International Journal of Industrial Organization 45 (2016) 38–46

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Position auctions with dynamic resizing

Patrick Hummel

Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA

ARTICLE INFO

Article history:
Received 20 May 2015
Received in revised form 24 November 2015
Accepted 4 December 2015
Available online 2 January 2016

JEL classifications:
D44
D47
D82

Keywords:
Position auctions
Dynamic resizing
Reserve prices
Mechanism design
Online advertising
Vickrey–Clarke–Groves

ABSTRACT

This paper analyzes mechanisms for selling advertising inventory in a position auction in which displaying less than the maximal number of ads means the ads that are shown can be dynamically resized and displayed more prominently. I characterize the optimal mechanism with and without dynamic resizing, and illustrate how the optimal reserve prices in a Vickrey–Clarke–Groves mechanism vary with the amount of dynamic resizing and the number of bidders.

1. Introduction

In standard Internet auctions where multiple positions for advertising opportunities are being auctioned at the same time, it is typical for the size and the prominence of the ads that are displayed to be independent of the number of ads that are shown. This is the case in standard sponsored search auctions where a fixed amount of space is allocated for each ad on the side of the search page regardless of the number of ads that end up being displayed.

But while it is standard for the size and the prominence of the ads that are displayed to be fixed in most standard Internet position auctions, this is not the case for all such auctions. The Google Display Network (GDN) helps independent publishers monetize their websites by finding appropriate advertisers for their websites and then running an auction to select ads that should be displayed next to the content on a publisher's website. Many of these auctions are for multiple advertising positions on the website, and have the feature that displaying fewer ads will enable the ads that are displayed to be dynamically resized and shown more prominently, and thus receive a larger number of clicks than if more ads had been displayed on the site. In such a setting, if the most valuable ads are significantly more valuable than some of the less valuable ads that could be displayed, one may prefer to show a smaller number of the most valuable ads so that these ads will receive more clicks.

While position auctions with dynamic resizing are used to auction off a wide amount of advertising inventory on GDN, to the best of my knowledge there has been no published work that theoretically analyzes the properties of these auctions. This paper fills this gap in the literature by analyzing the properties of mechanisms that could be used by Google to sell content ads on GDN. Much of the paper analyzes the properties of the Vickrey–Clarke–Groves (VCG) mechanism that is currently used on GDN.

I first characterize properties of the optimal reserve prices in a VCG mechanism. In a standard position auction without dynamic resizing, the optimal reserve price is the same as the optimal reserve price in a standard Vickrey auction for a single object. However, under dynamic resizing, the optimal reserve price for the publisher is greater than this.

Next, I address the question of whether this VCG mechanism with reserve prices is optimal. In a standard position auction without dynamic resizing, I show that there is no feasible mechanism that will lead to greater revenue for the publisher than running a VCG mechanism with the optimally chosen reserve price. But with dynamic resizing of ads, the optimal mechanism will instead be a direct revelation mechanism that maximizes efficiency with respect to the virtual valuations rather than the actual valuations. This optimal mechanism will typically display fewer ads than the VCG mechanism with reserve prices.

Finally, I give some comparative statics results for a special case of the model in which no more than two ads may be displayed. In this setting, I illustrate that the optimal reserve price in the VCG mechanism

E-mail address: phummel@google.com.

http://dx.doi.org/10.1016/j.ijindorg.2015.12.002
0167-7187/

will be increasing in the number of bidders and vary non-monotonically with the amount of dynamic resizing that takes place.

Position auctions with dynamic resizing are a type of auction with variable supply. Auctions with variable supply have been analyzed in the literature (*e.g.* Back and Zender, 2001; Damianov and Becker, 2010; Damianov et al., 2010; Lengwiler, 1999; LiCalzi and Pavan, 2005), but none of these papers can be applied to the Internet position auction setting since these papers focus on models with perfectly divisible goods, and they also explicitly restrict attention to discriminatory-price and uniform-price auctions, neither of which are used in Internet position auctions. Perhaps the most closely related paper of these is Ausubel and Cramton (2004), which considers the possibility of Vickrey auctions with reserve prices in an auction setting with variable supply. However, since this paper again focuses on perfectly divisible goods, it cannot be applied to the Internet position auction setting.

The fact that ads will be less likely to receive clicks if displayed alongside other ads is also a form of an allocative externality. Auctions with allocative externalities have been studied by Katz and Shapiro (1985, 1986) in the context of auctioning licenses to operate a technology or use a patent. Jehiel and Moldovanu (1996) and Jehiel et al. (1996, 1999) have also studied models with externalities in which the losing bidders' payoffs may be affected by the identity of the winning bidder. And Chen and Potipiti (2010) studies a setting with externalities in which a losing bidder's payoff may be decreasing in that bidder's type. These papers all differ from my paper in that they focus on auctions for a single object.

More closely related to the present paper, Jehiel and Moldovanu (2001) and Figueroa and Skreta (2011) have both considered auctions with externalities in which the auctioneer may sell multiple different objects in the same auction. Jehiel and Moldovanu (2001) derive necessary conditions for an efficient and incentive compatible mechanism to exist, and Figueroa and Skreta (2011) illustrate a number of insights about optimal mechanisms, including that revenue-maximizing reserve prices may depend on the bids of other buyers, and revenue-maximizing mechanisms may sell too often. However, these papers do not attempt to derive results that are specific to a position auction with dynamic resizing of ads.

Finally, there has been a variety of work analyzing mechanisms for selling multiple advertising opportunities for various positions at the same time. Some of these papers, such as Aggarwal et al. (2008); Athey and Ellison (2011); Chen and He (2011); Giotis and Karlin (2008); Gomes et al. (2009); Kempe and Mahdian (2008), and Kuminov and Tennenholtz (2009) explicitly model consumer search which may result in externalities arising from ad click-through rates being influenced by which other ads are shown on the page. Fotakis et al. (2011); Ghosh and Mahdian (2008), and Hummel and McAfee (2014) present alternative models of position auctions with externalities in which the click-through rates of an ad may be influenced by which other ads are on the page.[1] However, this literature has not considered scenarios in which displaying less than the maximal number of ads may result in dynamic resizing of ads which causes the ads that are displayed to receive a greater number of clicks. Furthermore, most papers on position auctions (*e.g.* Aggarwal et al., 2006; Börgers et al., 2013; Chen et al., 2009; Edelman et al., 2007; Edelman and Schwarz, 2010; Fukuda et al., 2013; Gomes and Sweeney, 2014; Lahaie, 2006; Varian, 2007, and Yenmez, 2014) do not allow for the possibility that the click-through rates of ads may be influenced by how many other ads appear on the page.

The paper proceeds as follows. Section 2 presents the model. Section 3 characterizes the properties of the optimal reserve prices for the VCG mechanism and the optimal mechanism and compares the properties of VCG with reserve prices to the optimal mechanism. Section 4 presents some comparative statics results for the case where a maximum of two ads may be displayed. Finally, Section 5 presents some simulation results comparing different mechanisms and Section 6 concludes.

[1] Also see Burguet et al. (2015) and White (2013) for work on the interaction between organic search results and advertisements.

2. The model

There are n advertisers in the set $N=\{1,\ldots,n\}$, and each advertiser i has a value v_i for a click on one of the advertiser's advertisements that is an independent draw from the cumulative distribution function $F_i(\cdot)$ with corresponding continuous density $f_i(\cdot)$. Each advertiser i simultaneously submits a bid $b_i \geq 0$ for the right to place an ad on a publisher. After receiving these bids, a system decides the order in which the ads should be displayed on the publisher's site and possibly also decides how many ads should be displayed. Throughout I assume that no more than $s \leq n$ ads can be displayed on the publisher's site.

The number of clicks that an advertiser receives will be affected both by the total number of ads displayed and the order in which these ads are displayed. In particular, if a total of k ads are displayed on the publisher's site, then the ad in the jth-highest position receives a total of $x_{j,k}$ clicks, where $x_{j,k}$ is nonincreasing in j for all k and $x_{j,k}=0$ for all $j>k$. The total payoff to an advertiser i who receives x clicks is then $x(v_i - c_i)$, where c_i represents the cost an advertiser must pay per click.

Throughout the paper I consider situations in which the publisher only shows an advertiser's ad on a page if the advertiser pays a certain minimum reserve price r for each click. In order to achieve this, I focus on a mechanism that I refer to as *VCG with reserve prices*. This mechanism will ensure that all advertisers have an incentive to bid truthfully and that any advertiser who has an ad displayed is required to pay a price of at least r per click.

