Position Auctions and Non-uniform Conversion Rates

Liad Blumrosen
Microsoft Research
Mountain View, CA 94043
liadbl@microsoft.com

Jason D. Hartline
Electrical Engineering and Computer Science
Northwestern University
Evanston, IL 60208
hartline@northwestern.edu

Shuzhen Nong
Microsoft AdCenter
Redmond, WA 98052
shunong@microsoft.com

ABSTRACT

The generalized second-price auction (GSP) is used predominantly for sponsored search by leading search engines like Google, MSN-Live Search and Yahoo!. Previous results showed, in a model where all clicks on an ad gain the advertiser the same benefit, that GSP maximizes the social welfare in equilibrium. In practice, however, the probability that a click will convert to a sale for the advertiser depends on the position (a.k.a. slot) of the ad on the search results page. We support this observation by empirical results collected in MSN-Live adCenter. We then prove that with non-uniform conversion rates GSP does not admit an optimal equilibrium; none-the-less, we are still able to bound the incurred loss. Finally, we devise an incremental change in the GSP mechanism that achieves socially-optimal results in equilibrium, while maintaining the same interface for advertisers and the same pay-per-click business model.

1. INTRODUCTION

The generalized second price auction (GSP) sells trillions of advertising impressions to millions of advertisers generating gross revenues in the billions annually. Its ability as an auction mechanism to effectively allocate the available supply of advertising impressions to advertisers is of crucial importance – if the allocation produced is only two-thirds the value of the optimal allocation it represents billions of dollars lost!

In this paper we combine an empirical study of the sponsored search auction market place with a theoretical study of the GSP auction. Our study suggests, contrary to the implication of initial analyses of GSP's equilibria in a model intended to represent the auctioning of ad slots for a single search query [8, 6], that GSP is not ideally suited to the slot auction problem. The total value (welfare) of the equilibrium of GSP can be as small as two-thirds of optimal. This non-optimality of GSP comes from a failure of GSP to account for the reality that a click on an ad in a top slot is not worth the same as a click on the same ad in a bottom slot.

Workshop on Ad Auctions '08 Chicago, Illinois USA

A conclusion can be made for pay-per-click online advertising in general (including, for example, banner ads); namely, that the value generated by clicks may vary due to various reasons (e.g., the position on the web page) and that this should taken into account in the design of the advertising mechanism.

GSP's inadequacies are many and, in response, other studies have proposed extending GSP in non-trivial ways to allow advertisers a more expressive bidding language in which to specify their preferences. However, in a market where already advertisers are required to solve daunting problems of specifying bids across a huge and diverse inventory of search impressions, making the bidding language more complex seems unlikely to significantly improve the market's efficiency. We propose a solution that takes a different approach. For many advertisers, the result of a successful advertisement is a *conversion*, e.g., the user both clicks on the ad and then subsequently buys one of the advertiser's products, fills out a form, or otherwise completes some electronically trackable action. All major web search engines allow advertisers to opt-in to a *conversion tracking system* where the search engine tracks and keeps statistics on how well the advertiser's ads are converting. These conversion tracking systems can be employed to automatically learn the relative value of the clicks in each slot position. In our work, we show that GSP can be modified, retaining its original bidding interface, to take into account relative differences in conversion rates across advertisers and slots. The resulting mechanism, in a model that is consistent with the results of our empirical study, has optimal equilibria.

GSP works as follows. GSP assumes that the probability that an ad will be clicked on is a product of a slot specific click-through rate and an advertiser specific click-through rate. It assumes that an advertiser's value for a click is the same for any ad slot. When a search query is entered into a search engine, all advertiser bids that match the query are entered into an auction. An advertiser's bid is interpreted as a bid-per-click. GSP computes the advertiser's bid-per-impression by multiplying the advertiser's bid by their click-through rate. The ads are then ranked by decreasing bid-per-impression and are assigned to the slots in this order (usually on the right hand side of the search results page). If an ad is clicked on, the advertiser is charged the minimum bid that would have been necessary to maintain their position. This is precisely the bid-per-impression of the next advertiser divided by this advertiser's click-through rate. If the advertisers bid their true value-per-click then the ranking rule makes sense: the total value of the ranking is the sum value-per-impression of each ad scaled by the click-through rate of the slot. Under the assumption that lower slots have lower click-through rates, this assignment of ads to slots maximizes the total value.

The prior analyses of GSP assumed a full information setting where each advertiser's actual value-per-click is publicly known. This full information assumption is reasonable for repeated games such as slot auctions. As implicit in the definition of GSP, these studies also assume that the value-per-click of an advertiser is the same for all slots. In this model, Varian [8] and Edelman, Ostrovsky, and Schwarz [6] show that in equilibrium GSP ranks the bidders in order of value-per-impression, even though the advertisers are probably not reporting their true value-per-click as their bid. Thus, in equilibrium GSP's assignment of ads to slots maximizes the total value.

The first part of our study is to empirically refute the assumption that the advertisers' value-per-click is uniform across slots. A large percentage of advertisers in MSN Search have conversion tracking enabled. With conversion tracking data we can calculate an advertiser's relative value-per-click for each slot. We find that these are never uniform. Instead, they often are ascending across slots, meaning that clicks from lower slots are often more valuable than clicks from top slots. For the same advertiser, the ratio in value-per-click between slots, according to our data, can be as much as a factor of three.

Motivated by this observation we revisit the theoretical model of Varian [8] and Edelman et al. [6], but in the more realistic setting where the advertisers' values for clicks are non-uniform across slots. Even in the special case where all advertisers have the same non-uniformity in value for clicks across slots, we show that there are price levels where advertisers prefer bottom slots to top slots. This induces equilibrium where bid-per-impressions are not always in the same order as value-per-impressions. Thus, there are scenarios where the rankings produced by GSP in equilibrium do not maximize the total value. We stress that this impossibility holds even when we assume that the optimal total value is obtained by assigning the bidders to slots in the order defined by their values (that is, the product of the click-through rate and the conversion rate decreases with the slot number). In fact, we give a price-of-stability[1] result that constructs settings where all the equilibria of GSP obtain at most two-thirds of the total value possible. We give the complementary result that the loss in total value cannot be drastically low; there always exist equilibria that capture at least two thirds of the total value. These upper and lower bounds are shown for 2-slot auctions, where it already requires a non-trivial analysis. We conjecture that a similar constant upper and lower bound holds more generally for $k > 2$ slots. Note that this price-of-stability result is in the same spirit of the work of [8] and [6] that proved that there always exists at least one welfare-maximizing equilibrium.

Our analyses assume that, like the click-through rates, the conversion rates are a product of an advertiser specific term and a slot specific term. If this is not the case (for either conversion rates or click-through rates) then no GSP variant can hope to optimally allocate advertisers to slots. GSP's ranking algorithm is inherently greedy; whereas, optimally allocating ads to slots when click-through rates (or conversion rates) do not factor requires a weighted matching algorithm. Fortunately, some of the relevant click-through/conversion rate learning algorithms compute these rates as products.

The final part of our study demonstrates a simple fix to GSP to take into account non-uniform conversion rates across slots and results in optimal equilibria. This fix is as simple as multiplying the price an advertiser pays for a click in GSP by the slot specific conversion rate of the slot that the ad was shown in. This effectively simulates a *pay-per-conversion* model without having to make a significant change away from the predominant pay-per-click business model of the sponsored search industry.

Related Work. Our work starts from the original analyses of GSP under uniform conversion rates by Varian [8] and Edelman et al. [6], and extends their study to the practically relevant case where conversion rates are non-uniform. In this case GSP equilibria are non-optimal, so to quantify the extent to which GSP is non-optimal we adopt the mathematical analysis techniques of the price-of-anarchy literature (see [7] for a survey).

