

Ranking and Tradeoffs in Sponsored Search Auctions

BEN ROBERTS, Microsoft Research Cambridge
 DINAN GUNAWARDENA, Microsoft Research Cambridge
 IAN A. KASH, Microsoft Research Cambridge
 PETER KEY, Microsoft Research Cambridge

In a sponsored search auction, decisions about how to rank ads impose tradeoffs between objectives such as revenue and welfare. In this paper, we examine how these tradeoffs should be made. We begin by arguing that the most natural solution concept to evaluate these tradeoffs is the lowest symmetric Nash equilibrium (SNE). As part of this argument, we generalise the well known connection between the lowest SNE and the VCG outcome. We then propose a new ranking algorithm, loosely based on the revenue-optimal auction, that uses a reserve price to order the ads (not just to filter them) and give conditions under which it raises more revenue than simply applying that reserve price. Finally, we conduct extensive simulations examining the tradeoffs enabled by different ranking algorithms and show that our proposed algorithm enables superior operating points by a variety of metrics.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social and Behavioral Sciences - Economics

General Terms: Economics, Theory, Experimentation

Additional Key Words and Phrases: Sponsored search, keyword auctions, generalized second price, reserve price

1. INTRODUCTION

In a sponsored search auction, the mechanism designer is given a search query from a user and bids from advertisers and faces the problem of deciding which ads to show in response and where to place them. His choices govern the balance achieved between the interests of stakeholders: the user wants meaningful content, the publisher wants revenue, and the advertisers want engagement. In current sponsored search auctions run by Google or Bing, ads' placement is determined by a ranking algorithm. This ranking then determines prices through the payment rule of the generalised second price (GSP) auction. Payment is made when a user clicks on an ad (pay-per-click). Hence an ad's ranking affects the user, both through a direct position effect and an indirect externality on non-sponsored or algorithmic content, has a consequent affect on the advertiser through the probability of engaging with or clicking on the ad, and affects the price paid by the advertiser and revenue generated to the publisher through the pricing mechanism. Higher positions typically receive more attention from users and more clicks. Different ranking algorithms enable different tradeoffs to be made among these stakeholders. In this paper, we examine the tradeoffs enabled by different ranking algorithms and use a combination of theoretical analysis and empirical study to gain insights into new ranking methods.

Historically, Yahoo! initially ranked advertisers by bid $\{b_i\}$, then Google adopted a ranking based on value per impression, or expected revenue, namely ranking by $\{b_i w_i\}$, where the *relevance* w_i is an estimate of the likelihood that advertiser i 's ad will be clicked. Since

Permission to make digital or hardcopies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

then, there has been a widespread adoption of this rank-by-expected revenue approach, subject to various enhancements such as setting reserves prices.

The framework for much of our paper is a single instance of an auction where advertisers bid for slots, which differ in their “quality” (such as click-through rate). This is an abstraction of live ad-auctions, where advertisers bid on keywords as part of a campaign, and keywords are matched to queries. We sidestep the matching issues and assume that the bids represent bids for the actual query, which is indeed the case for exact match bids on specific keywords. In the analysis phase, we assume many of the auction parameters are fixed, if unknown: for example quality scores, click-through rates and the number of bidders. Later in the paper we show simulation results using live data, where these assumptions certainly do not hold.

In order to examine the necessary tradeoffs, we need a solution concept that describes the outcome we expect. The standard analysis of GSP auctions [Aggarwal et al. 2006; Edelman et al. 2007; Varian 2007] looks at complete information Nash equilibria, and in particular their refinement to *symmetric* or *locally envy-free* Nash equilibria (SNE). However, the standard approach to analyzing the revenue of auction designs, based on Myerson’s work on optimal auctions [Myerson 1981], examines performance in Bayes-Nash equilibria.

Fortunately, it turns out that we do not need to choose between these two solution concepts. In a striking result, several groups of authors independently showed that, when ads are ranked by their bid multiplied by their click probability, the “lowest” SNE corresponds to the the VCG outcome [Aggarwal et al. 2006; Edelman et al. 2007; Varian 2007]. Aggarwal et al. showed that this continues to hold when these “rank scores” are multiplied by individualised weights, except that rather than VCG the results correspond to the outcome of a truthful mechanism they call the “laddered auction.” As this is a single parameter domain, this mechanism is a special case of Myerson’s general technique for transforming a monotone allocation rule into a truthful mechanism [Myerson 1981].

Our first result further generalises the connection between the Myerson outcome and the lowest SNE of the GSP auction. We show that the it holds for a broad class of ranking algorithms, which is large enough to include almost all of those that have been previously considered in the literature. We discuss how an argument due to Edelman and Schwartz [2010] provides additional justification for our adoption of the lowest SNE as our solution concept; it also enables us to bound the performance of ranking algorithms outside our class, where SNE may not exist. A notable algorithm of this type is the practically important case of simply imposing a minimum bid on the rank-by-expected revenue system.

Our second contribution is a proposal of a new ranking algorithm, inspired by features of the revenue-optimal auction, with provably good properties. Rather than using a reserve price simply as a minimum bid, we incorporate it directly into the ranking algorithm such that ads with bids near the reserve price receive low rank scores relative to ads with high bids but lower click probabilities. That is, we recommend the class of ranking algorithms $\{(b_i - r)w_i\}$, where r is the reserve price. This change from $\{b_i w_i\}$ achieves a similar “squashing” effect to introducing an exponent $\alpha < 1$ and ranking by $\{b_i w_i^\alpha\}$, proposed by Lahaie and Pennock [2007]. Hence both squashing and setting a reserve are achieved through a single parameter, as opposed to the effects being decoupled into a squashing exponent and a reserve.

We prove that, for sufficiently small reserve prices, incorporating the reserve price into the ranking algorithm raises more revenue than simply using that same reserve price purely as a filter (i.e. the additional revenue is due the change in ordering, not merely the introduction of a reserve price). The meaning of sufficiently small depends on the distribution of advertiser valuations, but for a number of natural distributions (e.g. uniform and exponential) it encompasses all choices of reserve price that do not exceed the revenue-optimal reserve. While our ranking is inspired purely by the revenue-optimal auction, this theorem provides

some insight into why it may enable favourable tradeoffs: to raise a given amount of revenue, a lower (and thus less distortionary) reserve price can be used.

Our third contribution is an extensive consideration of the tradeoffs enabled by different ranking algorithms through simulations. As the auctioneer cares about both short-term revenue and the long-run health of the search platform, the relevant operating points likely do not include the welfare-optimal and revenue-optimal designs. This is in contrast to previous work, which tends to focus on performance only at these extreme points. We use two types of simulation: one simulating a reactive environment where users react to changes in auction design by forming revised equilibria; the other using real auction logs which assumes advertisers do not react to changes, but which captures all the vagaries of real auctions. These include budget constraints, a changing set of advertisers, and stochastic quality factors. These simulations show that our proposed ranking algorithm enables favourable tradeoffs using a variety of metrics. Our results also show that our proposed ranking has several nice properties from an optimisation perspective.