Informally, this mechanism proceeds by considering all K bidders who submit a bid greater than the reserve price and restricting attention to allocations with no more than K ads. Prices are then set using VCG pricing under the assumption that there is an additional bidder who places a bid equal to the reserve price. In this situation, any advertiser who has an ad shown necessarily pays a price per click that is greater than or equal to the reserve. Similarly, since the VCG mechanism is a truthful mechanism, it follows that the advertisers have an incentive to bid truthfully in this framework.[2]

Formally, let a denote the number of advertisers who submitted a bid per click that is greater than or equal to the reserve price r, and let $K \equiv \min\{a,s\}$ be the smaller of the number of slots or the number of advertisers with a bid greater than or equal to the reserve. For $j \leq K$, define $b_{(j)}$ to be the jth-highest bid submitted by any of the advertisers. And for $j=K+1$, define $b_{(j)}$ to be the larger of the reserve price and the jth-highest bid.

The VCG mechanism with reserve price r proceeds as follows: The mechanism displays a total of k ads, where k is the positive integer in $[1, K]$ that has the highest value of $\sum_{j=1}^{k} x_{j,k} b_{(j)}$, by displaying the ad with the jth-highest bid in the jth-highest of these positions. That is, the mechanism selects the allocation that would maximize efficiency subject to the constraint that we can only show ads from advertisers who bid more than the reserve.

Then if $S_{j,K} \equiv \max_{k \in [1,K]} \sum_{i=1}^{j-1} x_{i,k} b_{(i)} + \sum_{i=j}^{k} x_{i,k} b_{(i+1)}$ and $R_{j,k} \equiv \sum_{i \neq j} x_{i,k} b_{(i)}$, the advertiser in the jth position is charged a total price per click of $\frac{1}{x_{j,k}}(S_{j,K} - R_{j,k})$. Here $S_{j,K}$ is a term representing the total welfare that could be achieved if the bidder with the jth-highest bid were not in the auction when there is an additional bidder who submits a bid of r and no more than K ads can be displayed. $R_{j,k}$ represents the total welfare of all advertisers except for the advertiser with the jth-highest bid. These prices are thus identical to the prices that would be selected under VCG if no more than K ads could be displayed and

[2] Another possible way to implement reserve prices in the VCG mechanism would be to increase bidder payments *ex post* if they are insufficient to meet the reserve. However, such an implementation might not preserve incentive compatibility for advertisers even if there is no dynamic resizing (Even-Dar et al., 2008).

there was an additional bidder who placed a bid equal to the reserve price.

I consider two possible settings. First I consider a standard position auction in which the number of ads that are displayed has no effect on the number of clicks that are received by the ad in the jth position (*i.e.* where $x_{j,k}$ is independent of k for all $k \geq j$). The other setting I consider is the setting on GDN in which there is dynamic resizing of ads. In this setting, if less than the maximal number of ads on the page are displayed, then the particular ads that are displayed on the page are shown more prominently and may receive more clicks as a result (*i.e.* $x_{j,k}$ is decreasing in k for all $k \geq j$).

Throughout the analysis I also sometimes make use of the following common assumptions on the distribution of advertiser values:

(IID) Each advertiser's value is an independent and identically distributed draw from the cumulative distribution function $F(v)$ with corresponding continuous density $f(v)$.
(MHR) Monotone Hazard Rate: Condition (IID) holds and the hazard rate $\frac{f(v)}{1-F(v)}$ is increasing in v.
(IVV) Increasing Virtual Valuations: The virtual valuation $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is increasing in v_i for all i.
(SUP) The distributions $F_i(v)$ all have the same upper bound on their support, $\overline{v} < \infty$.[3]

3. Optimal mechanisms and reserve prices

In this section I derive general results about the optimal reserve prices under the VCG mechanism and I also characterize the optimal mechanisms. Finally I compare the properties of the optimal mechanism with those of the VCG mechanism with reserve prices. First I present a result about the optimal reserve prices:

Theorem 1. *Suppose (MHR) holds. Then we have the following results:*

(a). Without dynamic resizing, the reserve price r that maximizes the publisher's revenue in the VCG mechanism is the unique r in the support of F satisfying $r = \frac{1-F(r)}{f(r)}$.[4]

(b). With dynamic resizing, the reserve price that maximizes the publisher's revenue in the VCG mechanism is greater than the reserve price that maximizes the publisher's revenue without dynamic resizing.

Theorem 1 indicates that the optimal reserve price for the publisher when there is no dynamic resizing is the same as the optimal reserve price in a standard single-object auction, but that this equivalence does not extend to settings with dynamic resizing of ads. To understand the intuition behind this result, note that when there is no dynamic resizing, the optimal reserve price can be found by finding the reserve price where the marginal benefits from setting a slightly higher reserve price equal the marginal costs. The marginal benefits come from the fact that if setting a slightly higher reserve price has no effect on the number of ads that are shown, then the advertisers who have their ads shown will have to pay a slightly higher price per click. The marginal costs come from the fact that setting a slightly higher reserve price may exclude some advertiser from having his ad shown.

When there is dynamic resizing, the optimal reserve price can again be found by finding the reserve price where the marginal benefits from setting a slightly higher reserve price equal the marginal costs. The marginal benefits are the same regardless of whether there is dynamic resizing. However, the marginal costs are not. With dynamic resizing, it is less costly to exclude some advertiser from having his ad shown because if some advertiser is excluded, the other advertisers who do have their ads shown will obtain more clicks and pay more money to the publisher. Since the marginal costs from setting a slightly higher reserve price are not as high with dynamic resizing, it follows that the revenue-maximizing reserve price is larger when there is dynamic resizing.

I now characterize the optimal mechanism and illustrate how this compares to the VCG mechanism with reserve prices:

Theorem 2. *Suppose (IVV) holds. Then the optimal mechanism is a direct revelation mechanism in which the jth position is assigned to the advertiser with the jth-highest value of* $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$*, and the mechanism shows a total of* k^* *ads, where* k^* *is the value of* $k \in [1,s]$ *that maximizes* $\sum_{j=1}^{k} x_{j,k}\left(v_{(j)} - \frac{1-F_{(j)}(v_{(j)})}{f_{(j)}(v_{(j)})}\right)$ *when (j) denotes the advertiser with the jth-highest value of* $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$.

The following is an immediate corollary of this result:

Corollary. *Suppose there is no dynamic resizing and (IID) and (IVV) hold. Then the optimal mechanism for the publisher will be to use VCG with reserve prices.*

The reason for this result is that under the conditions in the corollary, the mechanism in Theorem 2 assigns the jth position to the advertiser with the jth-highest value of $v_i - \frac{1-F(v_i)}{f(v_i)}$ if this value is non-negative, and does not assign the jth position if this value is less than zero. This is precisely the VCG mechanism with the reserve price in Theorem 1(a).

But it is also apparent from Theorem 2 that the VCG mechanism with reserve prices will cease to be the optimal mechanism when either the advertisers' values are drawn from different distributions or when there is dynamic resizing of ads. When the advertisers' values are drawn from different distributions, the publisher may not wish to show the ads from the advertisers with the highest values in the highest positions. And when there is dynamic resizing of ads, the publisher may wish to show a different number of ads than would be shown under the VCG mechanism with the optimal reserve price since the publisher wishes to show a number of ads k that maximizes $\sum_{j=1}^{k} x_{j,k}\left(v_{(j)} - \frac{1-F_{(j)}(v_{(j)})}{f_{(j)}(v_{(j)})}\right)$, a term that gives efficiency with respect to the virtual valuations $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ rather than the actual valuations. Since this cannot be guaranteed through the use of the VCG mechanism with reserve prices, the VCG mechanism ceases to be optimal when there is dynamic resizing of ads.

Since the optimal mechanism may differ from the VCG mechanism with reserve prices, it is natural to ask how the allocation of ads that are shown under these two mechanisms would differ. This is done in the following theorem:

Theorem 3. *Suppose (MHR) holds. Then the optimal mechanism in Theorem 2 will display no more ads and sometimes strictly fewer ads than the VCG mechanism with reserve price r satisfying* $r = \frac{1-F(r)}{f(r)}$.

