exclusive winner, the social welfare is a , otherwise, it is $a + bp$ or $b + ap$. As in GSP , we ignore the N equilibrium in which B gets the first slot. Therefore, the revenue is $\max(c, b, r + rp)$ or $\max(c, r + rp)$ and its social welfare is $a + bp$ or a , respectively. \square

The VCG_{2D} Auction

The VCG_{2D} auction is an extension of the one-dimensional VCG auction to two dimensions.

We assume that \max and \max^2 are the indices of the bidders with the highest and the second highest E -bids, i.e., $b_{\max}^E \geq b_i^E$ for any i , and $b_{\max^2}^E \geq b_i^E$ for any $i \neq \max$. Also, let \max_{-i} be the index of the bidder who has the highest E -bid other than i ; therefore, for $i \neq \max$ we have $\max_{-i} = \max$, and for $i = \max$ we have $\max_{-i} = \max^2$. We assume that the reserve price set by the auctioneer is r for both non-exclusive and exclusive bids.

Definition 3 (The VCG_{2D} auction) *The VCG_{2D} mechanism compares b_{\max}^E and $\sum_{i=1}^k \theta_i b_i^N$.*

- *If $b_{\max}^E \geq \sum_{i=1}^k \theta_i b_i^N$, VCG allocates the page to only one advertiser, namely \max , and charges him either the sum of the k highest $\theta_i b_i^N$'s (excluding himself) or the second highest E -bid,*

whichever is larger, i.e., the winner's payment is

$$\max(b_{\max^2}^E, \sum_{i=1}^{\max-1} \theta_i b_i^N + \sum_{i=\max}^k \theta_i b_{i+1}^N).$$

- If $b_{\max}^E < \sum_{i=1}^k \theta_i b_i^N$, then VCG allocation is N , but the expression for the payments is more complicated. When advertiser i is removed, the efficient reallocation can be either E or N . If it is E , the winner is \max_{-i} , and hence, the increase in the sum of the values of all advertisers other than i is $b_{\max_{-i}}^E - \sum_{j \neq i}^k \theta_j b_j^N$. If the efficient reallocation is N , all advertisers below i will move one slot up, therefore, the sum of their values increases by $\sum_{j=i}^k (\theta_j - \theta_{j+1}) b_{j+1}^N$. Therefore, the i -th advertiser's payment, $\theta_i p_i$, is (for $i \leq k$)

$$\max\left(\sum_{j=i}^k (\theta_j - \theta_{j+1}) b_{j+1}^N, b_{\max_{-i}}^E - \sum_{j=1}^k \theta_j b_j^N + \theta_i b_i^E\right).$$

VCG_{2D} has some notable characteristics. First, it is a truthful mechanism (i.e., the best strategy for the advertisers is to be truthful) which makes it a stable mechanism as well. Second, it maximizes social welfare because the outcome is chosen as E or N and slots are subsequently allocated to the advertisers based on who values them more. For the simplified model in Section 4, we obtain the following result for VCG_{2D} .

Lemma A3 *If $c > a + bp$, the revenue of VCG_{2D} is $a + bp$, otherwise, the revenue is $\max(c, a) + \max(c, b) - a - bp$. Its social welfare is $\max(c, a + bp)$.*

Proof: If $c \geq a + bp$, the outcome is E ; by the VCG_{2D} payment rule, the price is $a + bp$. If $c < a + bp$, the outcome is N ; by the VCG_{2D} payment rule, bidder A has to pay $\max(c, b) - bp$ and bidder B has to pay $\max(c, a) - a$. \square