2. PRELIMINARIES

We adopt a standard (Bayesian) model of a GSP auction:

- There are n advertisers (bidders) and m slots.
- If bidder i 's ad is displayed in slot k , its *click-through rate* (CTR) is $w_i s_k$. s_k is a slot effect, while w_i is an ad effect and can be interpreted as the *relevance* of bidder i 's ad. The slots are strictly heterogeneous, with effects $s_1 > s_2 > \dots$.
- Advertiser i has value θ_i for a click. Values are i.i.d. with cdf $F(\theta_i)$ and pdf $f(\theta_i)$.
- Advertisers are assigned to slots by a ranking algorithm. This can be represented by a *ranking function* $y(b, w) \geq 0$. The advertisers are sorted by $y(b_i, w_i)$ with the highest score receiving the first slot. Advertisers with $y(b_i, w_i) = 0$ are excluded (e.g. if they are below some reserve). We restrict y to be a monotone function with respect to b , but not necessarily w .
- Advertisers pay the *generalised second price* for their slot. Assuming for simplicity that advertisers are ordered such that advertiser i is in slot i , advertiser i 's payment is the minimum bid needed to keep his slot

$$p_i^y(b_{i+1}, w_{i+1}; w_i) = \inf\{b : y(b, w_i) > y(b_{i+1}, w_{i+1})\}. \quad (2.1)$$

- We ignore the possibility of non-trivial ties (i.e. $y(b_i, w_i) = y(b_j, w_j) > 0$) as they complicate analysis, and it is not clear how ties should be resolved in a GSP auction. Our analysis focuses on performance *in expectation*, and so we justify this oversight by noting that for all our considered ranking functions, non-trivial ties occur with probability zero and have no bearing on any expected quantity.
- In our analysis, it will be helpful to refer to the slot effect s_k assigned to advertiser i as his *allocation* x_i . In some instances it is appropriate to consider this a function of the realisation of advertiser types: $x_i(\boldsymbol{\theta}, \boldsymbol{w})$. If a ranking function $y(b, w)$ is used to assign the slots, then it is appropriate to consider an advertiser's allocation as a function of the bid and relevance vectors, which we write as $x_i^y(\mathbf{b}, \mathbf{w})$.
- We use x (or x^y) to denote an *allocation rule*, which comprises the set of allocation functions $\{x_i\}$ (or $\{x_i^y\}$).

The only new feature of this model is the use of a general ranking algorithm based on y rather than assuming a particular instantiation.

Much of our theoretical work utilises *virtual values*, a common concept in economic theory. Under general conditions, an advertiser's virtual value may depend on both his true value and his relevance. However in the interests of an easier analysis we assume independence between these two variables, which defines an advertiser's virtual value also to be

independent of his relevance:

$$\varphi(\theta_i) = \theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)}. \quad (2.2)$$

We further assume that the virtual value function is differentiable, and that the hazard rate $f(\theta_i)/(1 - F(\theta_i))$ is non-decreasing, conditions which hold for numerous common distributions. Again, these assumptions are to accommodate an easier analysis. The value at which an advertiser's virtual value becomes zero is $\bar{\theta}$. That is, $\varphi(\bar{\theta}) = 0$. Given that θ_i has distributional support at zero, our assumptions imply that $\bar{\theta}$ exists and is unique.

3. PROPOSED RANKING ALGORITHM

In order to compare ranking algorithms, we must make some assumption about bidder behaviour. A useful starting point is to assume a Bayes–Nash equilibrium (BNE) in which each advertiser submits a bid maximising his own benefit in expectation over the others' types and bids, and his own relevance. An advantage of working in the Bayesian setting is that we can use Myerson's [1981] theory to quickly calculate expected revenue R . In any BNE,

$$R = R(x) = \mathbb{E} \left[\sum_{i=1}^n \varphi(\theta_i) w_i x_i(\boldsymbol{\theta}, \mathbf{w}) \right], \quad (3.1)$$

where we write $R(x)$ instead of $R(x; f, \mathbf{w})$ to emphasise the dependence on the allocation rule x .

Hence using (3.1), one can simply characterise the revenue-optimal auction. That is, it ranks advertisers by $\varphi(\theta_i)w_i$, excluding any advertiser with negative virtual value (i.e. the auction has a reserve price of $r = \bar{\theta}$ where $\varphi(\bar{\theta}) = 0$). However, actually implementing this auction is unlikely to be feasible in practice. In particular, this simple form relies on our assumptions that bidders are symmetric, and that relevance and value are independent. Otherwise, virtual values (and hence the ranking and reserve price) depend on the identity and relevance of the bidder, which makes practical auction design difficult. Even if we could implement such a revenue-optimal auction, other considerations such as advertiser and user satisfaction would make doing so undesirable.

Instead, we note two qualitative features of the revenue-optimal auction. First, it uses a reserve price. Second, bidders with values barely above the reserve price are very low in the rankings. This inspires the new ranking algorithm we evaluate in Sections 5 and 6 which ranks ads by $\{(b_i - r)w_i\}$, which is perhaps the simplest ranking with these two features. Note that under our proposal the price paid for slot i is $b_{i+1}(w_{i+1}/w_i) + r(1 - w_{i+1}/w_i)$, which follows from (2.1), assuming advertiser i is allocated slot i .

4. THE LOWEST SYMMETRIC NASH EQUILIBRIUM

We now identify one specific Nash equilibrium that we shall use to compare ranking algorithms, which has appealing features. We also draw the connection to BNE.

Because of the difficulties involved in a full Bayes-Nash analysis for the GSP auction, a commonly used alternative is to assume a *symmetric Nash equilibrium* (SNE), an ex-post equilibrium concept proposed independently by Varian [2007] and Edelman et al. [2007], (who used the term *locally envy-free equilibrium*). A SNE requires the following inequalities to be satisfied:

$$(\theta_i - p_i^y(b_{i+1}, w_{i+1}; w_i))x_i \geq (\theta_i - p_i^y(b_{j+1}, w_{j+1}; w_i))x_j \quad \text{for all } i, j, \quad (4.1)$$

where p_i^y is the GSP payment (2.1). Note that the SNE inequalities (4.1) are stronger than those that define an ex-post Nash equilibrium, which for $j < i$ would replace the subscripts $j + 1$ with j in the right hand side of (4.1). Hence the set of SNE is a subset of the set of ex-post equilibria.

The striking result that both Varian [2007] and Edelman et al. [2007] realised is that under the ranking algorithm $\{b_i w_i\}$ any SNE yields an efficient outcome, and furthermore there exists a SNE — known as the *lowest* or *bidder-optimal* SNE — in which advertisers' positions and payments are identical to those imposed by the VCG mechanism. Aggarwal et al. [2006] showed a more general connection between the ranking algorithm $\{b_i w_i c_i\}$ (where c_i is a positive constant) and the corresponding “laddered auction”, a family of truthful mechanisms. This result is important for a number of reasons:

- It provides a focal outcome from the space of possible SNE.
- It creates a link between SNE behaviour and the Bayesian setting.
- It provides a natural lower bound on revenue, as every other SNE has higher revenue.

We show that this result is much more general. In particular, for any ranking function of the form

$$y(b, w) = (g(w)b - h(w))^+, \quad (4.2)$$

SNE always exist (g and h are arbitrary non-negative functions). Since y does not necessarily rank the best ad highest, the outcome is, in general, no longer efficient. However, it does respect y , in the sense that the ranking in all SNE is the same ranking that would be used if bidders reported their true values. Finally, the lowest SNE still has a very special structure. Recall that $x^y(\mathbf{b}, \mathbf{w}) = \{x_i^y(\mathbf{b}, \mathbf{w})\}$ are the allocations that result from bids \mathbf{b} and ranking function y . By (4.2), $x^y(\mathbf{b}, \mathbf{w})$ is a monotone allocation rule. Therefore there are unique payments that make x^y an ex-post direct revelation mechanism. The lowest SNE implements this mechanism in the exact same way that the standard ranking implements VCG. In particular, since ex-post direct revelation mechanisms are also BNE, this allows us to give a concise characterisation of the revenue in the lowest SNE.

THEOREM 4.1. *Consider a GSP auction subject to a ranking algorithm $y(b, w)$ within the class (4.2).¹ For any realisation $(\boldsymbol{\theta}, \mathbf{w})$, there exists a non-empty set of SNE and all SNE order bidders by $y(\theta_i, w_i)$. Furthermore, the lowest (revenue) SNE, defined by*

$$y(b_i, w_i)x_{i-1} = \sum_{j \geq i} y(\theta_j, w_j)(x_{j-1} - x_j), \quad (4.3)$$

generates expected revenue

$$R(x^y) = \mathbb{E} \left[\sum_{i=1}^n \varphi(\theta_i) w_i x_i^y(\boldsymbol{\theta}, \mathbf{w}) \right]. \quad (4.4)$$

PROOF. From the GSP payment rule (2.1), the price-per-click charged to bidder i is

$$p_i^y(b_{i+1}, w_{i+1}; w_i) = \frac{y(b_{i+1}, w_{i+1}) + h(w_i)}{g(w_i)}.$$

The SNE inequalities (4.1) are then

$$\left(\theta_i - \frac{y(b_{i+1}, w_{i+1}) + h(w_i)}{g(w_i)} \right) x_i \geq \left(\theta_i - \frac{y(b_{j+1}, w_{j+1}) + h(w_i)}{g(w_i)} \right) x_j,$$

which is equivalent to

$$(y(\theta_i, w_i) - y(b_{i+1}, w_{i+1}))x_i \geq (y(\theta_i, w_i) - y(b_{j+1}, w_{j+1}))x_j. \quad (4.5)$$

¹For simplicity, we assume that all bidders are ranked by the same algorithm y . However, our result still holds if each is ranked using an individualised algorithm y_i from the class (4.2). This enables our result to apply to settings where, for example, the mechanism design incorporates other factors into the rank score.