The reason for this result is that the VCG mechanism maximizes efficiency with respect to the actual valuations while the optimal mechanism maximizes efficiency with respect to the virtual valuations. Since the ratio of a higher-value advertiser's virtual valuation to that of a lower-value advertiser is greater than the ratio of their actual valuations, showing fewer ads is more likely to increase efficiency with respect to the virtual valuations than with respect to the actual valuations. Thus the optimal mechanism will display no more ads and sometimes strictly fewer ads than the VCG mechanism with the Myerson (1981) reserve price.[5]

[3] This assumption is automatically satisfied by bounded distributions under condition (IID) but can also be satisfied by other non-identical bounded distributions.

[4] Part (a) of this result also holds under the weaker assumption that (IID) and (IVV) hold.

[5] However, the optimal mechanism in Theorem 2 may display more ads than the VCG mechanism with the optimal reserve price. If r^* denotes the reserve price satisfying $r = \frac{1-F(r)}{f(r)}$, then we know from Theorem 1(b) that the optimal reserve price under dynamic resizing will be some $r^{**} > r^*$. Thus if all advertisers have values $v < r^{**}$, but one advertiser has value $v > r^*$, then an ad will be displayed under the optimal mechanism but not under the VCG mechanism with the optimal reserve price.

The fact that the optimal mechanism will display no more ads and sometimes strictly fewer ads than the VCG mechanism with the Myerson (1981) reserve price suggests that a publisher could have an incentive to restrict the number of slots available under a VCG auction. Suppose, for instance, there are exactly two advertisers and no reserve price. If only one ad is displayed, the publisher's revenue is $x_{1,1}b_{(2)}$ regardless of whether the publisher runs an auction with one or two slots. If two ads are displayed, then the publisher's expected revenue is $(x_{1,1}-x_{2,2})b_{(2)}+(x_{1,1}-x_{1,2})b_{(1)}$. But two ads are displayed if and only if $x_{1,1}b_{(1)}<x_{1,2}b_{(1)}+x_{2,2}b_{(2)}$, and in this case we have $(x_{1,1}-x_{2,2})b_{(2)}+(x_{1,1}-x_{1,2})b_{(1)}<x_{1,1}b_{(2)}$, the publisher's expected revenue under a single-slot auction. Thus in these cases the publisher's expected revenue is higher under a single-slot auction than under a two-slot auction. Moreover, even introducing a fixed reserve price may not be sufficient to overturn this conclusion.

But while the publisher may have an incentive to restrict the number of advertising slots the publisher makes available at the auction under a standard VCG auction, this incentive will vanish if the publisher is able to make use of position-specific reserve prices. Below I introduce the notion of a VCG auction with position-specific reserve prices and illustrate that if the publisher is able to use position-specific reserve prices, then the publisher will no longer want to restrict the number of slots that are available in the auction.

VCG auctions with position-specific reserve prices can be defined in a similar manner to VCG auctions with reserve prices. Let r_k denote the position-specific reserve price for the kth slot, where $r_k \le r_{k+1}$ for all k. Define K to be the largest integer $k \le s$ for which $b_{(k)} \ge r_k$ (*i.e.* the kth-highest bid is greater than the position-specific reserve price for position k). The VCG auction with position-specific reserve prices will then choose allocations and prices according to the VCG auction with K slots and a reserve price of r_K. This auction is again incentive compatible and ensures that if an advertiser is shown in position k, that advertiser will have to pay a price per click equal to at least r_k.

In this setting, we have the following result:

Proposition 1. *Suppose (SUP) holds and* $\sum_{j=1}^{s} x_{j,s} > \max_{k\in[1,s-1]} \sum_{j=1}^{k} x_{j,k}$.[6] *Then the publisher achieves greater expected revenue by running a VCG auction with s slots and the optimal position-specific reserve prices than by running a VCG auction with s − 1 slots and the optimal position-specific reserve prices.*

The intuition for this result is that if the publisher uses a VCG auction with the position-specific reserve price $r_s=\overline{v}$ for the $\overline{v}$ in condition (SUP), then the auction is equivalent to a VCG auction with $s-1$ slots and position-specific reserve prices. However, the publisher can achieve greater expected revenue by setting $r_s=\overline{v}-\epsilon$ for some small $\epsilon>0$ rather than $r_s=\overline{v}$. Thus the publisher will never want to artificially restrict the number of advertising slots available in the auction under position-specific reserve prices.

4. Comparative statics for two-slot case

We now consider the two-slot case in which $x_{1,1}$ represents the number of clicks an ad would receive if it were the only ad shown, and $x_{1,2}$ and $x_{2,2}$ represent the number of clicks the ads would receive in the top two positions if two ads are shown. From Theorem 1(b), we know the optimal reserve price in position auctions with dynamic resizing is different from that in standard single-slot auctions. Thus standard results about the optimal reserve price will fail to hold in the setting considered here. For instance, in standard second-price auctions, the optimal reserve price is independent of the number of bidders (Krishna, 2010). However, this independence will not hold when there is dynamic resizing. Furthermore, the optimal reserve price may also depend on the number of clicks an ad would receive in the various configurations. These points are illustrated in the following proposition:

Proposition 2. *Suppose there are s = 2 slots and (MHR) holds. Then if* $r(x_{1,1}, n)$ *denotes the reserve price that maximizes the publisher's revenue in the VCG mechanism for given values of* $x_{1,1}$ *and n (holding fixed* $x_{1,2}$ *and* $x_{2,2}$*), we have the following:*

(a). $r(x_{1,1}, n)$ *is increasing in the number of bidders n when there is dynamic resizing.*

(b). $r(x_{1,1}, n)$ *is increasing in* $x_{1,1}$ *for values of* $x_{1,1}$ *close to* $x_{1,2}$ *and decreasing in* $x_{1,1}$ *for values of* $x_{1,1}$ *close to* $x_{1,2}+x_{2,2}$.

The intuition for part (a) is as follows. When there are $s=2$ slots, a tiny increase in the reserve price will only affect the publisher's revenue if there are no more than two bidders with values greater than the reserve, so the publisher can condition on this event in setting the reserve price. Conditional on there being exactly one bidder with value greater than the reserve, the publisher will want to set a reserve that is equal to the optimal reserve price in a single-slot auction. But conditional on there being exactly two bidders with values greater than the reserve, the publisher will want to set a reserve price that is larger than the optimal reserve price in a single-slot auction for the same reasoning as in Theorem 1(b).

When there are a larger number of bidders, it is relatively more likely that there will be exactly two bidders with values greater than the reserve price than it is that there will be exactly one such bidder. Thus when there are a larger number of bidders, the publisher should weigh the possibility that there will be exactly two bidders with value greater than the reserve price relatively more heavily in setting the reserve price. Since the publisher wants to use a higher reserve price conditional on there being exactly two bidders with values greater than the reserve price, it then follows that the optimal reserve price is increasing in the number of bidders.

To understand part (b), note that when $x_{1,1}=x_{1,2}$, no dynamic resizing takes place, and the optimal reserve price will be the reserve price in Theorem 1(a), but when $x_{1,1}>x_{1,2}$, then we know from Theorem 1(b) that the optimal reserve price will be larger than this, so the optimal reserve price is increasing in $x_{1,1}$ for values of $x_{1,1}$ close to $x_{1,2}$. Similarly, when $x_{1,1}=x_{1,2}+x_{2,2}$, the mechanism will always display exactly one ad, and the optimal reserve price will be that in Theorem 1(a), but when $x_{1,1}<x_{1,2}+x_{2,2}$, then we know from Theorem 1(b) that the optimal reserve price will be larger than this, so the optimal reserve price is decreasing in $x_{1,1}$ for values of $x_{1,1}$ close to $x_{1,2}+x_{2,2}$.

5. Simulations

I now conduct some numerical simulations to address the question of how different the various mechanisms considered in this paper are. I focus on the case in which the bidders' values are all independent and identically distributed draws from the uniform distribution on [0,1], a case that was also considered by Ostrovsky and Schwarz (2009) in their empirical analysis of optimal reserve prices for sponsored search auctions at Yahoo!. First I consider the question of how different the optimal reserve prices are in position auctions with dynamic resizing compared to the optimal reserve prices for standard single-slot auctions.

Fig. 1 notes how the optimal reserve prices vary with the extent to which dynamic resizing affects the click-through rates of the ads ($x_{1,1}$) in a position auction with two slots when there are either $n=3$ bidders or $n=6$ bidders. This figure exhibits the properties noted in Proposition 2 in that the optimal reserve price is always larger when there are a larger number of bidders and that the optimal reserve price is initially increasing in the extent to which dynamic resizing affects the click-through rates of the ads and then decreasing in this

[6] This last assumption has the substantive interpretation that there exist advertiser values for which it is more efficient to show s ads than a smaller number of ads.