Our work is very similar in spirit to the recent work of Abrams et al. [1] that discusses the "cost of conciseness" in GSP, i.e., the possible effect of using a single bid to represent an advertiser's value per click when in fact the advertisers value per click is non-uniform across slots. [1] considers cases where the valuations are arbitrary complicated and not given formulaicly from known conversion rates. They give a general information theoretic lower bound of $1/k$ on the performance of any k-slot auction that allows only a single bid per advertiser. This result holds even in the special case where advertisers values are two-parameter: an advertiser i has a constant value v_i for the top k_i slots and no value for lower slots. We emphasize that their result is information theoretic and applies to advertiser preferences that are more general than our single-value-per-conversion preferences. Further more, since the conversion rates are known in our setting, there is a single-bid-per-advertiser auction that is optimal. Thus, we arrive at the opposite conclusion: there is no cost of conciseness, instead there is a cost due to failure of GSP to account for non-uniform conversion rates. A final note: the techniques necessary for the information theoretic lower bounds of [1] are completely difference from those necessary for our price-of-anarchy style results.

There have been many papers advocating alternative auction formats for slot auctions. For example, in the two-parameter case considered in [1], Aggarwal et al. [3] give a new mechanism with good equilibrium properties. The line of work in studying non-GSP-based auction formats for the next generation slot auctions is extremely important; however, it is also orthogonal to direction proposed in this paper that provides an easily implementable modification to GSP that can significantly improve the equilibrium performance of the current de facto standard slot auction.

In a very recent paper Milgrom [5] studies general properties of games where (where the set of actions is restricted), and analyses settings where such simplifications eliminate inefficient equilibria.

[1]The *price of stability* measures the ratio of the best equilibrium welfare to the optimal welfare. A related notion, *the price of anarchy*, measures the ratio of the worst equilibrium welfare to the optimal welfare.

2. SPONSORED SEARCH AUCTIONS: GSP AND CONVERSION RATES

Consider a set of n advertisers competing for k slots. Each player gains a value-per-conversion of v_i, i.e., v_i is their value for the transaction that results from their ad; what is exactly a "conversion" varies between advertisers, ranging between filling out an online form and buying an airline ticket for thousands of dollars.

A central ingredient of the sponsored search auction model, and what distinguishes it from standard multi-item auctions, is the special parameters that affect the utilities of the bidders. These parameters describe the behavior of numerous users that search for information on the web. The first important set of parameters are the *click-through rates*, where c_i^j denotes the probability that advertiser i's ad is clicked on when it is shown in the j'th slot. Our work focuses on *conversion rates*, where a_i^j denotes the probability that a click on advertiser i's ad will convert to a transaction or acquisition when the ad is shown in slot j.

2.1 Click-Through Rates and Conversion Rates

We wish to call explicit attention to several common assumptions that concern clicks on ads:

1. The click-through rates, c_i^j, are factorable as the product of an advertiser specific term, c_i, and a slot specific term, c^j. So, $c_i^j = c_i c^j$.

2. The click-through rates are monotone non-increasing in slot number. I.e., top (lower numbered) slots generate more clicks than bottom (higher numbered) slots. So, $c^j \geq c^{j+1}$.

A special case analyzed by prior studies of GSP's equilibrium makes an additional uniform conversion rate assumption: the advertiser's value for a click is the same regardless of the slot it comes from. Under this assumption the advertisers' value-per-click are equal to their value-per-conversion so we slightly abuse notation to let v_i represent both of these quantities for advertiser i. We show empirically that this assumption does not hold. In its place we provide the following more general assumptions:

3. The conversion rates, a_i^j, are factorable as the product of an advertiser specific term, a_i, and a slot specific term, a^j. So, $a_i^j = a_i a^j$.

4. The values *per impression* are monotone non-increasing in slot number. This is implied by the assumptions on factorability when the product of the slot specific click-through rate and the slot specific conversion rate is monotone. I.e., top (lower numbered) slots generate more conversions *per impression* than bottom (higher numbered) slots. So, $c^j a^j \geq c^{j+1} a^{j+1}$.

Notice that if the click-through rates are not factorable (Assumption 1) then GSP is not well defined. If the conversion rates are not factorable (Assumption 3) then no greedy ranking rule (like that of GSP) can give an economically efficient allocation. The assumption of monotone click-through rates (Assumption 2) is a consensus in the industry, and we observe it in our data. If the values per impression are not monotone (Assumption 4) then the ranking order used by GSP is inappropriate. (However, in such a case we could, without loss of generality, rename the slots so that slot j is the one with the jth highest number of conversions per impression.) Therefore, we are somewhat justified in these assumptions.

5. The advertiser specific click-through and conversion rates are uniform. So, $a_i^j = a^j$ and $c_i^j = c^j$.

This assumption is employed in our theoretical bounds only and is without loss of generality. Clearly for lower bounds it cannot weaken our results to add this assumption; more generally, equilibrium problems with general advertiser specific coefficients can be reduced to an analogous problem in the uniform case.

2.2 The Generalized Second Price Auction

We will now formally define GSP:

The Generalized Second Price Auction (GSP). The generalized second price auction, for k slots, is given input bids from advertisers. Let b_i represent the *bid-per-click* of advertiser i.

1. Let $b_i c_i$ be the *bid-per-impression* of advertiser i.

2. Sort the bidders by bid-per-impression:
   - Let i_j be the index of the bidder with the jth highest bid-per-impression.
   - For all j, $b_{i_j} c_{i_j} \geq b_{i_{j+1}} c_{i_{j+1}}$.
   - For all j, let p_j be the minimum bid necessary for advertiser i_j to maintain their position in slot j:
   $$p_j = \frac{b_{i_{j+1}} c_{i_{j+1}}}{c_{i_j}}.$$

3. For $j \leq k$, show advertiser i_j in slot j.

4. For $j \leq k$, if ad i_j is clicked, charge p_j to advertiser i_j.

Note that GSP does not explicitly take into account the conversion rates or slot specific click-through rates. We will make two important assumptions on implementations of GSP:

6. Ties are broken uniformly at random.

7. The bid space is discrete, e.g., restricted to be multiples of some minimal bid increment, ϵ.

The assumption that ties are broken uniformly at random (Assumption 6) is consistent with some implementations of GSP. It is also consistent with other equilibrium analyses of pricing games, e.g., Bertrand games. None-the-less, we do not believe changing this assumption significantly affect our results. The assumption of a discrete bid space (Assumption 7) is necessary for the existence of pure Nash equilibria in our model. It is well justified as every implementation of GSP we are aware of has a minimum bid increment of one penny. We do not require that the agent valuations be discrete.

Game-theoretic Analysis of GSP. A *truthful* auction is one where it is a dominant strategy for all bidders to submit bids equal to their true values. GSP is a generalization of the second price auction which is truthful [9]; however it is *not truthful* nor does it belong to the family of truthful VCG mechanisms. This means that we cannot expect advertiser i to submit a bid-per-click, that is equal to their value-per-click. None-the-less, the main theoretical result of [8] and [6] is that in equilibrium, under the uniform conversion rate assumption, when the advertisers have values-per-click of $v_1, \ldots, v_n$ and values-per-impression of $v_1c_1, \ldots, v_nc_n$, the sorted orders of advertisers by value-per-impression and by *bid*-per-impression are identical. Thus, we can conclude that in equilibrium the advertiser with the jth highest value-per-impression indeed gets the jth highest slot, which has the jth highest click-through rate. Such an equilibrium is optimal.