Varian's [2007] analysis can be directly reapplied to this generalisation, leading to the conclusion that there exists a non-empty set of SNE, and furthermore all SNE use the same allocation rule, ordering bidders by $y(\theta_i, w_i)$.

In the lowest SNE (4.3), advertiser i 's payment p_i satisfies

$$\begin{aligned} y(p_i, w_i) = y(b_{i+1}, w_{i+1}) &= \frac{1}{x_i} \sum_{j \geq i+1} y(\theta_j, w_j)(x_{j-1} - x_j) \\ &= y(\theta_i, w_i) - \frac{1}{x_i} \sum_{j \geq i} x_j (y(\theta_j, w_j) - y(\theta_{j+1}, w_{j+1})) \\ &= y(\theta_i, w_i) - \frac{1}{x_i} \int_0^{\theta_i} x_i^y(t, \boldsymbol{\theta}_{-i}, \mathbf{w}) dy(t, w_i). \end{aligned}$$

As $dy(t, w_i) = g(w_i) dt$,

$$p_i = \theta_i - \frac{1}{x_i} \int_0^{\theta_i} x_i^y(t, \boldsymbol{\theta}_{-i}, \mathbf{w}) dt,$$

which precisely describes the payment functions imposed by the ex-post direct revelation mechanism for the allocation rule $x^y(\boldsymbol{\theta}, \mathbf{w})$. Thus, the lowest SNE is also a BNE, and therefore generates expected revenue (4.4). \square

This generalisation of the lowest SNE to the class of rankings (4.2) includes ranking by bid $\{b_i\}$, by expected revenue $\{b_i w_i\}$, and the *squashed* ranking $\{b_i w_i^\alpha\}$ [Lahaie and Pennock 2007], all with a possible reserve score (i.e. a per-impression reserve). It also incorporates our proposed algorithm $\{(b_i - r)w_i\}$ with reserve price r (i.e. a per-click reserve).

Note, however, that the standard ranking algorithm $\{b_i w_i\}$ with reserve price r corresponds to the ranking function $z(b, w) = \mathbb{I}\{b \geq r\} bw$, which is *not* of the required form. This introduces some analytical complexities later when we wish to compare the properties of our proposed algorithm to this algorithm. While Theorem 4.1 guarantees SNE of ranking algorithms in the class (4.2) are well behaved, the same cannot be said of the standard ranking with a reserve price. In fact, we will see that this algorithm can be quite poorly behaved, in a sense that will be made clear later.

We conclude this section with an additional justification of our focus on the lowest equilibrium. This justification has the additional benefit of applying even for rankings outside of the class (4.2), for which SNE may not exist, a feature we exploit in the next section. Edelman and Schwarz [2010] argue that because SNE is a full information solution concept used to model the outcome of a game that is in reality one of incomplete information, one should only consider SNE that are in some sense "feasible" in the Bayesian setting. They defined what they called the *Non-Contradiction Criteria* (NCC), which deems a SNE implausible if it generates greater expected revenue than any BNE of the corresponding repeated game of incomplete information. Rather than characterising the BNE of the repeated game, they use the revenue of the optimal BNE as an upper bound. In their setting, this upper bound on revenue exactly matches the revenue of the lowest SNE, and therefore they argue it is the only reasonable equilibrium.

In our work, we are interested in understanding the behaviour of a large class of ranking algorithms, none of which need be optimal. So in our setting, the revenue of the optimal BNE, while still an upper bound, does not necessarily match the revenue given by the lowest SNE of an arbitrary ranking algorithm of the form (4.2). However, we know from Theorem 4.1 that given a ranking function y , all SNE share the same allocation rule $x^y(\boldsymbol{\theta}, \mathbf{w})$. Therefore, a natural comparison is to BNE that also share the same allocation rule. Following Edelman and Schwartz, rather than characterising such equilibria, we instead derive an upper bound on their revenue. Since we have fixed the allocations, Myerson's theory allows us to trivially derive such an upper bound.

PROPOSITION 4.2. *Given a ranking function y , the optimal BNE that ranks ads by $y(\theta_i, w_i)$ has expected revenue*

$$R(x^y) = \mathbb{E} \left[\sum_{i=1}^n \varphi(\theta_i) w_i x_i^y(\boldsymbol{\theta}, \mathbf{w}) \right].$$

By Theorem 4.1, this upper bound exactly matches the revenue of the lowest SNE, providing additional justification for our decision to use it as a focal outcome of a GSP auction. Further, any method of selecting an SNE given types $(\boldsymbol{\theta}, \mathbf{w})$ implicitly defines such a ranking function y , not necessarily within the class (4.2), so this upper bound remains useful even for ranking algorithms outside this class.

5. REVENUE DOMINANCE

In this section we compare the revenue generated by our proposed ranking algorithm $\{(b_i - r)w_i\}$ with the standard ranking $\{b_i w_i\}$, both employing the same per-click reserve price r . This comparison is of particular interest as the two algorithms exclude the same set of advertisers, thus isolating the effect of incorporating the reserve price into the ranking function. We find that for sufficiently small² reserve prices, our proposed algorithm is guaranteed to generate greater revenue. While in practice the designer may not be solely interested in revenue, this result helps to show how our algorithm may offer favourable tradeoffs between revenue and welfare. That is, for a given target revenue, a designer using our ranking algorithm needs to use a smaller (and thus less distortionary) reserve price than a designer employing the standard ranking.

We take our assumption of equilibrium behaviour in GSP auctions to be SNE whose revenue does not exceed the bound in Proposition 4.2. For our proposed ranking algorithm $\{(b_i - r)w_i\}$ this is equivalent to taking the lowest SNE. However, the standard ranking algorithm $\{b_i w_i\}$ with reserve price r has the corresponding ranking function $z(b, w) = \mathbb{I}\{b \geq r\} b w$, which is not within the class (4.2). With this ranking algorithm, we are not certain whether or not a SNE is guaranteed to exist. In Appendix A, we give an example where we can show that any SNE that does exist cannot always rank ads by $\theta_i w_i$. That is, ads do not necessarily appear in the desired order.

Despite the complexity of behaviour with this ranking algorithm, we present the following theorem which states that, for sufficiently small reserve prices, the lowest SNE of the GSP auction subject to our proposed ranking $\{(b_i - r)w_i\}$ generates greater expected revenue than *any* SNE under the standard ranking $\{b_i w_i\}$ (with the same reserve price r) that respects the revenue upper bound from Proposition 4.2.

THEOREM 5.1. *For $r \in (0, \bar{\theta}]$,³ define $R_1(r)$ and $R_2(r)$ to be the expected revenues from two allocation rules that select outcomes that are SNE and do not exceed the bound from Proposition 4.2 under the ranking algorithms $\{(b_i - r)w_i\}$ and $\{b_i w_i\}$ respectively. If*

$$r \leq \inf_{t \geq r} \left\{ t - \frac{\varphi(t)}{\varphi'(t)} \right\}, \tag{5.1}$$

then $R_1(r) > R_2(r)$.

Informally, condition (5.1) seems to hold for most reasonable distributions and for most $r \in (0, \bar{\theta}]$. More precisely, it will hold for all $r \in (0, \bar{\theta}]$ when $\varphi(\theta_i)$ is weakly convex. It is straightforward to show that sufficient conditions for $\varphi(\theta_i)$ to be weakly convex are that f is log-concave and non-increasing. Log-concavity is a property of many common distributions

²We give a detailed discussion on the meaning of “sufficiently small” later, however most reasonable reserve levels do suffice.