References

- Agarwal, Nikhil, Susan Athey and David Yang (2009), “Skewed Bidding in Pay Per Action Auctions for Online Advertising,” *American Economic Review Papers and Proceedings*, 99 (2), 441–447.
- Athey, Susan and Glenn Ellison (2011), “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, forthcoming.
- Bradley, Stephen P. (2001), “eBay, Inc.” Harvard Business School Case 9-700-007, June 21, 2001.
- Branco, Fernando (1997), “The Design of Multidimensional Auctions,” *RAND Journal of Economics*, 28 (1), 63–81.
- Cary, Matthew, Aparna Das, Benjamin Edelman, Ioannis Giotis, Kurtis Heimerl, Anna Karlin, Claire Mathieu, and Michael Schwarz (2008), “On Best-Response Bidding in GSP Auctions,” Working Paper, Harvard University.
- Che, Yeon-Koo (1993), “Design Competition Through Multidimensional Auctions,” *RAND Journal of Economics*, 24 (4), 668–680.
- Desai, Preyas, Woochoel Shin and Richard Staelin (2010), “The Company that You Keep: When to Buy a Competitors Keyword,” Working Paper, Duke University.
- Dukes, Anthony and Esther Gal-Or (2003), “Negotiations and Exclusivity Contracts for Advertising,” *Marketing Science*, 22 (2), 222–245.
- Edelman, Benjamin and Michael Ostrovsky (2007), “Strategic Bidder Behavior in Sponsored Search Auctions,” *Decision Support Systems*, 43 (1), 192–198.
- Edelman, Benjamin, Michael Ostrovsky and Michael Schwarz (2007), “Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97 (1), 242–259.
- Edelman, Benjamin, and Michael Schwarz (2010), “Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions,” *American Economic Review*, 100 (2), 597–602.
- Feng, Juan, Hemant K. Bhargava and David M. Pennock (2007), “Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms,” *INFORMS Journal on Computing*, 19 (1), 137–148.
- Ghosh, Arpita and Amin Sayedi (2010), “Expressive Auctions for Externalities in Online Advertising,” *Proceedings of the Nineteenth International World Wide Web Conference*, 371–380.
- Hallowell, Roger (2002), “Monster.com: Success Beyond the Bubble,” Harvard Business School Case 9-802-024, January 7, 2002.
- Jerath, Kinshuk, Liye Ma, Young-Hoon Park and Kannan Srinivasan (2011), “A “Position Paradox” in Sponsored Search Auctions,” *Marketing Science*, forthcoming.
- Katona, Zsolt and Miklos Sarvary (2010), “The Race for Sponsored Links: Bidding Patterns for Search Advertising,” *Marketing Science*, 29 (2), 199–215.

- Liu, De, Jianqing Chen and Andrew B. Whinston (2010), “Ex Ante Information and the Design of Keyword Auctions,” *Information Systems Research*, 21 (1), 133–153.
- Metz, Cade (2011), “Google exec announces departure as Page takes reins,” *The Registrar*, April 5 2011. http://www.theregister.co.uk/2011/04/05/senior_jonathan_rosenberg_to_leave_google.
- Mori, Pier Angelo (2006), “Design of Reasonable Multidimensional Auctions Under Regulator Ignorance,” *International Public Procurement Conference Proceedings*, 21–23 September 2006, 781–811.
- Mullin, Joe (2010), “Trademark Lawsuit Against Googles AdWords Hits Appeals Court,” <http://paidcontent.org/article/419-trademark-lawsuit-against-googles-adwords-hits-the-appeals-court/>.
- Muthukrishnan, S. (2009), “Bidding on Configurations in Internet Ad Auctions,” *Proceedings of the 15th Annual International Conference on Computing and Combinatorics*, 1–6.
- Rayport, Jeffrey F. (2000), “Monster.com,” Harvard Business School Case 9-801-145, August 11, 2000.
- Sterling, Greg (2010), “Google Wins Legal Battle Over AdWords Trademark Issue In Europe,” <http://searchengineland.com/google-wins-major-legal-battle-over-adwords-trademark-issue-in-europe-38628>.
- Thiel, Stuart E. (1988), “Multidimensional Auctions,” *Economics Letters*, 28 (1), 37–40.
- Varian, Hal R. (2007), “Position Auctions,” *International Journal of Industrial Organization*, 25 (6), 1163–1178.
- Zhu, Yi and Kenneth C. Wilbur (2011), “Hybrid Advertising Auctions,” *Marketing Science*, 30 (2), 249–273.