

Mechanism Design for Mixed Ads

Yoram Bachrach, Microsoft Research
Sofia Ceppi, Microsoft Research
Ian A. Kash, Microsoft Research
Peter Key, Microsoft Research
Mohammad Reza Khani, University of Maryland

We propose a hybrid auction mechanism for sponsored search, where bidders can be truthful or not, and are accordingly treated differently. Our class of hybrid mechanisms give incentives for non-truthful bidders to bid truthfully, while behaving as a non-truthful auction if no bidders are truthful.

Our motivation is that the Generalized Second Price (GSP) auction (the current mechanism of choice) has appealing properties when ads are simple (text based and identical in size). But GSP does not generalize to richer ad settings, whereas truthful mechanisms, such as VCG do. Hence there are incentives for search platforms to migrate to truthful mechanisms, but a straight switch from GSP to VCG either requires all bidders instantly bid truthfully or incurs significant revenue loss.

We introduce a transitional mechanism which encourages advertisers to update their bids to their valuations, while mitigating revenue loss. The mechanism is equivalent to GSP when nobody has updated their bid, is equivalent to VCG when everybody has updated, and it has the same allocation and payments of the original GSP if bids were in the minimum symmetric Nash equilibrium. In settings where both GSP ads and truthful (TF) ads exist, it is easier to propose a payment function than an allocation function. We give a general framework for these settings to characterize payment functions which guarantee incentive compatibility of truthful ads, by requiring that the payment functions satisfy two properties.

Finally, we compare the revenue of our transitional mechanism with revenues of GSP and VCG mechanisms when run on a sample of Bing data.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Algorithmic Game Theory, Mechanism Design, Revenue Optimization, Experimentation

Additional Key Words and Phrases: online advertising, Generalized Second Price Auction, Truthful Auctions

ACM Reference Format:

Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, Mohammad Reza Khani, 2015. Mechanism Design for Mixed ads *ACM X*, X, Article X (February 2015), 20 pages.

DOI : <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Sponsored search is the main source of revenue for most search engines, such as Google, Yahoo! or Bing. In the classic online ad auction or ‘position auction’, all ad slots and sizes are the same. Search engines typically use a variant of the Generalized Second Price (GSP) mechanism to select and price ads. In GSP, advertisers are rank-ordered by decreasing bids (more generally, by expected revenue or rank score) and slots are assigned in this order. The price of a slot is the minimum bid an advertiser has to make in order to maintain that position, which equates to the next highest bid in the simplest form of GSP. Payment is made when an ad is clicked. The GSP auction’s equilibria, bidding strategies, and other properties are well studied (see, e.g., [Edelman et al. 2005; Varian 2007]).

However, online advertising is becoming more complex. There may be multiple page-templates for search results, different ad formats (e.g., text-ads or image-ads) with different sizes, and several other constraints on showing ads. For these settings, GSP is not well defined and if generalized can be ill-behaved [Bachrach et al. 2014]. Therefore, there is an incentive for migrating from GSP to another mechanism. Truthful mechanisms (such as VCG) are attractive (see [Varian and Harris 2014]) because they

allow externalities to be calculated which makes handling more complex ad scenarios easier. Truthful mechanisms also remove the computational burden of calculating the optimum bid from advertisers, make the whole system more transparent, and the same bid (valuation) of an advertiser can be used across multiple auctions. Moreover, analyzing the market is easier with true valuations as opposed to bids.

One big obstacle to migration from GSP to VCG is the requirement that advertisers update their bids by increasing them up to their true value. Indeed, as Varian and Harris note, Google “thought very seriously about changing the GSP auction to a VCG auction during the summer of 2002.” However, there were several problems, including that “the VCG auction required advertisers to raise their bids above those they had become accustomed to in the GSP auction.” As GSP has only gotten more entrenched over the past decade, this issue has only grown larger. If such a switch were made today, advertisers may not update their bids quickly and even if they do update, it might not be their true valuation. Thus, for a long time there will be a mixed set of ads: (i) advertisers who have updated their bids to true valuations and (ii) advertisers with GSP bids. Therefore, it would be desirable to have a transitional mechanism which selects and prices the set of winners from such a mixed set of ads.

In this work we introduce a transitional mechanism in order to migrate GSP bids to true valuations. At the start of the transition the mechanism behaves the same as GSP. Then (in one implementation of our mechanism) when an advertiser tries to update her bid, we notify her that the mechanism will optimize her allocation and payment assuming that her new bid is her true valuation. This means that the best response for the advertiser is to bid her true valuation. We show that since this optimization will not happen for old GSP bids, the mechanism actually encourages advertisers to update their bids to their valuations as soon as possible (this is exactly what we want them to do). In the middle of the transition, the mechanism maintains a classification of bids as either truthful (updated bids) or GSP so that they can be treated differently, respecting incentives for truthful ads, but not GSP ads. Finally when all advertisers update their bids, the mechanism behaves as VCG and the transition finishes.