The equilibrium concept used in [8, 6] and this paper, is Nash equilibrium in the full information game induced on GSP by the click-through rates, conversion rates, and advertiser valuations. This concept is well suited for repeated games where private information will not stay private for long. This is the predominant model in the literature on sponsored search and on price-of-anarchy questions. In this game, the actions (pure strategies) of the bidders are their bids, and their payoffs are determined by their utility from the GSP outcome, $u_i^j(p) = c_i^j(a_i^j v_i - p)$ (where bidder i wins slot j and pays p, and v_i denotes their value per conversion).

DEFINITION 1. *A profile of bids $b_1, ..., b_n$ is a (pure Nash)* equilibrium *of GSP, if no bidder i can gain a better utility by deviating to any other bid. Formally, fixing the bids b_{-i} of the other bidders, and assuming that by bidding b_i bidder i wins slot j and pays p, then for every other bid b_i' for which bidder i is assigned to slot j' and pays p' we have $u_i^j(p) \geq u_i^{j'}(p')$.*[2]

Throughout the paper we will use the term equilibrium to denote a pure Nash equilibrium. We will not consider mixed Nash equilibria in this paper, unless explicitly mentioned.

Social Welfare. Let $i_1, \ldots, i_k$ be the ordering defined by the bids-per-impression in GSP, then the *social welfare* is $\sum_{j=1}^{k} v_{i_j} c_{i_j}^j a_{i_j}^j$. The maximum sum is achieved when the ordering of values-per-impression coincides with the ordering of bids-per-impression. Notice that the social welfare includes both the welfare of the advertisers (which is their total value less their total payments) and the payoff of the search engine (which is the total payments).

Finally, we formally note the result of [8, 6] on the economic efficiency of GSP (assuming uniform conversion rates):

THEOREM 1. *[8, 6] For uniform conversion rates, GSP always admits a welfare maximizing equilibrium.*

3. EMPIRICAL RESULTS

In this section, we present empirical data collected in Microsoft adCenter during August 2007. Some of the advertisers that advertise in Microsoft-Live Search use a conversion-tracking system to monitor conversions that are done electronically. We present data collected from all advertisers that are affiliated to an industry (e.g., travel, finance, insurance, etc.). This study is preliminary and meant only as motivation for our theoretical results in subsequent sections. A more thorough empirical study will be included with the full paper.

[2]We can assume that slot $k+1$ exists and gains the bidders values of zero to handle losing players.

Slot	Adv. 1	Adv. 2	Ind. 1
Main Line 1	100%	100%	100%
Main Line 2	90%	84%	110%
Main Line 3	87%	151%	117%
Side Bar 1	83%	151%	122%
Side Bar 2	70%	77%	112%
Side Bar 3	71%	109%	125%
Side Bar 4	51%	114%	132%
Side Bar 5	90%	195%	152%

Figure 1: Data on the relative conversion rates of two individual advertisers over a period of month, and aggregate data for all the advertisers from the same industry over a period of two months.

Slot	CR	CTR	CR*CTR
Main Line 1	100%	100%	100%
Main Line 2	80%	49%	40%
Main Line 3	87%	33%	28%
Side Bar 1	78%	9%	7.2%
Side Bar 2	82%	6%	5.6%
Side Bar 3	96%	5%	5.0%
Side Bar 4	110%	4%	4.6%
Side Bar 5	106%	4%	4.1%

Figure 2: Aggregate data on conversion rates (CR) and on click-through rates (CTR) in different slots, relative to those parameters in the top slot in the main line. The data was collected in Microsoft adCenter over a period of a month, and consists of an aggregation of many advertisers from different industries.

The Live search engine currently shows up to 8 ads in one page. Three appear above the true ("organic") search results and these slots are referred to as the *main line*. Up to five additional ads appear to the right of the organic search result, and we denote these slots as the *side bar*. We denote the slots in the main line as slots 1-3, and in the side bar as slots 4-8. We emphasize that we only collected data for search queries where all the 8 ads were presented.

Our main goal in this section is to exhibit aggregate data that describes how conversion rates tend to change over slots. This aggregation is not a straightforward task due to several reasons. First, the concept of "conversion rate" does not refer to the same thing for different advertisers and different industries, and we observe from the data that conversion rates vary considerably over different advertisers. Second, even after normalizing, the slot effect on conversion rates varies from advertiser to advertiser and industry to industry (Figure 1). It is not clear that an aggregate slot conversion rate is meaningful. Third, most of the data is sparse, representing queries in the "long-tail" of the distribution. A full description of our empirical methodology is given in Appendix C.

The aggregate results in Figure 2 show that the conversion rates tends to increase with the slot number. The highest

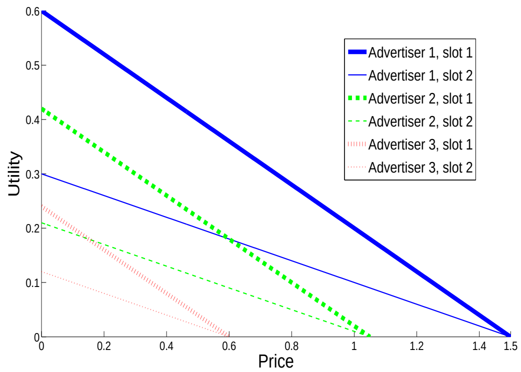

Figure 3: The utilities of the bidders as a function of the price with uniform conversion rates in a 2-slot auction. The figure depicts the preferences of 3 bidders with two curves per bidder, one per each slot. The figure shows that with uniform conversion rates, bidders prefer having the first slot over the second slot at every price level.

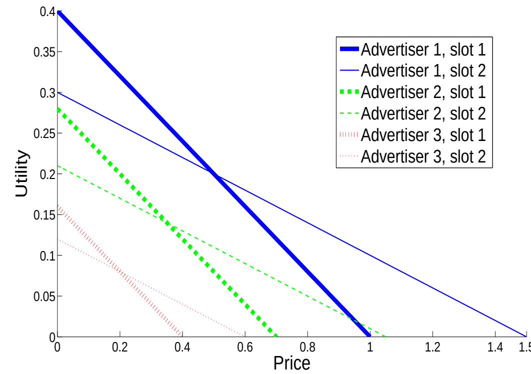

Figure 4: The figure describes the utility for three advertisers in GSP as a function of the price they pay. Each advertiser has a different utility curve for each one of the two slot. At high prices, bidders will prefer having the second slot over the first one. It turns out that in this example, no efficient Nash equilibrium exists.

conversion rate, on average, is achieved in the two lowest slots on the side bar (slots seven and eight).

4. GSP IS SUBOPTIMAL

In this section, we show that with non-uniform conversion rates there are profiles of valuations for which no welfare-maximizing equilibria exist in GSP. We prove this claim for the special case of 2-slot 3-bidder auctions. We will analyze this 2-slot case in detail as it highlights significant differences between the non-uniform conversion rate case and the uniform case studied by previous work of [8, 6].

We will first give an intuition for the mathematical effect of non-uniform conversion rates. Figure 3 shows the utility of bidder i when allocated to slot j ($j \in \{1, 2\}$) as a function of his payment in the standard model with *uniform* conversion rates. This utility is a linear function of the price with a negative slope, $c^j(v_i - p)$, where v_i here represents bidder i's value per click. An immediate property of this model, as shown in Figure 3, is that at all price levels, the utility from slot one is at least that of slot two. With non-uniform conversion rates, however, this property no longer holds. As Figure 4 shows that with non-uniform conversion rates the curves shift; for sufficiently high price levels, bidders will actually prefer the lower slot! Intuitively, when the price of a click is high, the bidder may prefer to have a smaller number of clicks in the lower slot when these clicks have higher quality (that is, higher conversion rates).

Theorem 2. *2-slot GSP admits a welfare-maximizing equilibrium if and only if* $v_3 \le \frac{a^1c^1 - a^2c^2}{a^2c^1 - a^2c^2} v_1$ *and the three highest valuations satisfy* $v_1 > v_2 > v_3$.