³Recall that $\bar{\theta}$ is the theoretical revenue-optimal reserve price — that is, the value at which an advertiser’s virtual value becomes zero: $\varphi(\bar{\theta}) = 0$.

and is a standard assumption in economic analysis [Bagnoli and Bergstrom 2005]. Requiring f to be non-increasing is somewhat restrictive, but permits, for example, the uniform or exponential distribution. Conversely, If $\varphi(\theta_i)$ is concave then it is likely that condition (5.1) does not hold for some choices of r . For example, consider $\theta_i \sim \text{Beta}(2, 2)$ which has a monotone hazard rate and defines $\varphi(\theta_i)$ to be concave. In this case $\bar{\theta} = 0.4215$, and for all r the RHS of (5.1) is minimised at $t = 1$ to the value $\frac{1}{3}$. Thus, there exists an interval $(\frac{1}{3}, \bar{\theta}]$ in which r does not satisfy condition (5.1).

For choices of r that do not satisfy condition (5.1), it does not follow that our ranking algorithm therefore generates less revenue than the standard. On the contrary, we expect our ranking algorithm to generate more revenue in most cases. To give an intuitive explanation, our proof of Theorem 5.1 involves showing that one can apply a large number of pairwise allocation swaps to transform the allocation rule arising from the standard ranking to that of our proposed ranking, each of which increases revenue. If r is slightly greater than the RHS of (5.1) then a small proportion of swaps will decrease revenue, while most will still cause an increase. In many such cases, the net result will still be a revenue increase. Indeed, we see such behaviour in our simulations given below.

We work up to Theorem 5.1 through a series of lemmas. As previously mentioned, the upper bound from Proposition 4.2 is well defined for arbitrary monotone allocation rules. Let $R(x)$ be the value of this bound for the allocation rule x :

$$R(x) = \mathbb{E} \left[\sum_{i=1}^n \varphi(\theta_i) w_i x_i(\boldsymbol{\theta}, \mathbf{w}) \right]. \quad (5.2)$$

If x is not monotone then Proposition 4.2 no longer holds, however $R(x)$ is still well defined as the same functional form. The only difference in this case is that $R(x)$ does not translate as an achievable revenue.

Our first lemma shows how one can increase the integrand of (5.2) for a given realisation $(\boldsymbol{\theta}, \mathbf{w})$. This is achieved by performing a simple adjustment or *swap* to the allocation rule x .

LEMMA 5.2. *Suppose x is an allocation rule for which there exists a realisation $(\boldsymbol{\theta}, \mathbf{w})$ and specific i, j such that*

$$\begin{aligned} \varphi(\theta_i) w_i &> \varphi(\theta_j) w_j, \\ x_i(\boldsymbol{\theta}, \mathbf{w}) &< x_j(\boldsymbol{\theta}, \mathbf{w}). \end{aligned}$$

Define the adjusted allocation rule \tilde{x} which is identical to x except for the single swap $\tilde{x}_i(\boldsymbol{\theta}, \mathbf{w}) = x_j(\boldsymbol{\theta}, \mathbf{w})$ and vice versa. Then,

$$\sum_{i=1}^n \varphi(\theta_i) w_i \tilde{x}_i(\boldsymbol{\theta}, \mathbf{w}) > \sum_{i=1}^n \varphi(\theta_i) w_i x_i(\boldsymbol{\theta}, \mathbf{w}).$$

PROOF. Direct from the conditions. \square

Our next lemma follows as a corollary to Lemma 5.2, extending it to situations where improvements are possible through a sequence of swaps.

LEMMA 5.3. *Let x^y and x^z be two allocation rules such that the following properties hold for all $\theta_i, w_i, \theta_j, w_j$:*

$$y(\theta_i, w_i) = 0 \quad \Leftrightarrow \quad z(\theta_i, w_i) = 0, \quad (5.3)$$

$$\left\{ y(\theta_i, w_i) > y(\theta_j, w_j) \quad \text{AND} \quad z(\theta_i, w_i) < z(\theta_j, w_j) \right\} \Rightarrow \varphi(\theta_i) w_i > \varphi(\theta_j) w_j. \quad (5.4)$$

Then, $R(x^y) \geq R(x^z)$. Furthermore, if x^y and x^z differ with positive probability then $R(x^y) > R(x^z)$.

PROOF. The intuition behind Lemma 5.3 is clear — if it holds that any time y and z disagree about the ranking of two advertisers then y is ‘correct’, then it should hold that $R(x^y) > R(x^z)$. The proof involves showing that one can perform a sequence of swaps to transform x^z into x^y , where each swap satisfies the conditions of Lemma 5.2 (see Appendix B). \square

Our third lemma is the heart of the proof. It shows that for a sufficiently small reserve price r , the two ranking functions defined by our proposed algorithm and the standard satisfy conditions (5.3) and (5.4) from the previous lemma.

LEMMA 5.4. *Given a reserve price $r \in (0, \bar{\theta}]$, let*

$$y(\theta_i, w_i) = (\theta_i - r)^+ w_i, \tag{5.5}$$

$$z(\theta_i, w_i) = \mathbb{I}\{\theta_i \geq r\} \theta_i w_i. \tag{5.6}$$

If

$$r \leq \inf_{t \geq r} \left\{ t - \frac{\varphi(t)}{\varphi'(t)} \right\}, \tag{5.7}$$

then $R(x^y) > R(x^z)$.

PROOF. The proof of Lemma 5.4 involves showing that conditions (5.3) and (5.4) are satisfied (see Appendix C). \square

The only remaining technical detail is that the ranking function $z(b, w) = \mathbb{I}\{b \geq r\} bw$ is not within the class (4.2). As previously discussed, this means that we need to consider the possibility that there may exist SNE with a different ranking from $z(\theta_i, w_i)$, to which we cannot directly apply the upper bound $R(x^z)$. However, our final lemma shows that any such alternate rankings can only reduce our upper bound on revenue.

LEMMA 5.5. *Let x be an allocation rule that selects a SNE of a GSP auction with the ranking function $z(b, w) = \mathbb{I}\{b \geq r\} bw$. Then $R(x^z) \geq R(x)$.*

PROOF. See Appendix D. \square

The addition of Lemma 5.5 is sufficient to complete our proof of Theorem 5.1. By Theorem 4.1, $R_1(r) = R(x^y)$; by Lemma 5.4, $R(x^y) > R(x^z)$; and by Lemma 5.5, $R(x^z) \geq R_2(r)$. Thus we have $R_1(r) > R_2(r)$.

6. SIMULATIONS

In this section, we use simulations to examine the performance of our ranking algorithm and show that it generally dominates existing ranking algorithms. We examine three metrics: revenue, welfare, and click yield. Revenue is what the auctioneer cares about (at least in the short term). Welfare is the total value created for advertisers ($\sum \theta_i w_i x_i$), and the auctioneer also cares about this for the long-term health of the platform. Similarly, click yield (i.e. the total number of clicks $\sum w_i x_i$) can be thought of as a proxy for the value created for the users who are clicking on (presumably) useful ads.

There is one technical detail relevant to Figs. 1-6. As previously discussed, the standard ranking algorithm coupled with a reserve price r corresponds to the ranking function $z(b, w) = \mathbb{I}\{b \geq r\} bw$, which is not within the class (4.2). As a consequence, any existing SNE may not be well-behaved. Instead of trying to characterise such equilibria, we use the relevant statistics of the optimal BNE which ranks ads by $z(\theta_i, w_i)$. By Lemma 5.5, the BNE revenue $R(x^z)$ is an upper bound for the corresponding SNE revenue. Thus, the (red) curves may display overestimates of the true revenues.

The figures’ legends in this section use a simple notational shorthand to signify the type of reserve — that is, r denotes a reserve price and ρ a reserve score (i.e. a per-impression reserve).




Fig. 1. Revenue comparison between our proposed algorithm and the standard ranking in a simple setting.




Fig. 2. Feasible welfare-revenue operating points in a simple setting.