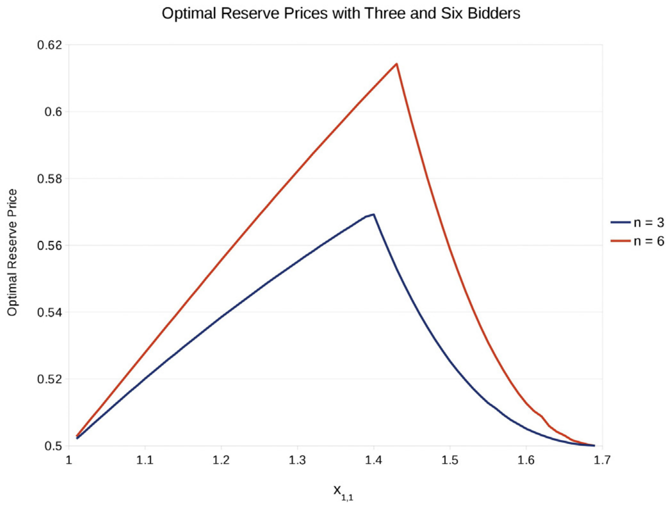

Fig. 1. Optimal reserve prices for various values of $x_{1,1}$ in a two-slot auction when $x_{1,2}=1$ and $x_{2,2}=0.7$ and bidder values are drawn from the uniform distribution on [0,1] for $n=3$ and $n=6$ bidders.

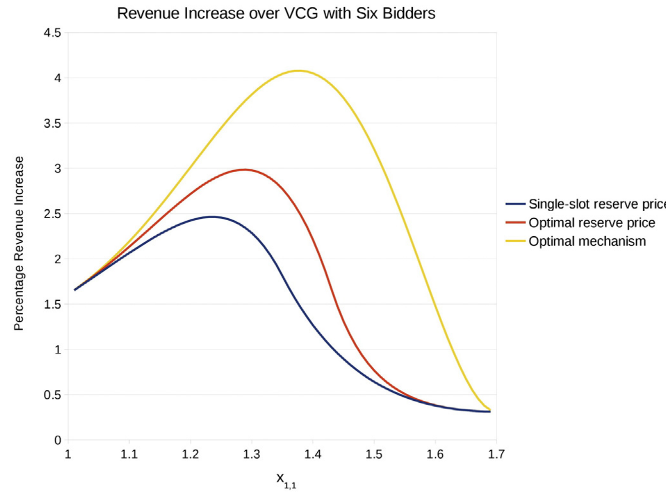

Fig. 3. Revenue increase over the VCG mechanism for various mechanisms and values of $x_{1,1}$ in a two-slot auction when $x_{1,2}=1$, $x_{2,2}=0.7$, there are three bidders, and bidder values are drawn from the uniform distribution.

quantity. We also see from this figure that whether the optimal reserve price differs significantly from the optimal reserve price under single-slot auctions $\left(r=\frac{1}{2}\right)$ depends on the other parameters of the model. When $x_{1,1}$ is close to $x_{1,2}+x_{2,2}$ and dynamic resizing takes place most of the time, then the optimal reserve price is no more than 1–2% larger in a position auction with dynamic resizing than it is in a standard single-slot auction. But for more moderate values of $x_{1,1}$, the optimal reserve prices can be 15–20% larger in a position auction with dynamic resizing than they are in standard single-slot auctions.

Figs. 2 and 3 note how the revenue gains that can be achieved by using either VCG with the best reserve price for single-slot auctions, VCG with the optimal reserve price, or the optimal mechanism vary with the parameters of the model in a position auction with two slots when compared to a standard VCG mechanism without reserve prices. There are some clear patterns that come from these figures. First, the percentage revenue increase from using a mechanism that differs from a standard VCG mechanism is larger when there are a smaller number of bidders, much as in standard single-slot auctions. Second, at least when there are a smaller number of bidders, the gain from using a different mechanism is larger when dynamic resizing plays a smaller role in the sense that $x_{1,1}$ is relatively lower.

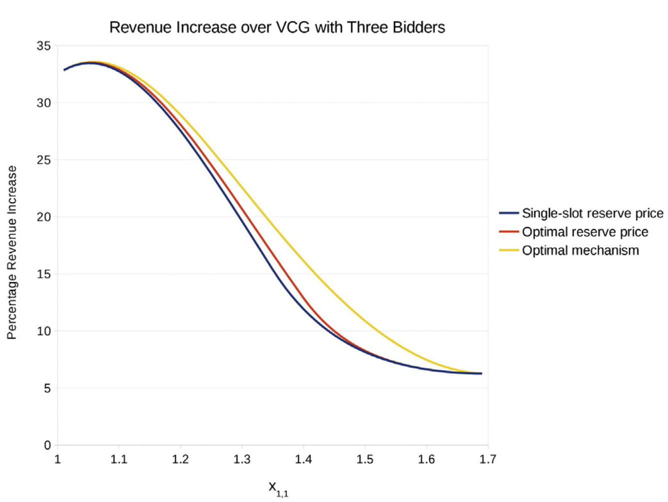

Fig. 2. Revenue increase over the VCG mechanism for various mechanisms and values of $x_{1,1}$ in a two-slot auction when $x_{1,2}=1$, $x_{2,2}=0.7$, there are three bidders, and bidder values are drawn from the uniform distribution.

These figures also contain some interesting insights about the relative performance of different mechanisms. When there are a small number of bidders, simply using the reserve price that is optimal for the single-slot auction typically captures most of the gains over the standard VCG mechanism that could be achieved, although the optimal mechanism still sometimes results in gains that are up to 35% larger than the revenue gains that are achieved by using the VCG mechanism with the best reserve price for single-slot auctions. But when there are a larger number of bidders, the revenue gains that can be achieved by using the optimal mechanism can be up to five times larger than the revenue gains that would be achieved by using the VCG mechanism with the best reserve price for single-slot auctions.

We can also see from these figures that there can be significant gains from using the optimal mechanism rather than simply using the best reserve price for single-slot auctions beyond those that can be achieved by using the optimal reserve price for the VCG mechanism. For all but the smallest values of $x_{1,1}$ plotted in the figures, the additional revenue gains that are achieved by using the optimal mechanism rather than the VCG mechanism with the optimal reserve price are at least as large and sometimes significantly larger than the revenue gains that are achieved by using the optimal reserve price rather than the reserve price that would be best for a single-slot auction.

6. Conclusion

This paper has studied the properties of mechanisms for selling content ads on the Google Display Network where multiple advertising opportunities on a page are auctioned off at the same time and displaying less than the maximal number of ads means the ads that are displayed can be dynamically resized and shown more prominently, thereby inducing these ads to receive more clicks. These position auctions with dynamic resizing have not been previously considered in the academic literature.

In a standard position auction with no dynamic resizing, the optimal mechanism is a VCG mechanism with a reserve price equal to the optimal reserve price for single-slot auctions. With dynamic resizing the publisher will instead want to select a configuration that is efficient with respect to the virtual valuations rather than the actual valuations. The optimal mechanism under position auctions with dynamic resizing will typically result in fewer ads being displayed than would be displayed under the VCG mechanism with reserve prices. Furthermore, if the publisher uses the VCG mechanism with reserve prices, the publisher will want to set a higher reserve price under dynamic resizing than when there is no dynamic resizing, and the optimal reserve price will vary with the number of bidders and the amount of dynamic resizing that takes place.

Acknowledgments

I thank Joshua Dillon, Chris Harris, Tim Lipus, Guy Maor, Preston McAfee, the anonymous referees, and the editor for helpful comments and discussions.

Appendix A. Proofs of Theorems

Lemma 1. *The publisher's expected revenue in a direct revelation mechanism in which a bidder with zero value obtains zero payoff is equal to* $\sum_{i=1}^{n}\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\int_0^\infty p_{i,j,k}(v)\left(v-\frac{1-F_i(v)}{f_i(v)}\right)f_i(v)\,dv$, *where* $p_{i,j,k}(v)$ *denotes the probability that if bidder i has value v, then exactly k ads are displayed and bidder i obtains the jth position.*

Proof. In a direct revelation mechanism, each advertiser follows a symmetric and strictly monotonic bidding strategy $b(v)$ in equilibrium, where $b(v)=v$ for all v. Thus we know from the Integral Form Envelope Theorem in Milgrom (2004) that if $u_i(b,v)$ denotes the expected utility that a bidder i with value v obtains from making a bid of b and $U_i(v)\equiv\sup_{b\in[0,\infty)}u_i(b,v)$, then $U_i(v)=u_i(b(v),v)=u_i(b(0),0)+\int_0^v u_{i,2}(b(y),y)\,dy$.