Proof. Consider two price levels. Let β be the maximal price that advertiser three would be willing to pay for slot two, i.e., $\beta = v_3 \cdot a^2$. Let α be the minimal price for which advertiser one prefers slot two over slot one, i.e., the payment α for which $c^1(a^1v_1 - \alpha) = c^2(a^2v_1 - \alpha)$, i.e., $\alpha = \frac{c^1a^1 - c^2a^2}{c^1 - c^2} \cdot v_1$. Notice that $\beta \le \alpha$ is the condition of the lemma.

First, we show that $\alpha < \beta$ implies no welfare maximizing equilibria. Since $v_1 > v_2 > v_3$, for maximizing the social welfare, we must have that $b_1 > b_2 > b_3$. We will now use the fact that advertiser three would be willing to be assigned to the second slot with any payment less than β and advertiser one prefers slot two over slot one for payments higher than β. Therefore, if $b_2 < \beta$, advertiser three is better off by bidding b_2 since his utility for the second slot is non-negative. If $b_2 \ge \beta$, then advertiser one is better off by bidding b_2, since for such prices he prefers a lottery of the second and first slots over receiving the first slot for sure. Thus, neither case could be an equilibrium.

Second, if $\alpha \ge \beta$ then the following set of bids is an equilibrium: $b_3 = a^2v_3$, $b_2 = b_3 + \epsilon$ and b_1 is any bid greater than p^*, where p^* is the price for which winning slot one at price p^* gains advertiser two exactly the same utility as winning slot two and paying b_2. □

The proof implies that GSP will lack welfare-maximizing equilibria whenever the crossing point of the two utility curves of advertiser one (denoted in the proof by α) is smaller than $\beta = v_3a^2$ (the highest payment for which advertiser three desires slot two). It is this observation that enables our bounds on inefficiency of GSP's equilibria; when GSP is inefficient, it must be that the advertiser values are not too far apart, so it cannot be too inefficient.

5. BOUNDS ON THE WELFARE LOSS

In the previous section, we saw that sometimes GSP does not admit an efficient equilibrium with non-uniform conversion rates. How do equilibria look like in GSP in such cases? We present a simple characterization of such equilibria. It turns out that given any equilibrium, there also exists an *order-preserving* equilibrium where the advertisers' bids are weakly ordered by the advertisers' values. Note that these equilibria will not maximize social welfare, since some of the bids may be tied and we break ties uniformly. It turns out that such ties are inevitable under non-uniform conversion rates, and they will be central in our analysis.

Definition 2. *An equilibrium* $\mathbf{b} = (b_1, ..., b_n)$ *is* order preserving, *if for every two bidders* i_1, i_2, *if* $v_{i_1} > v_{i_2}$ *then* $b_{i_1} \geq b_{i_2}$

LEMMA 1. *Consider a profile of valuations* $v_1, ..., v_n$ *for which an equilibrium exists, denoted by* $\mathbf{b} = (b_1, ..., b_n)$. *Then, there also exists an* order preserving *equilibrium* $\mathbf{b}^*$. *Moreover,* $\mathbf{b}^*$ *achieves at least the same social welfare as* $\mathbf{b}$.

PROOF. Notation: Let $u_i(v_i, \mathbf{b})$ denote the expected utility of bidder i with value v_i when the bidders bid $\mathbf{b}$. That is, the expected utility is (let $n_i(\mathbf{b})$ denote the number of bidders bidding exactly b_i and let the consecutive slots that they are assigned be $\underline{t}, ..., \overline{t}$).

$$u_i(v_i, \mathbf{b}) = \frac{1}{|n_i(\mathbf{b})|} \left(\sum_{\underline{t} \leq j \leq \overline{t}} c^j a^j v_i - (|n_i(\mathbf{b})| - 1) b_i - b_{i'} \right)$$

where $b_{i'}$ denotes the highest bid in $\mathbf{b}$ which is smaller than b_i.

We will now prove the lemma. Consider an equilibrium $\mathbf{b}^* = (b_1^*, ..., b_n^*)$. Consider two bidders h and l such that $b_h^* > b_l^*$ but $h > l$. We will show that the bid profile where bidder l bids b_k^* and bidder k bids b_l^* is also an equilibrium. It follows that we can continue swapping bidders until reaching an order-preserving equilibrium.

Let $\mathbf{b}'$ be a profile of bids identical to $\mathbf{b}^*$ except bidder h bids b_l^*. Since $\mathbf{b}^*$ is an equilibrium, clearly $u_h(v_h, \mathbf{b}^*) \geq u_h(v_h, \mathbf{b}')$. Note that both $u_h(v_h, \mathbf{b}^*)$ and $u_h(v_h, \mathbf{b}')$ are linear functions of v_h. The slope of $u_h(v_h, \mathbf{b}^*)$ is $\frac{\sum_{j=t,...,t+n_h(\mathbf{b})} c^j a^j}{n_h(\mathbf{b})}$ (the first slot assigned to bidders that bid b_h is t) and the slope of $u_h(v_h, \mathbf{b}')$ is $\frac{\sum_{j=t',...,t'+s} c^j a^j}{n_h(\mathbf{b})}$ (where s is the number of slots assigned to bidders who bid b_l^* and t' is the first such slot), which is clearly smaller since we assume that $a^j c^j$ is decreasing in j.

It follows, that if we increase v_h to be v_l, we will have that $u_h(v_l, \mathbf{b}^*) \geq u_h(v_l, \mathbf{b}')$. That is, for the bid profile $\mathbf{b}^{swap}$ which is similar to $\mathbf{b}^*$ except h bids v_l and l bids v_h, bidder l will not be willing to bid b_l. Since $\mathbf{b}^*$ was an equilibrium, bidder l will not be willing to bid any other bid level in $\mathbf{b}^{swap}$ for the same reason. Similar arguments show that h will not be willing to change his bid as well. $\square$

5.1 The Social Welfare Loss Might be Substantial

We now give our first bound on the welfare loss in GSP with non-uniform conversion rates. For given parameters (CTR's and conversion rates), we construct a profile of valuation such that *all equilibria* incur a large welfare loss. We show that as the parameters of the problem become extreme (but still agree with all the assumptions in our model), the welfare loss can be arbitrarily close to $\frac{1}{3}$ of the total welfare.

The reason for GSP inefficiency is that advertisers with distinct values bid in pools. [3] To analyze the equilibria we distinguish between the pooling outcomes of the three advertisers with the highest values.

DEFINITION 3. *An order preserving equilibrium is*

- non-pooling *if the top three advertisers have distinct bids, i.e.,* $b_1 > b_2 > b_3$,
- 3-pooling *if the top three advertisers bid the same, i.e.,* $b_1 = b_2 = b_3$,
- $(1,2)$-pooling *if the top advertiser beats the second two who pool, i.e.,* $b_1 > b_2 = b_3$, *and*
- $(2,1)$-pooling *if the top two advertisers pool and beat the third, i.e.,* $b_1 = b_2 > b_3$.

The above definitions include the case where the fourth and lower valued advertisers pool with the third highest value advertiser.

For 2-slot auctions, we simplify notation and denote $c^1/c^2 = x$ and $a^1/a^2 = y$. Notice that by our previous monotonicity assumptions $x \geq 1$ and $y \leq 1$. We also denote m to be the number of players whose values v_i are at least the third highest value b_3 ($3 \leq m \leq n$).

THEOREM 3. *For every* $\epsilon > 0$, *and given any set of click-through rates and conversion rates, there exist valuations profile for which the social welfare achieved by* every *equilibrium of GSP is at most a* $\frac{xy+y}{xy+\frac{xy+y}{xy+1}} + \epsilon$ *fraction of the optimal welfare. When* $xy = 1$ *and* $y \to 0$, *this fraction approaches* $\frac{2}{3} + \epsilon$.