We begin with a simple (albeit unrealistic) example, which satisfies the distributional assumptions made in Sect. 2. There are eight advertisers bidding for three slots. Advertisers have i.i.d. types (θ_i, w_i) where θ_i and w_i are independent and both uniformly distributed on $[0, 1]$. Figure 1 illustrates Theorem 5.1 in this setting: for all $r \leq 0.5 (= \bar{\theta})$, our proposed algorithm of incorporating the reserve price into the ranking function raises more revenue than the standard ranking. In this simple setting, we can actually achieve the optimal revenue at $r = 0.5$.

However, Theorem 5.1 does not tell us what the cost of this added revenue is in terms of welfare. Figure 2 shows us that this revenue is essentially free: for any welfare we desire, we can achieve more revenue with our ranking algorithm. Note that this does not mean that with the same reserve price our algorithm is more efficient. Instead, if separate reserve prices are chosen such that both algorithms have the same welfare, our algorithm has higher revenue.

In Fig. 2, we also compare performance against Lahaie and Pennock’s [2007] squashed ranking algorithm with reserve score ϱ (i.e. $y(b, w) = (bw^\alpha - \varrho)^+$). Since there are two parameters, the operating points form the entire shaded region. Our algorithm leads to a set of operating points that dominates this algorithm as well. Two special cases of this ranking algorithm are the squashing ranking algorithm with no reserve ($\varrho = 0$) and the standard ranking algorithm with a reserve score ($\alpha = 1$). The latter is particularly interesting to compare to the standard ranking algorithm with a reserve price. We observe that for identical pre-reserve rankings, the addition of a reserve price dominates the alternative option of a reserve score. Despite the fact that the plotted revenues of the standard ranking with a reserve price may be overestimates, this still suggests that it is generally better to use reserve prices than reserve scores. Figure 3 shows that these results do not change if we examine click yield rather than welfare.

Lahaie and Pennock [2007] examined the performance of the squashed ranking in a more realistic setting, which they selected by fitting Yahoo! data from a particular query. This distribution violates several of our assumptions. Bidder valuations have a lognormal distribution, which does not have a monotone hazard rate. Values are also correlated with relevance. Nevertheless, Figs. 4-6 show that the results from our simple setting are essen-




Fig. 3. Feasible click yield-revenue operating points in a simple setting.




Fig. 4. Revenue comparison between our proposed algorithm and the standard ranking in Lahaie and Pennock's setting.

tially unchanged, with our proposed algorithm of incorporating a reserve price into the ranking function offering superior tradeoffs.




Fig. 5. Feasible welfare-revenue operating points in Lahaie and Pennock's setting.




Fig. 6. Feasible click yield-revenue operating points in Lahaie and Pennock's setting.

Finally, all of these results are based on the assumption that bidders are in equilibrium. In reality, if parameters are changed, the algorithm may take some time to reach the new equilibrium, and there is empirical evidence that some advertisers react quite slowly to changes. Therefore, a natural question is what happens when the ranking algorithm is changed but advertisers do not react? If the short-term effect is revenue-positive or revenue-neutral, it is much easier for the auctioneer to be patient. Furthermore, by not requiring

an equilibrium analysis, we can examine the performance of different ranking algorithms on historical data, which has many realistic features not captured by our simple model (e.g. changing bidders, matching of bids to multiple queries, and stochastic quality scores).

Figure 7 shows the effect on revenue of changing from the standard ranking algorithm to our proposed ranking algorithm while keeping the reserve price fixed on Microsoft historical data for a keyword with over 500 bidders, which we have selected as representative of a “thick” market. The data has been normalised, but the exact values are not relevant for our purposes. In such markets, incorporating the reserve price into the ranking function seems to consistently increase revenue.




Fig. 7. Revenue comparison between our proposed algorithm and the standard ranking in thick market auction replays.




Fig. 8. Revenue comparison between our proposed algorithm and the standard ranking in thin market auction replays.

Figure 8 is a similar plot of a “thin” market with fewer than 10 bidders. Here, at certain values, the standard ranking raises somewhat more revenue. The included histogram of bid frequencies suggests an explanation for this: setting the reserve price at a common bid makes those bidders pay their full value, while the standard ordering ranks them highly to extract as much revenue as possible. In practice, such reserve prices are unlikely to be chosen, as setting a reserve price to match common bids would essentially make that auction first price, as well as being very sensitive to small changes in bid. At more reasonable choices of reserve price, our proposed algorithm of incorporating it into the ranking function yields greater revenue.

Figures 7 and 8 also demonstrate several advantages of our ranking algorithm from an optimisation perspective. First, the blue lines are “smoother”, which creates a somewhat easier problem. Second, the fact that bidders near the reserve price have low rank scores means that the revenue from an advertiser begins to decrease before the reserve price is actually raised past his bid. This reduces the tendency of optimisation to overfit and choose a reserve price directly below an advertiser’s bid.

To examine the tradeoff between revenue and click yield, we tested a subset of the ranking algorithms using global parameter settings on a sample of a week’s worth of Microsoft data across all queries. As there is a minimum bid of 5 cents in the actual system, an implicit reserve price of 5 cents is applied to all ranking algorithms in addition to any other parameters. Figure 9 shows that incorporating the reserve price into the ranking function




Fig. 9. Feasible click yield-revenue operating points on Microsoft data.

results in a better tradeoff than using the reserve price solely as a minimum bid. As in the single query case, the standard ranking experiences bigger peaks and drops as the reserve price approaches common bids (in this figure increasing the reserve corresponds to moving right to left). Interestingly, a reserve score does not generate a useful tradeoff (at least when set globally) as increasing it reduces both revenue and clicks. Any gains from raising the price the last ad shown pays are more than offset by the lower number of clicks. We cannot investigate the tradeoff between revenue and welfare as we do not know the true values of advertisers.

7. RELATED WORK

A number of papers have examined auctions other than GSP designed for sponsored search auctions. Abrams and Ghosh [2007] considered the design of auctions that guarantee revenue to be at least some fixed proportion of the optimal revenue that can be generated by an omniscient designer, something that is not achieved by the GSP auction. Yenmez [2010] studied equilibrium properties of general auction designs, including looking at conditions under which the VCG outcome is implementable in equilibrium.

GSP auction literature is generally divided between Bayesian and non-Bayesian analysis. In the former, Gomes and Sweeney [2009] pursued existence and uniqueness results regarding BNE of the GSP auction. Pin and Key [2011] and Athey and Nekipelov [2012] proposed separate models from which advertisers' values can be inferred from bid data. Pin and Key derive an advertiser best-response in the face of uncertainty in a repeated auction setting, and make a connection to certain scenarios when the BNE of [Gomes and Sweeney 2009] exist; whereas Athey and Nekipelov start directly from a BNE setting. Both models can be interpreted in a Bayesian setting, but differ in the precise information available to advertisers. Athey and Nekipelov do consider the the tradeoffs from introducing various changing to the ranking algorithm, but only consider one particular choice of parameters rather than the full range of operating points. Athey and Ellison [2011] examined tradeoffs in the context of a model of consumer search where reserve prices can actually increase welfare. In their model a reserve price changes consumer perceptions of an ad's relevance, which can lead to more overall clicks. Vorobeychik [2009] presented a simulation-based technique to approximate a BNE of the GSP auction, made feasible by restricting the space of possible

bidding strategies. Liu and Chen [2006] and Liu et al. [2010] compare the designs of the welfare-optimal and revenue-optimal auctions in a simplified settings where there are only two types of advertisers, but also provide some insight into settings with more than two types.

Analysis of the GSP auction in the complete information setting was pioneered independently by Edelman et al. [2007], Varian [2007], and Aggarwal et al. [2006]. All three papers note the striking equivalence result between the VCG outcome and one specific ex-post equilibrium of the GSP auction — that is, the ‘lowest SNE’ or ‘bidder-optimal locally envy free equilibrium’. Ashlagi et al. [2009] showed that this result holds more generally. That is, for many auction designs in which advertisers’ payments are calculated as functions of lower-ranked bids, there exists a specific equilibrium with payoffs equivalent to the VCG outcome. Paes Leme and Tardos [2010] note that the set of ex-post equilibria of the GSP auction is larger than the set of SNE, and may contain inefficient equilibria. For both the complete-information and Bayesian settings, they investigated the maximum possible loss of welfare — that is, they found upper bounds on the *price of anarchy*. Their results have since been strengthened in a line of work that includes [Lucier et al. 2012; Caragiannis et al. 2011, 2012].