The key difficulty for a transition mechanism is to decide how bids of one group affect allocation and payment of the other group. For example, if the bid of a truthful ad is higher than the bid of a GSP ad, it does not necessarily mean that the truthful ad has to have a better allocation, since we know the GSP ad is shading its bid downward. Moreover, if a truthful ad is given an ad slot with click probability f and the ad below him is a GSP ad with bid g , we may want the expected payment of the truthful ad to be at least $g \cdot f$. The usual approach for designing truthful mechanisms is to give a monotone allocation function and then derive the unique payments using Myerson’s lemma [Myerson 1981]. However, in our setting it is not easy to first give an allocation function, as the bid of a GSP ad implies that *payments* of ads above him should be altered. Hence, designing an allocation function which takes into account the payment constraints is harder.

In this paper we give a general framework, not limited to only GSP or VCG ads¹, for designing mechanisms in settings similar to the above where it is easier to specify a payment function than an allocation function. In practice, this flexibility allows our framework to apply in the presence of some of the additional complexity that exists in actual auction systems. To describe this framework more formally, assume ad slots are indexed from 1, the top slot, to n , the bottom one, and there are n advertisers. The designer just has to specify a payment function $p^i : \mathbb{R}^{n-i} \mapsto \mathbb{R}$ for each $i \in [n]$ which specifies the payment of the advertiser assigned to position i given types of advertisers

¹The set of ads can have arbitrary number of groups. For example in addition to GSP and truthful ads, it may contain advertisers who have pre-established contracts and pay a fixed amount if their ad gets clicked.

assigned to positions below him. Our framework requires that any payment function satisfies two simple properties: (i) Minimum Marginal Increase (MMI): the payment has to be high enough so that truthful bidders assigned to lower slots do not envy the winner of a higher slot and (ii) Exact Marginal Increase (EMI): the marginal payment increase of the slot directly above a truthful ad has to be equal to the truthful ad's bid. Given a set of payment functions satisfying MMI and EMI, our framework shows how to construct an allocation rule with payments that are exactly those given by the payment functions applied to the realized allocation.

While we primarily focus on deriving a practical rule for transitioning from GSP to VCG, the set of ad auction mechanisms that fit in our framework is quite general. In fact, we prove that by using our framework one can design any truthful mechanism in which the payment of an ad is derived solely from ads below that ad (subject to a few additional requirements). Equivalently, our framework encapsulates mechanisms where raising the bid of an ad does not affect the allocation of ads that were previously allocated below it. More broadly, while the question of what properties of allocation rules lead to truthful mechanisms has been intensively studied (see, e.g., [Archer and Kleinberg 2008; Ashlagi et al. 2010; Frongillo and and 2014]), the question of what properties of payment rules lead to truthful mechanisms has not. Indeed, the only prior characterization we know of is the taxation principle for single-agent mechanisms. Since EMI and MMI are more inspired by Myerson than by the details of our auction setting, this approach may be of independent interest.

Having designed a large class of candidate mechanisms, and selected a representative member of the class, we analyze it both theoretically and in simulations based on Bing data. On the theory side, we show that transitional behavior is particularly nice if the system starts in the lowest symmetric Nash equilibrium. In particular, allocations and prices do not change regardless of the order of updates, in principle leading to a painless transition.

This conclusion relies on a number of strong (and false) assumptions. In our simulations, we analyze the consequences of relaxing them. We find that essentially all the costs of a transition are in terms of revenue—the welfare effects for both advertisers and users are small. We also see that the hybrid mechanism can have significant revenue benefits if bidders directly update to true valuations, but these benefits decrease if the bidders fail to do that for various reasons such as not knowing the true valuations, not being utility maximizers, or not trusting the system. So, rather than simply telling advertisers that they will be treated differently once they change their bid (which may not even be feasible in practice), it may make sense to attempt to use a learning approach to identify which bidders to treat as truthful. We simulate a few behaviors, and consequently identification models, and analyze the performance of the hybrid, GSP, and VCG mechanisms.

To summarize, our three main contributions in this paper are:

- (1) a new framework for deriving truthful mechanisms from payment rules that satisfy MMI and EMI (Section 4),
- (2) a specific hybrid mechanism to enable transitioning from GSP to VCG (Section 5), and
- (3) an evaluation of the mechanism based on Bing data (Section 6).