PROOF. Consider the following set of m bidders in a 2-slot auction. We first denote $\beta = \frac{c^1 a^1 + c^2 a^2}{c^1 + c^2}$, and note that βv is exactly the price where bidders with value v are indifferent between losing and winning the lottery of slot 1 and 2. Bidder 1 has value v_1, and all the other $n-1$ bidders have the same value v_2, such that $a^1 v_1 < \beta v_2$, or equivalently $v_2 > \frac{xy+y}{xy+1} v_1$. Denote $\delta = v_2 - \frac{xy+y}{xy+1} v_1$, and note that δ can be made arbitrarily small and that it is defined independently of m. Figure 5 in the appendix describes utility functions for this construction. Assumption: the exact bid level βv_2 is not in the bid space for any discretization. Denote the highest winning bid by p^*.

For the above valuations, every 3-pooling equilibrium will gain social welfare which is arbitrarily close to $\frac{2}{3}$ of the social welfare in the optimal slot allocation when $xy = 1$ and $y \to 0$.

$$\begin{aligned}
&\frac{\frac{1}{m}(a+b)(v_1 + (m-1)v_2)}{av_1 + bv_2} \\
&\leq \frac{\frac{1}{m}(xy+1)\left(v_1 + (m-1)\frac{xy+y}{xy+1}v_1 + (m-1)\delta\right)}{xyv_1 + \frac{xy+y}{xy+1}v_1 + \delta} \\
&= \frac{(xy+y)v_1}{xyv_1 + \frac{xy+y}{xy+1}v_1 + \delta} \\
&\quad + \frac{\frac{xy+1}{m}v_1 + \frac{m-1}{m}(xy+1)\delta - \frac{1}{m}(xy+y)v_1}{xyv_1 + \frac{xy+y}{xy+1}v_1 + \delta} \\
&< \frac{xy+y}{xy + \frac{xy+y}{xy+1}} + O(\max\{\frac{1}{m}, \delta\})
\end{aligned}$$

Claims 1–5, below, establish that the only equilibria in the above setting are 3-pooling. Finally, if the parameters are chosen such that $xy = 1$ and y approaches zero, then clearly $\frac{xy+y}{xy+\frac{xy+y}{xy+1}}$ approaches $\frac{2}{3}$. $\square$

CLAIM 1. *For the above valuations, GSP admits no non-pooling equilibrium.*

[3] Identifying pooling and separating equilibria plays a central role in microeconomics, mainly in the context of signaling games (see [4]).

This above claim follows immediately from Theorem 2.

Claim 2. *For the above valuations, GSP admits no* $(2,1)$*-pooling equilibrium.*

Proof. If $p^* > \beta v_2$, then bidder 2 prefers losing over his current payoff. If $p^* < \beta v_2$, then all the other losing bidders will gain positive utility from bidding p^*. Due to the above assumption, $p^* = \beta v_2$ cannot be an equilibrium since it is not in the bid space. □

Claim 3. *For the above valuations, GSP admits no* $(1,2)$*-pooling equilibrium.*

Proof. First, $p^* \leq a^1 v_1$, otherwise bidder 1 will prefer losing. Since we selected v_2 such that $a^1 v_1 < \beta v_2$, it follows that each one of the bidders with the value v_2 would earn some positive utility by bidding p^* and sharing Slots 1 and 2 with bidder 1. Since this positive utility is independent of the number of bidders m, for sufficiently large m bidder 2 will prefer bidding p^* over sharing Slot 2 with the other $m-2$ bidders, in contradiction to the existence of an equilibrium. (Note that the bidders that win together Slot 2 gain positive utility, thus if such an equilibrium had existed, then all the bidders except bidder 1 would have shared Slot 2.) □

Claim 4. *For the above valuations, GSP admits a 3-pooling equilibrium.*

Proof. It is easy to see that bidding the highest available bid p^* such that $p^* \leq a^1 v_1$ by all bidders is an equilibrium. Since $a^1 v_1 < \beta^* v_2$, every bidder with value v_2 gains positive utility, and thus bidder 1 also gains positive utility. However, all bidders would gain zero utility by under-bidding p^* and gain negative utility by over-bidding. □

Claim 5. *For the above valuations, GSP admits no non-order-preserving equilibria (for the top three advertisers).*

Proof. This claim is a direct corollary of Lemma 1. A 3-pooling equilibria is the only order-preserving equilibria that exists and it cannot be a result of a swapping process starting from other equilibria. The other order-preserving equilibria do not exist and therefore neither to non-order-preserving equilibria that can be converted to them via Lemma 1. □

5.2 Existence of Equilibria with a Bounded Loss

In the proof of Theorem 3 we constructed a family of bidder profiles for which all the equilibria achieve social welfare which is arbitrarily close to $\frac{2}{3}$ of the optimum. In this section we show that this bound is almost tight, at least for 2-slot auctions. We show that there always exists an equilibrium that achieves $\min\{\frac{2}{3}, 1/2 + \frac{a^1}{a^2}\}$ of the optimal social welfare; this does not rule out the existence of other equilibria with worse performance. This kind of results is commonly referred to as the *price of stability* [2], and it is similar in spirit to the work of [8, 6] that showed the existence of at least one welfare-maximizing equilibrium in GSP. We were able to prove this result only for 2-slot auctions, and we conjecture that a similar constant lower bound on the sub-optimality of GSP may be proved for any number of slots; however, we leave this problem open. Note that it is reasonable in this model to study a small constant number of slot, rather than asymptotically growing number of slots, therefore the 2-slot case can give us good intuition on the more general case.

We will start by proving that pure Nash equilibria always exist, even under the assumption of non-uniform conversion rates. Again, we were only able to prove this for the 2-slot case. The main reason is that our current proof relies on a case analysis that does not extend to a general number of slots. A new technique should probably be used for generalizing this lemma.

Lemma 2. *With a discrete bid space, GSP for two slots always admits a (pure Nash) equilibrium, even with non-uniform conversion rates.*

Proof. (sketch) Consider the following iterative ascending-price process. All bids are initialized to zero. Each bidder in his turn ($i = 1, ..., n$) raises his bid by ϵ (the minimal increment allowed) if it improves his utility. Also, losing players will raise their bid in their turn as long as they do not decrease their utility by doing so. We claim that this process will always end up in an equilibrium.

Consider the three bidders with the highest values (bidders 1,2,3), and consider the first iteration where one of them will not increase his bid. If the three bidders stop bidding at the same iteration, then it is clearly an equilibrium (they do not to bid higher bids since they stopped, but they do earn positive utilities; if they bid a lower bid they will get zero). If player 2 is the first to stop bidding, this is an equilibrium as well, since bidder one preferred the higher bid, and player 1 and 2 preferred not to raise their bid. Finally, if bidder 3 stopped increasing his bid, it is not necessarily an equilibrium since the iterative process may proceed, but since we require that the losing bidders will continue raising their bids in the algorithm, all we have to show is that bidder 1 will not prefer decreasing his bid when he shares the highest bid with bidder 2. This is clear, since bidder 1 preferred moving to the current bid over sharing slot 2 with the others, and thus gain non-negative value from slot 1. He will therefore prefer the lottery of slots 1,2 over having slot 2 with probability of at most 1/2 (since losers will always continue bidding). □

Corollary 1. *GSP with 2 slots always admits an order preserving (pure) equilibrium.*

Proof. Immediate from Lemmas 2 and 1. □

The following theorem shows that the equilibrium social welfare is not fatally low. As long as the conversion rates in the slots are not extremely far from each other (i.e., $a^2 \leq 3a^1$ and this is compatible with what we observed from the data), GSP achieves at least $\frac{2}{3}$ of the optimal welfare (with 2 slots).