Now let $t_i(v)$ denote the expected payment that bidder i makes in the mechanism if this bidder has a value v. In this case we know that $u_i(b(v),v)=\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-t_i(v)$. We also know that $u_{i,2}(b(y),y)=\frac{d}{dv}\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)v-t_i(y)$ evaluated at $v=y$ or $\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)$. Combining these results with the results in the previous paragraph shows that $\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-t_i(v)=u_i(b(0),0)+\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)\,dy$. But this means that $t_i(v)=\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-u_i(b(0),0)-\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)\,dy$, which in turn implies that advertiser i's expected payment, $\int_0^\infty t_i(v)f_i(v)dv$, is equal to $\int_0^\infty\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-u_i(b(0),0)-\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)\,dy\right]f_i(v)\,dv$. Thus the total revenue for the publisher, $\sum_{i=1}^{n}t_i(v)$, is equal to $\sum_{i=1}^{n}\left[\int_0^\infty\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-u_i(b(0),0)-\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)\,dy\right]f_i(v)\,dv\right]$. And if a bidder with zero value obtains zero payoff, then $u_i(b(0),0)=0$ for all i, and total revenue for the publisher is $\sum_{i=1}^{n}\left[\int_0^\infty\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\,p_{i,j,k}(y)\,dy\right]f_i(v)\,dv\right]$.

Now we can rewrite the publisher's expected revenue as

$$
\begin{aligned}
&\sum_{i=1}^{n}\left[\int_0^\infty\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(v)v-\int_0^v\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}p_{i,j,k}(y)\,dy\right]f_i(v)\,dv\right]\\
&=\sum_{i=1}^{n}\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_0^\infty p_{i,j,k}(v)vf_i(v)dv-\int_0^\infty\int_0^v p_{i,j,k}(y)f_i(v)dy\ dv\right]\right]\\
&=\sum_{i=1}^{n}\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_0^\infty p_{i,j,k}(v)vf_i(v)dv-\int_0^\infty\int_y^\infty p_{i,j,k}(y)f_i(v)dv\ dy\right]\right]\\
&=\sum_{i=1}^{n}\left[\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_0^\infty p_{i,j,k}(v)vf_i(v)dv-\int_0^\infty p_{i,j,k}(y)(1-F_i(y))\,dy\right]\right]\\
&=\sum_{i=1}^{n}\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\int_0^\infty p_{i,j,k}(v)(vf_i(v)-(1-F_i(v)))\,dv\\
&=\sum_{i=1}^{n}\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\int_0^\infty p_{i,j,k}(v)\left(v-\frac{1-F_i(v)}{f_i(v)}\right)f_i(v)\,dv.
\end{aligned}
$$

The result then follows. □

Corollary 1. *Suppose (IID) holds. Then the publisher's expected revenue in the VCG mechanism with a reserve price of r is* $n\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_r^\infty p_{j,k}(v,r)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)dv\right]$, *where* $p_{j,k}(v,r)$ *denotes the probability that if a bidder has a value v and the reserve price is r, then exactly k ads are displayed and the bidder with value v obtains the jth position.*

Proof. In the special case where $F_i(v)=F(v)$, $f_i(v)=f(v)$, $p_{i,j,k}(v,r)=p_{j,k}(v,r)$ for all i, and $p_{j,k}(v,r)=0$ for $v<r$, the expression for the publisher's expected revenue in Lemma 1 simplifies to $n\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_r^\infty p_{j,k}(v,r)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)\,dv\right]$. □

Proof of Theorem 1. (a). When there is no dynamic resizing, $x_{j,k}$ is independent of k for all $k\geq j$. Thus if x_j denotes the value of $x_{j,k}$ for all $k\geq j$, and $p_j(v)$ denotes the probability that an advertiser with value v has the jth-highest valuation, then we know from Corollary 1 that the publisher's revenue in the dominant strategy equilibrium of the VCG mechanism with a reserve price of r is $n\sum_{j=1}^{s}x_j\left[\int_r^\infty p_j(v)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)dv\right]$. Now the virtual valuation $v-\frac{1-F(v)}{f(v)}$ is increasing in v, $v-\frac{1-F(v)}{f(v)}<0$ when $v=0$, and $v-\frac{1-F(v)}{f(v)}>0$ for sufficiently large values of v in the support of F. Thus by the intermediate value theorem, we know there is some unique value of $v>0$ in the support of F, say v^*, for which $v-\frac{1-F(v)}{f(v)}=0$. For this v^*, it is necessarily the case that $v-\frac{1-F(v)}{f(v)}>0$ for all $v>v^*$ and $v-\frac{1-F(v)}{f(v)}<0$ for all $v<v^*$, so the integral $\int_r^\infty p_j(v)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)dv$ is maximized when $r=v^*$. Thus the reserve price r that maximizes the publisher's revenue when there is no dynamic resizing is the unique r satisfying $r=\frac{1-F(r)}{f(r)}$.

(b). From part (a) we know the reserve price that maximizes the publisher's revenue when there is no dynamic resizing is the unique r satisfying $r=\frac{1-F(r)}{f(r)}$. Thus to prove that the revenue-maximizing reserve price under dynamic resizing is greater than the revenue-maximizing reserve price without dynamic resizing, it suffices to show that the derivative of the publisher's revenue under dynamic resizing with respect to r is strictly positive when evaluated at the r satisfying $r=\frac{1-F(r)}{f(r)}$.

We know from Corollary 1 that the publisher's revenue in the dominant strategy equilibrium of the VCG mechanism with a reserve price of r is $n\sum_{k=1}^{s}\sum_{j=1}^{k}x_{j,k}\left[\int_r^\infty p_{j,k}(v,r)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)dv\right]$. Note that we can rewrite this as $\int_V\sum_{j=1}^{k^*(\vec{v},r)}x_{j,k^*(\vec{v},r)}\left(v_{(j)}-\frac{1-F(v_{(j)})}{f(v_{(j)})}\right)h(\vec{v})d\vec{v}$, where $\vec{v}\equiv(v_1,v_2,\ldots,v_n)$ denotes the advertisers' values for a click, $h(\vec{v})\equiv f(v_1)f(v_2)\ldots f(v_n)$ denotes the density at each possible vector of the advertisers' valuations, $k^*(\vec{v},r)$ denotes the number of ads that will be shown on the publisher's page when the advertisers' values are $\vec{v}$ and there is a reserve price of r, $v_{(j)}$ denotes the jth-highest value of the advertisers, and V denotes the set of all possible advertiser valuations.

Note that small changes in the reserve price r can only affect the value of this expression by changing the value of $k^*(\vec{v},r)$ for some realizations of the advertisers' values, $\vec{v}$. Thus if $V(r,\epsilon)$ denotes the set of advertiser valuations, $\vec{v}$, for which the value of $k^*(\vec{v},r)$ is different when the reserve price is r than when the reserve price is $r+\epsilon$ for some small $\epsilon>0$, then the difference between the publisher's revenue when the reserve price is $r+\epsilon$ and the publisher's revenue when the reserve price is r is $\int_{V(r,\epsilon)}\sum_{j=1}^{k^*(\vec{v},r)}\left(x_{j,k^*(\vec{v},r+\epsilon)}-x_{j,k^*(\vec{v},r)}\right)\left(v_{(j)}-\frac{1-F(v_{(j)})}{f(v_{(j)})}\right)h(\vec{v})d\vec{v}$.