Edelman and Schwarz [2010] examined the revenue of various SNE and argue that the ‘lowest’ is the only reasonable one based on a criterion that compares SNE revenue to the optimal revenue possible in the Bayesian setting. They use the resulting insight to argue that that reserve prices raise more revenue through their indirect effects (increasing equilibrium bids) than through their direct effects (raising prices paid) and that going from a zero reserve price to the optimal reserve price leads to only a small welfare loss while significantly increasing revenue.

Ostrovsky and Schwarz [2011] presented the results of a field experiment aimed at testing the effects of employing Myerson’s [1981] optimal reserve price in GSP auctions. Working in conjunction with Yahoo!, they used historical bid data for a large number of queries to estimate distributions of advertisers’ values and subsequently optimal reserve prices. After employing the new reserve prices, they observed substantial increases in revenue. However, their reserves were implemented as minimum scores and were not used to change to ordering of ads.

Lahaie and Pennock [2007] suggest a different method of increasing revenue. They found that incorporating a squashing exponent $\alpha < 1$ into the ranking algorithm $\{b_i w_i^\alpha\}$ generally has a positive effect on revenue. They used simulations to show that the optimal value of α is sensitive to the correlation between advertisers’ values and relevances. They observed that introducing a squashing exponent to increase revenue incurs a smaller welfare loss than that of setting a reserve score. Thus, they argue their method to be preferential. Further strengthening this justification, Lahaie and McAfee [2011] showed that when there is uncertainty in the estimates of advertisers’ relevances, introducing a squashing exponent can actually *increase* welfare by reducing the weight placed on these uncertain estimates. While we have not explored this issue for our proposed ranking algorithm, it certainly has the same effect, which makes it plausible that their conclusion would apply to our algorithm as well.

8. CONCLUSION

In this paper, we have examined the tradeoffs between revenue, welfare, and click yield enabled by the choice of algorithm to rank sponsored search ads. In developing our solution concept, we have extended the connection between the lowest SNE and the VCG mechanism [Varian 2007; Edelman et al. 2007; Aggarwal et al. 2006] to a much wider class of ranking algorithms. As this class includes inefficient rankings, our extension instead establishes an equivalence between the lowest SNE and the corresponding Myerson [1981] outcome, of which the SNE/VCG result is a special case.

We then proposed a new class of ranking algorithms that incorporates a reserve price/minimum bid into the ranking algorithm that governs the ordering of advertisers (not just whether they appear at all), a system which shares several qualitative features with the revenue-optimal auction. Our proposed ranking essentially combines two methods of increasing revenue previously discussed in literature — namely, squashing [Lahaie and Pennock 2007] and setting a reserve [Edelman et al. 2007; Ostrovsky and Schwarz 2011]. Where previously these two methods have been implemented using separate parameters, our algorithm achieves the aims of both with a single parameter.

We derived conditions under which, for a fixed reserve price, our proposed ranking algorithm generates greater revenue than the standard ranking. This comparison is particularly informative, as it isolates the effect of incorporating the reserve price into the ranking function. It also provides intuition for why our ranking enables good tradeoffs: to raise a given amount of revenue our ranking algorithm can use a lower (less distortionary) reserve price. Our theorem work relies on various distributional assumptions that may not hold in reality, however we expect the comparison to be very similar in many cases which violate the underlying assumptions.

We finished with extensive simulations of the tradeoffs enabled by different ranking algorithms. We used two types of simulation: one simulating an reactive environment where users react to changes in auction design by forming revised equilibria; the other using real auction logs which assumes advertisers do not have time to react to changes, but which captures all the vagaries of real auctions, where bids change in the light of budget constraints and changing users, and where underlying quality factors are stochastic. These simulations show that our proposed ranking algorithm enables superior tradeoffs using a variety of metrics. In particular, to achieve a fixed revenue increase, our ranking algorithm incurred a smaller loss of welfare and click yield than the alternative rankings. These simulations also showed that our proposed ranking has several nice properties from an optimisation perspective.

REFERENCES

- ABRAMS, Z. AND GHOSH, A. 2007. Auctions with revenue guarantees for sponsored search. In *Proc. 3rd International Conference on Internet and Network Economics*. WINE'07. San Diego, California, 143–154.
- AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *ACM Conference on Electronic Commerce*. 1–7.
- ASHLAGI, I., MONDERER, D., AND TENNENHOLTZ, M. 2009. Mediators in position auctions. *Games and Economic Behavior* 67, 1, 2–21.
- ATHEY, S. AND ELLISON, G. 2011. Position auctions with consumer search. *The Quarterly Journal of Economics* 126, 3, 1213–1270.
- ATHEY, S. AND NEKIPELOV, D. 2012. A structural model of sponsored search advertising auctions. Working paper.
- BAGNOLI, M. AND BERGSTROM, T. 2005. Log-concave probability and its applications. *Economic Theory* 26, 445–469.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., AND KYROPOULOU, M. 2011. On the efficiency of equilibria in generalized second price auctions. In *Proc. 12th ACM conference on Electronic Commerce*. EC'11. San Jose, California, 81–90.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND TARDOS, É. 2012. On the efficiency of equilibria in generalized second price auctions. *CoRR abs/1201.6429*.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the Generalized Second-Price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1, 242–259.
- EDELMAN, B. AND SCHWARZ, M. 2010. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review* 100, 2.

- GOMES, R. AND SWEENEY, K. 2009. Bayes–Nash equilibria of the Generalized Second-Price auction. In *Proc. 10th ACM conference on Electronic Commerce*. EC’09. Stanford, California, 107–108.
- LAHAIE, S. AND MCAFEE, R. P. 2011. Efficient ranking in sponsored search. In *Proc. 7th international conference on Internet and Network Economics*. WINE’11. Singapore, 254–265.
- LAHAIE, S. AND PENNOCK, D. M. 2007. Revenue analysis of a family of ranking rules for keyword auctions. In *Proc. 8th ACM conference on Electronic Commerce*. EC’07. San Diego, California, 50–56.
- LIU, D. AND CHEN, J. 2006. Designing online auctions with past performance information. *Decision Support Systems* 42, 3, 1307–1320.
- LIU, D., CHEN, J., AND WHINSTON, A. B. 2010. Ex ante information and the design of keyword auctions. *Information Systems Research* 21, 1, 133–153.
- LUCIER, B., LEME, R. P., AND TARDOS, É. 2012. On revenue in the generalized second price auction. In *WWW*. 361–370.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- OSTROVSKY, M. AND SCHWARZ, M. 2011. Reserve prices in internet advertising auctions: a field experiment. In *Proc. 12th ACM conference on Electronic Commerce*. EC’11. San Jose, California, 59–60.
- PAES LEME, R. AND TARDOS, É. 2010. Pure and Bayes–Nash Price of Anarchy for Generalized Second Price auction. In *Proc. 51st Annual IEEE Symposium on Foundations of Computer Science*. FOCS’10. Las Vegas, Nevada, 735–744.
- PIN, F. AND KEY, P. 2011. Stochastic variability in sponsored search auctions. In *Proc. 12th ACM conference on Electronic Commerce*. EC’11. San Jose, California, 61–70.
- VARIAN, H. R. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6, 1163–1178.
- VOROBAYCHIK, Y. 2009. Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies. In *Proc. 25th Annual Conference on Uncertainty in Artificial Intelligence*. UAI’09. Corvallis, Oregon, 583–590.
- YENMEZ, M. B. 2010. Pricing in position auctions and online advertising. Working paper.