Theorem 4. *In GSP for 2 slots and any number of bidders there always exists an equilibrium that achieves at least a fraction of* $\min\{\frac{2}{3}, \frac{1}{2} + \frac{a^1}{a^2}\}$ *of the optimal social welfare.*

The theorem is immediately concluded from the following lemma proved in Appendix 4. This is proved by a case analysis that bounds the social welfare loss in each equilibria for each possible pooling. Recall that the loss in social value in order-preserving equilibria is due to the randomized assignments in case of ties.

An important part of the analysis is the use of the parameterized constraints on the values of the bidders that are derived by the equilibrium constraints. We formulate these constraints to show that the values of the bidders that bid the same bid level cannot have values that are too far from each other, and therefore the loss in social welfare is limited. One subtle issue is that in the third part of the lemma, we require a more strict requirement than the equilibrium requirements that correspond to 3-pooling equilibria, and we use the fact that we can assume that no (2,1)-pooling equilibrium exists (if such an equilibrium exists, we have a bound by the first item of the lemma).

LEMMA 3.

1. *Every (2-1)-pooling equilibrium achieves at least $\frac{3}{4}$ of the optimal social welfare.*

2. *Every (1,2)-pooling equilibrium achieves at least $\frac{1}{2}+y$ of the optimal social welfare.*[4]

3. *Every 3-pooling equilibrium achieves at least $\frac{2}{3}$ of the optimal social welfare (when (2,1)-pooling equilibria do not exist).*

Actually, the proof of Lemma 3 gives us a better lower bound that depends on all the parameters of the model (CTR's, conversion rates, number of bidders). The $\min\{\frac{2}{3}, \frac{1}{2}+\frac{a^1}{a^2}\}$ bound is only a simplification. The better bound is a minimum of several parameterized functions, and is given in Appendix B.

6. GASP: CONVERSION RATE AWARE GSP AUCTIONS

Conversion tracking allows us to estimate the likelihood that a click will generate a conversion when the user visits the advertiser's website. We presented empirical data (in Section 3) that showed that the conversion rates on bottom slots are non-uniform across slots.

Unfortunately, the equilibrium results of [8, 6] which show that GSP maximizes the social welfare (i.e., performs optimally) fails when conversion rates are not uniform across slots. In Section 5, we have shown that for reasonable click-through rates, conversion rates, and advertiser valuations; the social welfare of GSP can be as much as 33% less than optimal.

In this section we demonstrate how we can modify the payment rule in GSP using conversion rates that are automatically calculated by the conversion tracking system. The equilibrium in the resulting auction maximizes the social welfare.

One immediate suggestion might be to move from a pay-per-click auction to a pay-per-conversion model. However, this transition may be too drastic and fast for such large businesses; an incremental improvement of the current system, which will maintain the same pay-per-click business model and the same user interface but will take into account the conversion tracking data seems to be desirable. This is exactly what we suggest in the following variant of GSP which we call GASP. The difference from standard GSP is in how the payments are normalized. Note that that since GASP preserves the pay-per-click business model, it can concurrently sell ads to bidders with and without conversion rate tracking systems (e.g., for advertiser with offline conversions).

The Generalized Acquisition-aware Second Price Auction. The generalized acquisition-aware second price (GASP) auction, for k slots, is given input bids from advertisers. Let b_i represent the *bid-per-conversion* of advertiser i.

1. Let $b_i c_i a_i$ be the *bid-per-impression* of advertiser i.

2. Sort the bidders by bid-per-impression:

   - Let i_j be the index of the bidder with the jth highest bid-per-impression.
   - For all j, $b_{i_j} c_{i_j} a_{i_j} \geq b_{i_{j+1}} c_{i_{j+1}} a_{i_{j+1}}$.
   - For all j, let p_j be the minimum bid necessary for advertiser i_j to maintain their position in slot j:

$$p_j = \frac{b_{i_{j+1}} c_{i_{j+1}} a_{i_{j+1}}}{c_{i_j} a_{i_j}}.$$

3. For $j \leq k$, show advertiser i_j in slot j.

4. For $j \leq k$, if ad i_j is clicked, charge $a^j_{i_j} p_j$ to advertiser i_j.[5]

Game-theoretic Analysis of GASP. A game-theoretic analysis of GASP follows directly from the analysis of GSP in [8, 6]. Under the assumption that click-through rates and conversion rates are factorable (which was also required in GSP, see Section 2), and that the impressions in top (lower numbered) slots lead to more conversions than bottom (higher numbered) slots; in equilibrium, GASP maximizes the social welfare.

PROPOSITION 1. *GASP always admits a welfare-maximizing equilibrium, even with non-uniform conversion rates.*

The proof of this claim is based on the following observation about an advertiser's utility. Using the prices in the definition of the GASP mechanism, actually the conversion rates can be factored out in the utility functions of the bidders, returning to the same utility structure discussed for uniform conversion rates by [8, 6]. In other words, we move from a utility structure as in Figure 4 back to the structure in Figure 3. Formally, advertiser i's utility for slot j at price $a^j_i p$ is now

$$u^j_i(p) = v_i c^j_i a^j_i (v_i - p),$$

since if there is a conversion (with probability $c^j_i a^j_i$), the advertiser gets their value v_i, but if there is a click (with probability c^j_i), the advertiser must pay $a^j_i p$. Notice that our assumption on the monotonicity of values per impression across the slots allows us to assume that at any p all advertisers have higher utility for top (lower numbered) slots.

[4]We conjecture that a better analysis may improve this bound to $\frac{2}{3}$.

[5]First $a^j_{i_j} p_j$ simplifies to $b_{i_{j+1}} c_{i_{j+1}} a^j_{i_{j+1}} / c_{i_j}$. Second, if wanted a pay-per-conversion auction instead of a pay-per-click auction we would simply charge p_j advertiser i_j upon conversion. Constrained to a pay-per-click model, we scale this payment by i_j's conversion rate for slot j to get the appropriate payment.

7. REFERENCES

[1] Z. Abrams, A. Ghosh, and E. Vee. Cost of conciseness in sponsored search auctions. In *Proc. 3rd International Workshop on Internet and Network Economics*, 2007.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS 04*, pages 295–304, 2004.

[3] S. Muthukrishnan Gagan Aggarwal, Jon Feldman. Bidding to the top: Vcg and equilibria of position-based auctions. In *Proceedings of the Fourth Workshop on Approximation and Online Algorithms*, 2006.

[4] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford university press, 1995.

[5] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions, 2007. Working paper, Stanford University.

[6] Michael Ostrovsky, Benjamin Edelman, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Forthcoming in American Economic Review*, 2006.

[7] Tim Roughgarden and Eva Tardos. *In Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 17. Introduction to the Inefficiency of Equillibria.* Cambridge University Press., 2007.

[8] Hal Varian. Position auctions. *To appear in International Journal of Industrial Organization*, 2006.

[9] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

APPENDIX

A. PROOF OF LEMMA 3

Notations: recall that m denotes be the number of bidders that have at least the 3rd-highest value (i.e., where $v_i \geq v_3$) and that $a = c^1a^1$, $b = a^2c^2$, $x = \frac{c^1}{c^2}$ and $y = \frac{a^1}{a^2}$, and let p be the lowest winning bid. In the following three subsections we prove the items in Lemma 3.

A.1 A bound on the welfare loss in 3-pooling equilibria

This subsection proves the following lemma.