Now suppose that r is the reserve price that satisfies $r=\frac{1-F(r)}{f(r)}$. Note that the only way that the final allocation of advertising opportunities can be different when the reserve price is r than it is when the reserve price is $r+\epsilon$ for some small $\epsilon>0$ is if there is some advertiser with a value per click in $[r,r+\varepsilon)$ who would have his ad shown if the reserve price is r but not if the reserve price is $r+\epsilon$. Thus if r satisfies $r=\frac{1-F(r)}{f(r)}$, then it is necessarily the case that for any vector of advertiser valuations,

$\overrightarrow{v} \in V(r,\epsilon)$, it must be the case that $v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}$ is arbitrarily close to zero when $j = k^*\left(\overrightarrow{v}, r\right)$. From this it follows that when r satisfies $r = \frac{1-F(r)}{f(r)}$, then in the limit as $\epsilon > 0$ becomes arbitrarily close to zero, the difference between the publisher's revenue when the reserve price is $r+\epsilon$ and the publisher's revenue when the reserve price is r is arbitrarily close to $\int_{V(r,\epsilon)} \sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} \left(x_{j,k^*(\overrightarrow{v}, r+\epsilon)} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) h\left(\overrightarrow{v}\right) d\overrightarrow{v}$.

Now there are two ways in which the values of $x_{j,k^*(\overrightarrow{v}, r+\epsilon)}$ may differ from $x_{j,k^*(\overrightarrow{v}, r)}$ for $j \le k^*\left(\overrightarrow{v}, r\right) - 1$. For some values of $\overrightarrow{v} \in V(r,\epsilon)$, it may be the case that $x_{j,k^*(\overrightarrow{v}, r+\epsilon)}$ is greater than $x_{j,k^*(\overrightarrow{v}, r)}$ for all values of $j \le k^*\left(\overrightarrow{v}, r\right) - 1$ (in particular, this will happen whenever $k^*\left(\overrightarrow{v}, r+\epsilon\right) = k^*\left(\overrightarrow{v}, r\right) - 1$ and there is only one less ad that is shown when the reserve price is $r+\epsilon$ than when the reserve price is r). For such values of $\overrightarrow{v}$, it is necessarily the case that $\left(x_{j,k^*(\overrightarrow{v}, r+\epsilon)} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > 0$ for all $j \le k^*\left(\overrightarrow{v}, r\right) - 1$ and $\sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} \left(x_{j,k^*(\overrightarrow{v}, r+\epsilon)} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > 0$ as well.

If $x_{j,k^*(\overrightarrow{v}, r+\epsilon)}$ is not greater than $x_{j,k^*(\overrightarrow{v}, r)}$ for all values of $j \le k^*\left(\overrightarrow{v}, r\right) - 1$, then there is some value of $j \le k^*\left(\overrightarrow{v}, r\right) - 1$ for which $x_{j,k^*(\overrightarrow{v}, r+\epsilon)} \le x_{j,k^*(\overrightarrow{v}, r)}$. Note that this can only take place if $k^*\left(\overrightarrow{v}, r+\epsilon\right) \le k^*\left(\overrightarrow{v}, r\right) - 2$, and there is some $j \le k^*\left(\overrightarrow{v}, r\right) - 1$ for which $x_{j,k^*(\overrightarrow{v}, r+\epsilon)} = 0$. But if the mechanism prefers to show $k^*\left(\overrightarrow{v}, r+\epsilon\right) \le k^*\left(\overrightarrow{v}, r\right) - 2$ ads than to show $k^*\left(\overrightarrow{v}, r\right) - 1$ ads when the reserve price is $r+\epsilon$, then it must be the case that if $k \equiv k^*\left(\overrightarrow{v}, r+\epsilon\right)$ and $K \equiv k^*\left(\overrightarrow{v}, r\right) - 1$, then $\sum_{j=1}^{k} x_{j,k} v_{(j)} \ge \sum_{j=1}^{K} x_{j,K} v_{(j)}$.

From this it follows that $\sum_{j=1}^{k} \left(x_{j,k} - x_{j,K}\right) v_{(j)} \ge \sum_{j=k+1}^{K} x_{j,K} v_{(j)}$. This implies that $\sum_{j=1}^{k} \left(x_{j,k} - x_{j,K}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) \ge \sum_{j=k+1}^{K} x_{j,K} \left(v_{(j)}, -, \frac{1-F(v_{(j)})}{f(v_{(j)})}\right)$ because $v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}$ is a larger fraction of $v_{(j)}$ for larger values of $v_{(j)}$ due to the fact that $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ is smaller for larger values of $v_{(j)}$, and the values of $v_{(j)}$ for $j \le k$ are larger than the values of $v_{(j)}$ for $j \ge k+1$.

From this it follows that $\sum_{j=1}^{k} x_{j,k} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) \ge \sum_{j=1}^{K} x_{j,K} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right)$. But we already know from an earlier part of the proof that $\sum_{j=1}^{K} \left(x_{j,K} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > 0$, which implies that $\sum_{j=1}^{K} x_{j,K} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > \sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} x_{j,k^*(\overrightarrow{v}, r)} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right)$. Thus it is also the case that $\sum_{j=1}^{k} x_{j,k} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > \sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} x_{j,k^*(\overrightarrow{v}, r)} \left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right)$, meaning $\sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} \left(x_{j,k^*(\overrightarrow{v}, r+\epsilon)} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > 0$.

Thus it is necessarily the case that $\sum_{j=1}^{k^*(\overrightarrow{v}, r)-1} \left(x_{j,k^*(\overrightarrow{v}, r+\epsilon)} - x_{j,k^*(\overrightarrow{v}, r)}\right)\left(v_{(j)} - \frac{1-F(v_{(j)})}{f(v_{(j)})}\right) > 0$ for all $\overrightarrow{v} \in V(r,\epsilon)$ regardless of whether $x_{j,k^*(\overrightarrow{v}, r+\epsilon)}$ is greater than $x_{j,k^*(\overrightarrow{v}, r)}$ for all values of $j \le k^*\left(\overrightarrow{v}, r\right) - 1$. From this it follows that if r satisfies $r = \frac{1-F(r)}{f(r)}$, then the publisher obtains strictly greater revenue by setting a reserve price of $r+\epsilon$ for some small $\epsilon > 0$ than by setting a reserve price of r. Thus the reserve price that maximizes the publisher's revenue in the VCG mechanism with dynamic resizing is greater than the reserve price that maximizes the publisher's revenue when dynamic resizing is not present. □

Proof of Theorem 2. Note that we can rewrite the publisher's expected revenue in Lemma 1 as $\int_V \sum_{i=1}^{n} \sum_{k=1}^{s} \sum_{j=1}^{k} x_{j,k} p_{i,j,k}\left(\overrightarrow{v}\right)\left(v_i - \frac{1-F_i(v_i)}{f_i(v_i)}\right) h\left(\overrightarrow{v}\right) d\overrightarrow{v}$, where $\overrightarrow{v} \equiv (v_1, v_2, \ldots, v_n)$ denotes the advertisers' valuations, $h\left(\overrightarrow{v}\right) \equiv f_1(v_1) f_2(v_2) \ldots f_n(v_n)$ denotes the density for each possible realization of the advertisers' valuations, $p_{i,j,k}\left(\overrightarrow{v}\right)$ denotes the probability that advertiser i obtains the jth position and k ads are displayed if the advertisers' valuations are given by $\overrightarrow{v}$, and V denotes the set of all possible advertiser valuations.

It is apparent that this expression is maximized by assigning the jth position to the advertiser with the jth-highest value of $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. It is also apparent that this expression is maximized by always displaying $k^*\left(\overrightarrow{v}\right)$ ads, where $k^*\left(\overrightarrow{v}\right)$ denotes the value of $k \in [1,s]$ that maximizes $\sum_{j=1}^{k} x_{j,k} \left(v_{(j)} - \frac{1-F_j(v_{(j)})}{f_j(v_{(j)})}\right)$ when $v_{(j)}$ denotes the value of the advertiser in position j. Thus the optimal mechanism is a direct revelation mechanism in which the jth position is assigned to the advertiser with the jth-highest value, and the mechanism shows k^* ads, where k^* is the value of $k \in [1,s]$ that maximizes $\sum_{j=1}^{k} x_{j,k} \left(v_{(j)} - \frac{1-F_j(v_{(j)})}{f_j(v_{(j)})}\right)$. □

Proof of Theorem 3. Note that both the optimal mechanism in Theorem 2 as well as the VCG mechanism with a reserve price of r will never display any ads from advertisers with bids $b < r$ because these advertisers fail to meet the reserve price in the VCG mechanism with reserve price r and they also have negative virtual valuations $b - \frac{1-F(b)}{f(b)} < 0$ and therefore will not be shown by the mechanism in Theorem 2. Thus we can exclusively focus on advertisers with bids $b > r$ in assessing whether the mechanism in Theorem 2 or the VCG mechanism with a reserve price of r will display more ads.