Appendix

A. NON-EXISTENCE OF AN ORDER-PRESERVING SNE

For the standard ranking algorithm with a reserve price (i.e. $z(b, w) = \mathbb{I}\{b \geq r\}bw$), we show that a SNE cannot always rank ads by $z(\theta_i, w_i)$, in contrast to SNE under ranking algorithms within the class (4.2). The SNE inequalities (4.1) can be written as

$$(\theta_i w_i - \max\{r w_i, b_{i+1} w_{i+1}\}) x_i \geq (\theta_i w_i - \max\{r w_i, b_{j+1} w_{j+1}\}) x_j \quad \text{for all } i, j. \quad (\text{A.1})$$

Consider the following realisation. There are precisely two advertisers who submit qualifying bids ($b_i \geq r$), with bidder 1 being awarded the top slot and bidder 2 the second ($b_1 w_1 \geq b_2 w_2$ and $x_1 > x_2$). Bidder 1 is less relevant ($w_1 < w_2$).

Suppose a SNE always ranks ads by $z(\theta_i, w_i)$, so that $\theta_1 w_1 \geq \theta_2 w_2$. The bids (b_1, b_2) must satisfy the inequalities

$$(\theta_1 w_1 - \max\{r w_1, b_2 w_2\}) x_1 \geq (\theta_1 w_1 - r w_1) x_2, \quad (\text{A.2})$$

$$(\theta_2 w_2 - r w_2) x_2 \geq (\theta_2 w_2 - \max\{r w_2, b_2 w_2\}) x_1. \quad (\text{A.3})$$

It is necessary that $b_2 > r$ in order to satisfy (A.3), and as $w_1 < w_2$, we have $\max\{r w_1, b_2 w_2\} = \max\{r w_2, b_2 w_2\} = b_2 w_2$. Then inequalities (A.2) and (A.3) can be rewritten:

$$b_2 w_2 x_1 \leq \theta_1 w_1 (x_1 - x_2) + r w_1 x_2, \quad (\text{A.4})$$

$$b_2 w_2 x_1 \geq \theta_2 w_2 (x_1 - x_2) + r w_2 x_2. \quad (\text{A.5})$$

We need the RHS of (A.4) to be at least as large as the RHS of (A.5). However, this is not always the case. For example, suppose $(\theta_1, w_1) = (1, 0.7)$, $(\theta_2, w_2) = (0.6, 1)$, $r = 0.5$, and $(x_1, x_2) = (1, 0.5)$. We find the bounds on advertiser 2's bid to be $b_2 \geq 0.55$ and $b_2 \leq 0.525$. Thus, a SNE under the standard ranking algorithm with a reserve price does not necessarily rank ads by $\theta_i w_i$.

B. PROOF OF LEMMA 5.3

Lemma 5.3. Let x^y and x^z be two allocation rules such that the following properties hold for all $\theta_i, w_i, \theta_j, w_j$:

$$y(\theta_i, w_i) = 0 \quad \Leftrightarrow \quad z(\theta_i, w_i) = 0, \quad (\text{B.1})$$

$$\left\{ y(\theta_i, w_i) > y(\theta_j, w_j) \quad \text{AND} \quad z(\theta_i, w_i) < z(\theta_j, w_j) \right\} \Rightarrow \varphi(\theta_i) w_i > \varphi(\theta_j) w_j. \quad (\text{B.2})$$

Then, $R(x^y) \geq R(x^z)$. Furthermore, if x^y and x^z differ with positive probability then $R(x^y) > R(x^z)$.

PROOF. Given a realisation $(\boldsymbol{\theta}, \boldsymbol{w})$, suppose there are k advertisers who receive positive scores. Take the labelling of advertisers:

$$\begin{aligned} y(\theta_1, w_1) &> y(\theta_2, w_2) > \dots > y(\theta_k, w_k) \\ x_1^y(\boldsymbol{\theta}, \boldsymbol{w}) &\geq x_2^y(\boldsymbol{\theta}, \boldsymbol{w}) \geq \dots \geq x_k^y(\boldsymbol{\theta}, \boldsymbol{w}). \end{aligned} \quad (\text{B.3})$$

Recall that in Sect. 2 we make the assumption of strict heterogeneity of slot effects ($s_1 > s_2 > \dots$). We specify the weak ordering in (B.3) because the number of available slots may be less than k . However, if advertiser i receives a positive allocation then the strict inequality $x_i^y > x_{i+1}^y$ holds.

From now on we will use the shorthand notation $y(\theta_1, w_1) = y_1$, $x_1^y(\boldsymbol{\theta}, \boldsymbol{w}) = x_1^y$ etc. Let Γ be the permutation of indices such that

$$\begin{aligned} z_{\Gamma(1)} &> z_{\Gamma(2)} > \dots > z_{\Gamma(k)} \\ x_{\Gamma(1)}^z &\geq x_{\Gamma(2)}^z \geq \dots \geq x_{\Gamma(k)}^z. \end{aligned}$$

That is, if an advertiser has the i 'th highest score w.r.t. z , he has the $\Gamma(i)$ 'th highest score w.r.t. y .

We need to show that Γ can be reordered through a sequence of swaps, each of which either satisfies the conditions of Lemma 5.2, or is a trivial swap. Let \mathcal{S} be the set of inversions

$$\mathcal{S} = \{(\Gamma(i), \Gamma(j)) : i < j \text{ and } \Gamma(i) > \Gamma(j)\}.$$

We use the fact that Γ (and Γ^{-1}) can be decomposed into a product of $|\mathcal{S}|$ adjacent transpositions, where each transposition resolves precisely one of the inversions in \mathcal{S} . Applying such a decomposition to the allocations x^z , each non-trivial swap resolves some inversion $(\Gamma(i), \Gamma(j))$. Note that

- $z_{\Gamma(i)} > z_{\Gamma(j)}$ as $i < j$.
- $y_{\Gamma(i)} < y_{\Gamma(j)}$ as $\Gamma(i) > \Gamma(j)$.
- $\varphi(\theta_{\Gamma(i)})w_{\Gamma(i)} < \varphi(\theta_{\Gamma(j)})w_{\Gamma(j)}$ as (B.2) holds.
- $x_{\Gamma(i)} > x_{\Gamma(j)}$ as the inversion has not been previously resolved, and the swap is non-trivial.

By the repeated application of Lemma 5.2, given an arbitrary realisation $(\boldsymbol{\theta}, \mathbf{w})$ at which the allocation rules x^y and x^z differ,

$$\sum_{i=1}^n \varphi(\theta_i)w_i x_i^y(\boldsymbol{\theta}, \mathbf{w}) > \sum_{i=1}^n \varphi(\theta_i)w_i x_i^z(\boldsymbol{\theta}, \mathbf{w}).$$

This process can be applied to all realisations, showing that $R(x^y) \geq R(x^z)$ pointwise. Furthermore, if x^y and x^z differ with positive probability then $R(x^y) > R(x^z)$. \square

C. PROOF OF LEMMA 5.4

Lemma 5.4. Given a reserve price $r \in (0, \bar{\theta}]$, let

$$y(\theta_i, w_i) = (\theta_i - r)^+ w_i, \tag{C.1}$$

$$z(\theta_i, w_i) = \mathbb{I}\{\theta_i \geq r\} \theta_i w_i. \tag{C.2}$$

If

$$r \leq \inf_{t \geq r} \left\{ t - \frac{\varphi(t)}{\varphi'(t)} \right\}, \tag{C.3}$$

then $R(x^y) > R(x^z)$.

PROOF. Consider two advertisers i and j where $\theta_i > \theta_j \geq r$. Using the shorthand notation $\varphi(\theta_i) = \varphi_i$, we write the ratio of their virtual values as

$$\frac{\varphi_i}{\varphi_j} = \frac{\theta_i - g(\theta_i, \theta_j)}{\theta_j - g(\theta_i, \theta_j)},$$

where

$$g(\theta_i, \theta_j) = \frac{\theta_j \varphi_i - \theta_i \varphi_j}{\varphi_i - \varphi_j}. \tag{C.4}$$

Note that

$$\frac{\partial}{\partial k} \left(\frac{\theta_i - k}{\theta_j - k} \right) = \frac{\theta_i - \theta_j}{(\theta_j - k)^2} > 0.$$

If $r \leq g(\theta_i, \theta_j)$,

$$\frac{\theta_i}{\theta_j} < \frac{\theta_i - r}{\theta_j - r} \leq \frac{\varphi_i}{\varphi_j}.$$

In this case the following properties hold:

$$\left\{ \frac{\theta_i - r}{\theta_j - r} > \frac{w_j}{w_i} \quad \text{AND} \quad \frac{\theta_i}{\theta_j} < \frac{w_j}{w_i} \right\} \Rightarrow \frac{\varphi_i}{\varphi_j} > \frac{w_j}{w_i}, \quad (\text{C.5})$$

$$\frac{\theta_i - r}{\theta_j - r} < \frac{w_j}{w_i} \Rightarrow \frac{\theta_i}{\theta_j} < \frac{w_j}{w_i}. \quad (\text{C.6})$$

We would like to find under what conditions $r \leq g(\theta_i, \theta_j)$ for all $\theta_i > \theta_j \geq r$.