LEMMA 4. *Every 3-pooling equilibrium achieves at least $\frac{2}{3}$ of the optimal social welfare (if (2-1)-pooling equilibria do not exist).*

Note that the optimal welfare is $av_1 + bv_2$, and the welfare when $m > 2$ bidders bid the highest bid is $\frac{1}{m}(a + b)(v_1 + v_2 + (m-2)v_3)$. In this proof, we will denote the ratio between them as WR (for Welfare Ratio), and we will bound it from below.

$$WR = \frac{\frac{1}{m}(a+b)(v_1 + v_2 + (m-2)v_3)}{av_1 + bv_2} \quad (1)$$

We first observe that the inefficiency ratio is either always increasing or always decreasing in v_2. It is easy to see (by first-order conditions) that the function is increasing in v_2 if $v_3 < \frac{(m-1)b}{a-b}v_1$ and decreasing otherwise. We will treat these two cases separately.

Case 1: *The function in Equation 1 is decreasing in v_2.*

We can now change v_2 to be v_3 and the ratio will not increase.

$$WR > \frac{\frac{1}{m}(a+b)(v_1 + (m-1)v_3)}{av_1 + bv_3} \quad (2)$$

Since this is a 3-pooling equilibrium, we will now write constraints derived by the equilibrium definition. These constraints will help us bound the WR function.

If bidding p is an equilibrium bid for the top m bidders, then bidder 1 will not benefit from deviating and bid $p + \epsilon$. Therefore,

$$\frac{1}{m}c^1(a^1v_1 - p) + \frac{1}{m}c^2(a^2v_1 - p) \geq c^1(a^1v_1 - p)$$

Since bidder 4 gains non negative utility,

$$\frac{1}{m}c^1(a^1v_3 - p) + \frac{1}{m}c^2(a^2v_3 - p) \geq 0 \quad (3)$$

The two above inequalities give an upper bound and a lower bound on p; taken together, we get that $v_3 \geq tv_1$ where

$$t = \frac{((m-1)a - b)\left(c^1 + c^2\right)}{((m-1)c^1 - c^2)(a+b)} = \frac{((m-1)xy - 1)(x+1)}{((m-1)x - 1)(xy+1)} \quad (4)$$

Since Equation 2 is increasing in v_3 (assuming $b < 1$ and $m > 2$), we get that:

$$WR > \frac{\frac{1}{m}(a+b)(v_1 + (m-1)tv_1)}{av_1 + btv_1} = \frac{\frac{1}{m}(a+b)(1 + (m-1)t)}{a + bt} = \frac{\frac{1}{m}(xy+1)(1+(m-1)t)}{xy + t} \quad (5)$$

It is easy to see that in this case $WR \geq \frac{xy+1}{xy+2} = \frac{a+b}{a+2b} > \frac{2}{3}$.[6]

Case 2: *The function in Equation 1 is increasing in v_2.*

We can now change the v_2 in equation 1 to be v_1 and the ratio will not increase.

$$WR > \frac{\frac{1}{m}(a+b)(2v_2 + (m-2)v_3)}{av_2 + bv_2} \quad (7)$$

[6]From Equation 4 clearly $t \geq \frac{(m-1)xy-1}{(m-1)(xy+1)}$. Substituting into Equation 5 we get:

$$\frac{\frac{1}{m}(xy+1)(1+(m-1)t)}{xy+t} = \frac{xy(m-1)(xy+1)}{(m-1)(x^2y^2+xy) + (m-1)xy - 1} \geq \frac{xy(xy+1)}{(x^2y^2+xy)+xy} = \frac{(xy+1)}{xy+2} \quad (6)$$

In the analysis of this part of the proof we will actually need a stronger constraints than the one that follows from the 3-pooling equilibrium properties. We will assume that there is no (2,1)-pooling equilibrium. (We can do this, since part 1 of Lemma 3 shows that (2,1)-pooling equilibria achieve at least $\frac{3}{4}$ of the social welfare.) It follows that bidder 2 prefers sharing the second slot at price p with the bidders with the value v_3 over sharing the highest bid with bidder 1:[7]

$$\frac{1}{m-1}\left(c^2(a^2v_2-p)\right) \quad > \quad \frac{1}{2}c^1(a^1v_2-(p+\epsilon))+\frac{1}{2}c^2(a^2v_2-p)$$

After rearranging the terms and together with Equation 3 we get:

$$v_3 \quad > \quad \frac{(a+\left(1-\frac{2}{m-1}\right)b)(c^1+c^2)}{(a+b)\left(c^1+\left(1-\frac{2}{m-1}\right)c^2\right)}v_2$$

And we denote $v_3 > tv_2$, and note that $t > \frac{a+\left(1-\frac{2}{m-1}\right)b}{a+b}v_2$.

Using Equation 7 we get the following bound on the inefficiency:

$$\begin{aligned} WR &> \frac{\frac{1}{m}(a+b)\left(2v_2+(m-2)tv_2\right)}{av_2+bv_2} \\ &= \frac{1}{m}\left(2+(m-2)t\right) \\ &> \frac{1}{m}\left(2+(m-2)\frac{a+\left(1-\frac{2}{m-1}\right)b}{a+b}\right) \\ &= 1-\frac{(m-2)2b}{m(m-1)(a+b)} \\ &> \frac{5}{6} \end{aligned}$$

Where the last equality holds since $\frac{2(m-2)}{m(m-1)} < \frac{1}{3}$ for $m > 2$ and since $b < a$.

A.2 A bound on the loss in (1,2)-pooling equilibria

Lemma 5. *Every (1,2)-pooling equilibrium achieves at least $\frac{1}{2}+\frac{a^1}{a^2}$ of the optimal social welfare.*

Proof. First, bidder 3 must have non zero utility from bidding p, where he receives Slot 2 with a positive probability and pays p. Thus, $p \leq a^2v_3$. On the other hand, bidder 2 prefers bidding p to bidding $p+\epsilon$. That is, he prefers no slot (zero utility) to Slot 1 at price $p+\epsilon$. Thus, $p+\epsilon > a^1v_2$. It follows that $a^2v_3+\epsilon > a^1v_2$,

$$v_3 > yv_2 - \frac{\epsilon}{a^2} \tag{8}$$

[7]If such an equilibrium does not exists, bidder 2 must want to deviate. The bidders with the value v_3 will not bid above p since p must be exactly the price for which bidder 3 is willing to get Slot 1 and Slot 2 at random, but this lottery will gain him negative utility for price $p+\epsilon$. Also, if bidder 1 will deviate, necessarily bidder 2 will deviate too.

$$\begin{aligned} WR &\geq \frac{av_1+\frac{1}{m-1}bv_2+\frac{m-2}{m-1}bv_3}{av_1+bv_2} \\ &\geq \frac{av_1+bv_2\left(\frac{1}{m-1}+\frac{m-2}{m-1}y\right)}{av_1+bv_2}-O(\epsilon) \qquad (9)\\ &\geq \frac{av_1+av_1\left(\frac{1}{m-1}+\frac{m-2}{m-1}y\right)}{av_1+av_1}-O(\epsilon) \\ &\geq \frac{1}{2}+\frac{1+(m-2)y}{2(m-1)}-O(\epsilon) \\ &> \frac{1}{2}+\frac{y}{2}-O(\epsilon) \qquad (10) \end{aligned}$$

Where the 2nd inequality is due to Equation 8, and the 3rd inequality is due to the fact that Equation 9 is decreasing in the value of bv_2 (since $\frac{1}{m-1}+\frac{m-2}{m-1}y<1$) , which is bounded by av_1.

□

A.3 A bound on the loss in (2,1)-pooling equilibrium

Lemma 6. *Every (2,1)-pooling equilibrium achieves at least $\frac{3}{4}$ of the optimal social welfare.*

Bidders 1 and 2 win by bidding p, and bidder 3 bids $p-\epsilon$. Again, we will consider constraints on the valuations of the bidders that follow from the equilibrium properties.