Note that if the VCG mechanism with a reserve price of r displays $k < K$ ads, then it must be the case that $\sum_{j=1}^{k} x_{j,k} b_{(j)} \ge \sum_{j=1}^{k'} x_{j,k'} b_{(j)}$ for all $k' \in (k, K]$, meaning we have $\sum_{j=1}^{k} \left(x_{j,k} - x_{j,k'}\right) b_{(j)} \ge \sum_{j=k+1}^{k'} x_{j,k'} b_{(j)}$. This in turn implies that $\sum_{j=1}^{k} \left(x_{j,k} - x_{j,k'}\right)\left(b_{(j)} - \frac{1-F(b_{(j)})}{f(b_{(j)})}\right) \ge \sum_{j=k+1}^{k'} x_{j,k'} \left(b_{(j)} - \frac{1-F(b_{(j)})}{f(b_{(j)})}\right)$ because $b_{(j)} - \frac{1-F(b_{(j)})}{f(b_{(j)})}$ is a larger fraction of $b_{(j)}$ for larger values of $b_{(j)}$ due to the fact that $\frac{1-F(b_{(j)})}{f(b_{(j)})}$ is smaller for larger values of $b_{(j)}$, and the values of $b_{(j)}$ for $j \le k$ are larger than the values of $b_{(j)}$ for $j \ge k+1$. But this means that $\sum_{j=1}^{k} x_{j,k} \left(b_{(j)} - \frac{1-F(b_{(j)})}{f(b_{(j)})}\right) \ge \sum_{j=1}^{k'} x_{j,k'} \left(b_{(j)} - \frac{1-F(b_{(j)})}{f(b_{(j)})}\right)$ for all $k' \in (k, K]$, meaning the optimal mechanism in Theorem 2 would also prefer to display k ads rather than display k' ads for all $k' \in (k, K]$. Thus the optimal mechanism in Theorem 2 will never display more ads than the VCG mechanism with a reserve price of r.

To see that the optimal mechanism in Theorem 2 will sometimes display strictly fewer ads than the VCG mechanism with a reserve price of r, note that if there is exactly one advertiser who submits a bid equal to r, exactly one advertiser who submits a bid $b > r$, and $x_{1,1} b < x_{1,2} b + x_{2,2} r$, then the VCG mechanism with a reserve price of r will display exactly two ads. But the optimal mechanism in Theorem 2 will display exactly one ad because the fact that $r - \frac{1-F(r)}{f(r)} = 0$ implies that $x_{1,1}\left(b - \frac{1-F(b)}{f(b)}\right) > x_{1,2}\left(b - \frac{1-F(b)}{f(b)}\right) + x_{2,2}\left(r - \frac{1-F(r)}{f(r)}\right)$. From this it follows that the optimal mechanism in Theorem 2 will sometimes display strictly fewer ads than the VCG mechanism with a reserve price of r. □

Proof of Proposition 1. Let $r_1, \ldots, r_{s-1}$ denote the optimal position-specific reserve prices for a VCG auction with $s-1$ slots. Note that if the publisher runs a VCG auction with s slots and reserve prices $r_1, \ldots, r_{s-1}, r_s$, where $r_s \equiv \overline{v}$, then this will always result in the same allocations and prices as the VCG auction with $s-1$ slots because no advertiser will ever bid $b \geq r_s$.

Now consider a VCG auction with s slots and reserve prices $r_1, \ldots, r_{s-1}, r_s$, where $r_s \equiv \overline{v}-\epsilon$ for some arbitrarily small $\epsilon > 0$. Setting $r_s = \overline{v}-\epsilon$ rather than $r_s = \overline{v}$ can only have an effect on allocations or prices if at least s advertisers have a value greater than $\overline{v}-\epsilon$. If all s advertisers have values greater than $\overline{v}-\epsilon$, then setting $r_s = \overline{v}-\epsilon$ results in revenue of at least $\sum_{j=1}^{s} x_{j,s}(\overline{v}-\epsilon)$. However, setting $r_s = \overline{v}$ would result in revenue no greater than $\max_{k \in [1,s-1]} \sum_{j=1}^{k} x_{j,k}\overline{v}$. Since $\sum_{j=1}^{s} x_{j,s} > \max_{k \in [1,s-1]} \sum_{j=1}^{k} x_{j,k}$, it follows that for sufficiently small $\epsilon > 0$, setting $r_s = \overline{v}-\epsilon$ results in greater revenue than setting $r_s = \overline{v}$. Thus the publisher achieves greater expected revenue by running a VCG auction with s slots and appropriately chosen position-specific reserve prices than by running a VCG auction with $s-1$ slots and the optimal position-specific reserve prices. □

Proof of Proposition 2. (a). We know from Corollary 1 that the publisher's revenue from the VCG mechanism with a reserve price of r is $n\sum_{k=1}^{s}\sum_{j=1}^{k} x_{j,k}\left[\int_r^{\infty} p_{j,k}(v,r)\left(v-\frac{1-F(v)}{f(v)}\right)f(v)\,dv\right]$ when each bidder's value is an independent and identically distributed draw from the cumulative distribution function $F(\cdot)$ with corresponding pdf $f(\cdot)$. To prove the result, I first calculate the values of $p_{1,1}(v,r)$, $p_{1,2}(v,r)$, and $p_{2,2}(v,r)$ for all values of $v \geq r$.

First note that a bidder with value v will be in the second slot if and only if this bidder has the second-highest value and the bidder with the highest value of $v_{(1)}$ has a value satisfying $x_{1,1}v_{(1)} \leq x_{1,2}v_{(1)} + x_{2,2}v$, which holds if and only if $v_{(1)} \leq \frac{x_{2,2}}{x_{1,1}-x_{1,2}}v$. Thus a bidder with value v will be shown in the second slot if and only if the bidder with the highest value has value in $\left(v, \frac{x_{2,2}}{x_{1,1}-x_{1,2}}v\right]$ and all other bidders have a value lower than v. This happens with probability $(n-1)\left(F\left(\frac{x_{2,2}}{x_{1,1}-x_{1,2}}v\right)-F(v)\right)F(v)^{n-2}$, so $p_{2,2}(v,r) = (n-1)\left(F\left(\frac{x_{2,2}}{x_{1,1}-x_{1,2}}v\right)-F(v)\right)F(v)^{n-2}$.

Now if a bidder with value v has the highest value, then this bidder will have the only ad shown if and only if the second-highest value $v_{(2)}$ satisfies either $v_{(2)} \leq r$ or $x_{1,2}v + x_{2,2}v_{(2)} \leq x_{1,1}v$, which holds if and only if either $v_{(2)} \leq r$ or $v_{(2)} \leq \frac{x_{1,1}-x_{1,2}}{x_{2,2}}v$. This happens with probability $F\left(\max\left\{r, \frac{x_{1,1}-x_{1,2}}{x_{2,2}}v\right\}\right)^{n-1}$. And the probability that a bidder with value v has the highest value in the first place is just $F(v)^{n-1}$. By combining these facts, it follows that $p_{1,2}(v,r) = F(v)^{n-1} - F\left(\max\left\{r, \frac{x_{1,1}-x_{1,2}}{x_{2,2}}v\right\}\right)^{n-1}$ and $p_{1,1}(v,r) = F\left(\max\left\{r, \frac{x_{1,1}-x_{1,2}}{x_{2,2}}v\right\}\right)^{n-1}$.

Now the derivative of the seller's revenue with respect to r is $n\sum_{k=1}^{s}\sum_{j=1}^{k} x_{j,k}\left[-p_{j,k}(r,r)\left(r-\frac{1-F(r)}{f(r)}\right)f(r) + \int_r^{\infty}\frac{\partial p_{j,k}(v,r)}{\partial r}\left(v-\frac{1-F(v)}{f(v)}\right)f(v)\,dv\right]$, which for the case $s=2$ is equal to $-n\left[x_{1,1}F(r)^{n-1} + x_{2,2}(n-1)\left(F\left(\frac{x_{2,2}}{x_{1,1}-x_{1,2}}r\right)-F(r)\right)F(r)^{n-2}\right]\left(r-\frac{1-F(r)}{f(r)}\right)f(r) + n\int_r^{\frac{x_{2,2}}{x_{1,1}-x_{1,2}}r}(n-1)f(r)F(r)^{n-2}(x_{1,1}-x_{1,2})\left(v-\frac{1-F(v)}{f(v)}\right)f(v)\,dv\Big]$. Thus the derivative of the seller's revenue with respect to r is non-negative if and only if $-\left[\frac{x_{1,1}}{n-1}F(r) + x_{2,2}\left(F\left(\frac{x_{2,2}}{x_{1,1}-x_{1,2}}r\right) - F(r)\right)\right]\left(r-\frac{1-F(r)}{f(r)}\right) + \int_r^{\frac{x_{2,2}}{x_{1,1}-x_{1,2}}r}(x_{1,1}-x_{1,2})\left(v-\frac{1-F(v)}{f(v)}\right)f(v)\,dv\Big] \geq 0$. But this equation is increasing in n for all r satisfying $r-\frac{1-F(r)}{f(r)} > 0$, so from this it follows that if the derivative of the seller's revenue with respect to r is zero for some n, then this derivative will still be positive for larger values of n. Thus the optimal reserve price is increasing in n.