Denote the upper bound of the range of θ_i by T (possibly infinite), and consider the infimum of $g(\theta_i, \theta_j)$. We argue that this must occur either at one of the limit points as $\theta_j \rightarrow \theta_i$ for some $\theta_i \in [r, T]$, or at $(\theta_i, \theta_j) = (T, r)$.

Consider the minimising value of θ_i given a fixed $\theta_j = t$. This is either (a) at $\theta_i = T$, (b) at some $\theta_i \in (t, T)$, or (c) at the limit as $\theta_i \rightarrow t$. In case (b), we must have $\partial g / \partial \theta_i = 0$ as g is continuous. From (C.4) this is equivalent to

$$\varphi'_i = \frac{\varphi_i - \varphi(t)}{\theta_i - t},$$

in which case g can be rewritten as

$$g(\theta_i, t) = \theta_i - \frac{\varphi_i}{\varphi'_i} = \lim_{\theta_j \rightarrow \theta_i} g(\theta_i, \theta_j).$$

Thus in both cases (b) and (c), the infimum of $g(\theta_i, t)$ w.r.t. t is some limit point of $g(\theta_i, \theta_j)$ as $\theta_j \rightarrow \theta_i$.

A similar argument can be made regarding the minimising value of θ_j given a fixed θ_i , leading us to the conclusion that the infimum value of $g(\theta_i, \theta_j)$ must occur either (i) at one of the limit points as $\theta_j \rightarrow \theta_i$ or (ii) at $(\theta_i, \theta_j) = (T, r)$. In case (ii)

$$r \leq \bar{\theta} \Rightarrow \varphi(r) \leq 0 \Rightarrow r \leq \inf \{g(\theta_i, \theta_j)\}.$$

In case (i), we require condition (C.3) to obtain $r \leq \inf \{g(\theta_i, \theta_j)\}$.

Therefore, if (C.3) holds then $r \leq g(\theta_i, \theta_j)$ for all $\theta_i > \theta_j \geq r$. Then properties (C.5) and (C.6) imply (B.2) holds and we can invoke Lemma 5.3 to show $R(x^y) \geq R(x^z)$. Furthermore, as θ_i has continuous support over its range, the allocation rules x^y and x^z must differ with positive probability, implying $R(x^y) > R(x^z)$. \square

D. PROOF OF LEMMA 5.5

Lemma 5.5. Let x be an allocation rule that selects a SNE of a GSP auction with the ranking function $z(b, w) = \mathbb{I}\{b \geq r\} bw$. Then $R(x^z) \geq R(x)$.

PROOF. Given a realisation $(\boldsymbol{\theta}, \mathbf{w})$, suppose the allocation rule x selects the SNE in which advertiser i bids $b_i(\boldsymbol{\theta}, \mathbf{w})$. We make the following intuitive assumptions about advertisers' bidding strategies:

- (1) $\theta_i \geq r \Rightarrow r \leq b_i \leq \theta_i$.
- (2) $x_i = 0 \Rightarrow b_i = \theta_i$.

Assumption 1 makes sense as if $\theta_i \geq r$, then $b_i = r$ is a dominant strategy over $b_i < r$ and $b_i = \theta_i$ is dominant over $b_i > \theta_i$. Assumption 2 is a little less intuitive, but is a common concept in auction theory — that is, any losing bidder submits the maximum bid without exposing himself to the possibility of a loss, which is clearly a (weakly) dominant strategy and further drives competition in the auction. Assumptions 1 and 2 imply

$$z(\theta_i, w_i) = 0 \Leftrightarrow z(b_i, w_i) = 0. \quad (\text{D.1})$$

Suppose there are k qualifying advertisers ($\theta_i, b_i \geq r$). Take the labelling of advertisers such that

$$\begin{aligned} b_1 w_1 &> b_2 w_2 > \dots > b_k w_k \\ x_1 &\geq x_2 \geq \dots \geq x_k. \end{aligned}$$

Consider any realisation at which the allocation rules x and x^z differ. That is, there exists a pair of advertisers $j < i$ such that $b_j w_j > b_i w_i$, $\theta_j w_j < \theta_i w_i$, and $x_j > x_i$. From the SNE inequalities (A.1) we have

$$(\theta_j w_j - \max\{r w_j, b_{j+1} w_{j+1}\}) x_j \geq (\theta_j w_j - \max\{r w_j, b_{i+1} w_{i+1}\}) x_i \quad (\text{D.2})$$

$$(\theta_i w_i - \max\{r w_i, b_{j+1} w_{j+1}\}) x_j \leq (\theta_i w_i - \max\{r w_i, b_{i+1} w_{i+1}\}) x_i. \quad (\text{D.3})$$

Taking (D.3) away from (D.2):

$$\begin{aligned} (\theta_j w_j - \theta_i w_i)(x_j - x_i) &\geq \max\{r w_i, b_{i+1} w_{i+1}\} - \max\{r w_i, b_{j+1} w_{j+1}\} \\ &\quad + \max\{r w_j, b_{j+1} w_{j+1}\} - \max\{r w_j, b_{i+1} w_{i+1}\}. \end{aligned}$$

As $x_j > x_i$ and $\theta_j w_j < \theta_i w_i$, the LHS (and thus the RHS also) is negative. In the interest of brevity, denote the four terms in the RHS by A , B , C , and D respectively. We know $C \geq D$ (as $b_{j+1} w_{j+1} > b_{i+1} w_{i+1}$), thus it is necessary that $A < B$. This implies $B = b_{j+1} w_{j+1}$ as $A \geq r w_i$. This implies $C \geq B$, and thus it is necessary that $A < D$. This implies $D = r w_j$ as $A \geq b_{i+1} w_{i+1}$. Now we have

$$\begin{aligned} \text{RHS} &= \max\{r w_i, b_{i+1} w_{i+1}\} - b_{j+1} w_{j+1} + \max\{r w_j, b_{j+1} w_{j+1}\} - r w_j \\ &= \max\{r w_i, b_{i+1} w_{i+1}\} - \min\{r w_j, b_{j+1} w_{j+1}\}. \end{aligned}$$

For this to be negative, it is necessary that $r w_j > r w_i$ and thus $w_j > w_i$. As $\theta_j w_j < \theta_i w_i$, we need $\theta_j < \theta_i$. As the hazard rate $f(\theta_i)/(1 - F(\theta_i))$ is non-decreasing,

$$\begin{aligned} \frac{1 - F(\theta_i)}{f(\theta_i)} &\leq \frac{1 - F(\theta_j)}{f(\theta_j)} \\ \frac{1 - F(\theta_i)}{f(\theta_i)} w_i &\leq \frac{1 - F(\theta_j)}{f(\theta_j)} w_j. \end{aligned}$$

Taking this inequality away from $\theta_i w_i > \theta_j w_j$, we get $\varphi(\theta_i) w_i > \varphi(\theta_j) w_j$. Thus,

$$\left\{ \theta_i w_i > \theta_j w_j \quad \text{AND} \quad b_i w_i < b_j w_j \right\} \Rightarrow \varphi(\theta_i) w_i > \varphi(\theta_j) w_j. \quad (\text{D.4})$$

Note that (D.4) coupled with (D.1) closely resemble the properties (B.1) and (B.2) required for Lemma 5.3. Indeed, one can follow the same process described in Appendix B and show that $R(x^z) \geq R(x)$. \square