Bidder 2 must gain non-negative utility by bidding p.

$$\frac{1}{2}c^1\left(a^1v_2-p\right)+\frac{1}{2}c^2\left(a^2v_2-(p-\epsilon)\right)\geq 0$$

Therefore:

$$p \leq \frac{a+b}{c^1+c^2}v_2+\frac{\epsilon c^2}{c^1+c^2} \tag{11}$$

Bidder 1 will not increase his utility from bidding above p and winning Slot 1 alone:

$$\frac{1}{2}c^1\left(a^1v_1-p\right)+\frac{1}{2}c^2\left(a^2v_1-(p-\epsilon)\right)\leq c^1(a^1v_1-p)$$

Therefore:

$$p \geq \frac{a-b}{c^1-c^2}v_1 \tag{12}$$

From Inequalities 11 and 12 we get:

$$\begin{aligned} v_2 &\geq \frac{(a-b)(c^1+c^2)}{(a+b)(c^1-c^2)}v_1-O(\epsilon) \\ &> \frac{xy-1}{xy+1}v_1-O(\epsilon) \end{aligned}$$

Now we can bound the social-welfare loss. Note that the expected welfare in $(2,1)$-pooling equilibrium is $\frac{1}{2}(a+b)(v_1+v_2)$, and recall that the optimal welfare is given by av_1+bv_2. We will use the above constraints to prove the following lower bound on the ratio between the equilibrium welfare and the optimal welfare.

$$
\begin{aligned}
WR &= \frac{\frac{1}{2}(a+b)(v_1+v_2)}{av_1+bv_2} \\
&> \frac{\frac{1}{2}(a+b)(v_1+\frac{xy-1}{xy+1}v_1-O(\epsilon))}{av_1+b\frac{xy-1}{xy+1}v_1-O(\epsilon)} \\
&> \frac{\frac{1}{2}(a+b)(1+\frac{xy-1}{xy+1}-O(\epsilon))}{a+b\frac{xy-1}{xy+1}} \\
&= \frac{\frac{1}{2}(xy+1)(1+\frac{xy-1}{xy+1})}{xy+\frac{xy-1}{xy+1}}-O(\epsilon) \\
&= \frac{x^2y^2+xy}{x^2y^2+2xy-1}-O(\epsilon) \qquad (13) \\
&> \tfrac{3}{4}-O(\epsilon) \qquad (14)
\end{aligned}
$$

Where Inequalities 13 and 14 follow from simple algebra.

B. BETTER LOWER BOUND ON THE SOCIAL WELFARE LOSS

The proof of Theorem 4 (and Lemma 3) actually directly derives the following lower bound on the social-value loss. Simplifying this bound proves a $\min\{\frac{2}{3}, \frac{1}{2}+\frac{a^1}{a^2}\}$ on the equilibrium social welfare.

Let $f_1, ..., f_4$ be the following functions:

$$
\begin{aligned}
f_1 &= \frac{1}{m}\left(2+(m-2)\frac{(xy+\left(1-\frac{2}{m-1}\right))(x+1)}{(xy+1)\left(x+\left(1-\frac{2}{m-1}\right)\right)}\right) \\
f_2 &= \frac{\frac{1}{m}(xy+1)\left(1+(m-1)\frac{((m-1)xy-1)(x+1)}{((m-1)x-1)(xy+1)}\right)}{xy+\frac{((m-1)xy-1)(x+1)}{((m-1)x-1)(xy+1)}} \\
f_3 &= \frac{1}{2}+\frac{1+(m-2)y}{2(m-1)} \\
f_4 &= \frac{\frac{1}{2}(xy+1)(1+\frac{(xy-1)(x+1)}{(xy+1)(x-1)})}{xy+\frac{(xy-1)(x+1)}{(xy+1)(x-1)}}
\end{aligned}
$$

THEOREM 5. *In GSP for 2 slots and any number of bidders there always exists an equilibrium that achieves at least a fraction of* $\min\{f_1, f_2, f_3, f_4\}$ *of the optimal social welfare.*

C. EMPIRICAL METHODOLOGY

We describe our experimental methodology for aggregating advertiser conversion rates. Recall that the concept of "conversion rate" does not refer to the same thing for different advertisers and different industries, and we observe from the data that conversion rates vary considerably over different advertisers. Therefore, we normalized the conversion rates of each advertiser using the top slot (Slot 1 in the main line) as a benchmark before aggregating. We then computed the weighted average of these normalized conversion rates per each slot, removing advertisers with extremely low number of conversion to avoid bias from statistically insignificant data.

Overall, the data aggregation process of the data presented in Figure was as follows:[8]

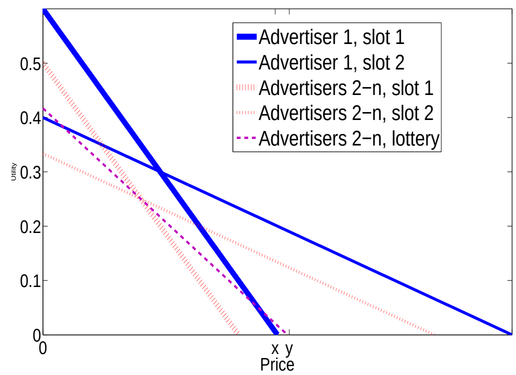

Figure 5: The figures show the utility structure of settings where the social welfare loss approaches $\frac{2}{3}$. The solid lines describe the utility of bidder 1, the dashed lines show the utility of the other bidders; the middle dashed line shows their utility from being allocation randomly to slot 1 or 2. Following the notations in the proof of theorem 3, $x = a^1 v_1$ and $y = \beta^* v_2$.

1. **Restriction to search queries with eight ads.** To control for conversion and click phenomena that result of a varying number of slots, we only considered search queries for which eight ads were shown.
2. **Normalization at the advertiser level.** For each advertiser and for each one of the eight slots, we calculated the conversion rate in this slot, relative to the conversion rate at slot one of the main line. I.e., the relative conversion rate at slot j for advertiser i is $\hat{a_i^j}/\hat{a_i^1}$ (where $\hat{a_i^j}$ stands for the empirical conversion rate).
3. **Ignoring data with very small number of samples.** For each advertiser we calculate the total empirical conversion rate $\hat{a_i}$ ("total number of conversions" / "total number of clicks"). For advertiser i we ignore the data in slot i if the expected number of clicks was too close to zero in terms of standard deviation.[9] Such ad impressions are be treated as if they received no clicks.
4. **Computing a weighted average of the relative conversion rates.** Finally, we computed a weighted average of the conversion rates separately for each slot ($j = 1, ..., 8$). The relative conversion rate of each advertiser at slot j is weighted by the number of clicks this advertiser received in slot j divided by the total number of clicks of all advertisers at slot j.

[8]For example, consider the following three advertisers in a three slot auction. Advertiser 1 had 1000, 900, and 500 clicks respectively on slots 1, 2 and 3, where 20%, 15% and 30% of the clicks were converted. Advertiser 2 had 2000, 800 and 400 clicks, and conversion rates of 5%, 7% and 12%. The third advertiser had 400, 200 and 30 clicks and conversion rates of 4%, 5% and 0%. The aggregate conversion rates are (100%, 108% and 190%. Note that the data of the slot 3 of the 3rd advertiser was disregarded since the number of clicks is too low.

[9]Specifically, we excluded data from slots where $E[conversions] - 2 \cdot STDEV < 0$ (STDEV is the standard deviation of the number of conversions in this slot). Namely, $(\#clicks) * \hat{a_i} - 2\sqrt{(\#clicks) * \hat{a_i} * (1-\hat{a_i})} < 0$ (where $\#clicks$ is the number of clicks for advertiser i in slot j.