(b). When $x_{1,1} = x_{1,2}$, we know from Theorem 1(a) that $r(x_{1,1},n)$ is the r satisfying $r = \frac{1-F(r)}{f(r)}$, but when $x_{1,1}$ is slightly greater than $x_{1,2}$, we know from Theorem 1(b) that $r(x_{1,1},n)$ will be greater than this. Thus $r(x_{1,1},n)$ will be increasing in $x_{1,1}$ for values of $x_{1,1}$ close to $x_{1,2}$. And when $x_{1,1} = x_{1,2} + x_{2,2}$, then the auction will always show exactly one ad, $r(x_{1,1},n)$ will be the same as the optimal reserve price for a single slot auction, and the optimal reserve price will be the r satisfying $r = \frac{1-F(r)}{f(r)}$. But when $x_{1,1}$ is slightly less than $x_{1,2} + x_{2,2}$, then we know from the result in Theorem 1(b) that $r(x_{1,1},n)$ will be greater than this r. Thus $r(x_{1,1},n)$ will be decreasing in $x_{1,1}$ for values of $x_{1,1}$ close to $x_{1,2} + x_{2,2}$. □

References

Aggarwal, G., Feldman, J., Muthukrishnan, S., 2006. Bidding to the top: VCG and equilibria of position-based auctions. Proceedings of the 4th Workshop on Approximation and Online Algorithms (WAOA) 15–28.

Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M., 2008. Sponsored search auctions with Markovian users. Proceedings of the 4th International Workshop on Internet and Network Economics (WINE) 621–628.

Athey, S., Ellison, G., 2011. Position auctions with consumer search. Q. J. Econ. 126 (3), 1213–1270.

Ausubel, L.M., Cramton, P., 2004. Vickrey auctions with reserve pricing. Economic Theory 23 (3), 493–505.

Back, K., Zender, J.F., 2001. Auctions of divisible goods with endogenous supply. Econ. Lett. 73 (1), 29–34.

Börgers, T., Cox, I., Pesendorfer, M., Petricek, V., 2013. Equilibrium bids in sponsored search auctions: theory and evidence. Am. Econ. J. Microecon. 5 (4), 163–187.

Burguet, R., Caminal, R., Ellman, M., 2015. In Google we trust? Int. J. Ind. Organ. 39, 44–55.

Chen, B., Potipiti, T., 2010. Optimal selling mechanisms with countervailing positive externalities and an application to tradable retaliation in the WTO. J. Math. Econ. 46 (5), 825–843.

Chen, J., Liu, D., Whinston, A.B., 2009. Auctioning keywords in online search. J. Mark. 73 (4), 125–141.

Chen, Y., He, C., 2011. Paid placement: advertising and search on the Internet. Econ. J. 121 (556), 301–328.

Damianov, D.S., Becker, J.G., 2010. Auctions with variable supply: uniform price versus discriminatory. Eur. Econ. Rev. 54 (4), 571–593.

Damianov, D.S., Oechssler, J., Becker, J.G., 2010. Uniform vs. discriminatory auctions with variable supply. Games Econ. Behav. 68 (1), 60–76.

Edelman, B., Schwarz, M., 2010. Optimal auction design and equilibrium selection in sponsored search auctions. Am. Econ. Rev. 100 (2), 597–602.

Edelman, B., Ostrovsky, M., Schwarz, M., 2007. Internet advertising and the generalized second price auction: selling billions of dollars of keywords. Am. Econ. Rev. 97 (1), 242–259.

Even-Dar, E., Feldman, J., Mansour, Y., Muthukrishnan, S., 2008. Position auctions with bidder-specific minimum prices. Proceedings of the 4th International Workshop on Internet and Network Economics (WINE) 577–584.

Figueroa, N., Skreta, V., 2011. Optimal allocation mechanisms with single-dimensional private information. Rev. Econ. Des. 15 (3), 213–243.

Fotakis, D., Krysta, P., Telelis, O., 2011. Externalities among advertisers in sponsored search. Proceedings of the 4th International Symposium on Algorithmic Game Theory (SAGT) 105-116.

Fukuda, E., Kamijo, Y., Takeuchi, A., Masui, M., Funaki, Y., 2013. Theoretical and experimental investigations of the performance of keyword auction mechanisms. RAND J. Econ. 44 (3), 438–461.

Ghosh, A., Mahdian, M., 2008. Externalities in online advertising. Proceedings of the 17th International World Wide Web Conference (WWW) 161-168.

Giotis, I., Karlin, A.R., 2008. On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. Proceedings of the 4th International Workshop on Internet and Network Economics (WINE) 629–638.

Gomes, R., Sweeney, K., 2014. Bayes–Nash equilibria of the generalized second price auction. Games Econ. Behav. 86, 421–437.

Gomes, R., Immorlica, N., Markakis, E., 2009. Externalities in keyword auctions: an empirical and theoretical assessment. Proceedings of the 5th International Workshop on Internet and Network Economics (WINE) 172–183.

Hummel, P., McAfee, R.P., 2014. Position auctions with externalities. Proceedings of the 10th Conference on Web and Internet Economics (WINE) 417–422.

Jehiel, P., Moldovanu, B., 1996. Strategic nonparticipation. RAND J. Econ. 27 (1), 84–98.

Jehiel, P., Moldovanu, B., 2001. Efficient design with interdependent evaluation. Econometrica 69 (5), 1237–1258.

Jehiel, P., Moldovanu, B., Stachetti, E., 1996. How (not) to sell nuclear weapons. Am. Econ. Rev. 86 (4), 814–829.

Jehiel, P., Moldovanu, B., Stachetti, E., 1999. Multidimensional mechanism design for auctions with externalities. J. Econ. Theory 85 (2), 258–293.

Katz, M.L., Shapiro, C., 1985. On the licensing of innovations. RAND J. Econ. 16 (4), 504–520.

Katz, M.L., Shapiro, C., 1986. How to license intangible property. Q. J. Econ. 101 (3), 567–590.

Kempe, D., Mahdian, M., 2008. A cascade model for externalities in sponsored search. Proceedings of the 4th International Workshop on Internet and Network Economics (WINE) 585–596.

Krishna, V., 2010. Auction Theory. Elsevier, London.

Kuminov, D., Tennenholtz, M., 2009. User modeling in position auctions: re-considering the GSP and VCG mechanisms. Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 273-280.

Lahaie, S., 2006. An analysis of alternative slot auction designs for sponsored search. Proceedings of the 7th ACM Conference on Electronic Commerce (EC) 218–227.
Lengwiler, Y., 1999. The multiple unit auction with variable supply. Economic Theory 14 (2), 373–392.
LiCalzi, M., Pavan, A., 2005. Tilting the supply schedule to enhance competition in uniform-price auctions. Eur. Econ. Rev. 49 (1), 227–250.
Milgrom, P.R., 2004. Putting Auction Theory to Work. Cambridge University Press, New York.
Myerson, R.B., 1981. Optimal auction design. Math. Oper. Res. 6 (1), 58–73.
Ostrovsky, M., Schwarz, M., 2009. Reserve Prices in Internet Advertising Auctions: A Field Experiment. Stanford Graduate School of Business Typescript.
Varian, H., 2007. Position auctions. Int. J. Ind. Organ. 25 (6), 1163–1178.
White, A., 2013. Search engines: left side quality versus right side profits. Int. J. Ind. Organ. 31 (6), 690–701.
Yenmez, M.B., 2014. Pricing in position auctions and online advertising. Economic Theory 55 (1), 243–256.