2. RELATED WORK

Sponsored search auctions are arguably the most successful recent application of auction theory to a business environment. As a result, much research has been conducted regarding the influence of the mechanism used for the auction on social welfare and the generated revenue. In the case where the VCG mechanism is used, truthfulness is

the dominant bidding behavior. However, the same does not hold for the GSP auction and predicting bids in this case is trickier.

A complete information analysis of GSP auctions is discussed by Edelman et al. [2005], Varian [2007], and Aggarwal et al. [2006]. A common theme in this line of work is the equivalence between the auctioneer’s revenue and bidders’ utility under a VCG auction and under the lowest symmetric Nash equilibrium of a GSP auction (which is sometimes referred to as the “bidder-optimal locally envy free equilibrium”). Ashlagi et al. [2007] generalize this, showing that in many auction types in which the payments are a function of the lowest ranked bids, there exists an equilibrium in which bidders’ utility is equivalent to their utility under the VCG auction. Roberts et al. [2013] generalize this along a different axis, showing that this result also holds for a variety of rank score functions other than simply ranking by highest bid.

Much of the research on equilibria in GSP auctions has focused on symmetric equilibria. Edelman and Schwarz [2010] examined the revenue of different symmetric Nash equilibria, noting that under a certain comparison to optimal revenue possible under the Bayesian setting, the “lowest” equilibrium is the reasonable one. A generalized auction proposed by Aggarwal et al. [2007] allows advertisers to specify not only a bid but also the positions they are interested in, ruling out the bottom positions. They show that this auction has a symmetric Nash equilibrium implementing the same outcome (i.e., allocation and pricing) as the VCG auction.

Complementary to studies on symmetric equilibria, several researchers have studied the inefficiency that can result from asymmetric equilibria [Leme and Tardos 2009; Caragiannis et al. 2011; Lucier et al. 2012]. Some studies of auction tuning have also explored the full set of equilibria [Thompson and Leyton-Brown 2013]

Taking a Bayesian perspective, Gomes and Sweeney [2014] examined the existence and uniqueness of Bayes-Nash equilibria in a GSP auction. Several models have also been proposed for inferring the valuations of advertisers based on the observed bid data [Pin and Key 2011; Athey and Nekipelov 2010]. The model by Pin and Key [2011] considers advertisers best responding in an uncertain environment in a repeated auction setting, relating the bidding behavior to scenarios when the Bayes-Nash Equilibria of Gomes and Sweeney [2014] are known to exist. The model of Athey and Nekipelov [2010] starts directly from the Bayes-Nash Equilibria, but has a different model of the information available to the bidders. Instead, Vorobeychik [2009] proposed a framework based on agent simulation to approximate the Bayes-Nash equilibria in GSP auctions, which relies on restricting the space of allowed bidding strategies.

The dynamics leading to the equilibrium outcomes in GSP auctions are less studied. Cary et al. [2007] consider dynamics under a greedy bidding strategy, where each bidder chooses the optimal bid for the next round assuming the other bidders do not change their bids. They show this bidding strategy has a unique fixed point, with payments identical to those of the VCG mechanism.

Closest to our work, Aggarwal et al. [2009] propose a framework that frames both GSP and VCG in terms of the assignment game with appropriate models of bidder utility. Their framework can be applied to derive a hybrid auction that incorporates both GSP and VCG bidders that is a special case of our more general framework. However, they do not explore this and their framework lacks the flexibility ours provides.

3. PRELIMINARIES

We study the standard model of a sponsored search auction. There is a set of $\{1, \dots, n\}$ of ads, denoted by $[n]$. We assume that there are k ad positions ($[k]$), also called slots, where position i has CTR (Click-Through-Rate) f_i . Without loss of generality we assume that the first position has the highest CTR and the k -th position has the lowest ($0 < f_k < \dots < f_1 \leq 1$). For notational convenience, and without loss of generality, we

take $k = n$ and assume that there are no ad quality scores—the probability of any ad being clicked in slot i is exactly f_i .

In our setting ads can have different characteristics. For example they could be,

- Truthful (TF) ads which have updated their bids to their valuations;
- Generalized Second Price (GSP) ads which assume that they are participating in a GSP auction.
- First price ads which have a contract to pay a fixed amount upon being clicked.

We are going to provide a framework for mechanism designers which only requires a payment rule to be specified, and which places special constraints on the payments of TF ads. Therefore, we specially differentiate between TF and non-TF ads and assume each ad has a type taken from set $\mathcal{T} = \{TF, \text{non-TF}\} \times \mathbb{R}^+$ that specifies the truthfulness attribute and bid. Here, non-TF ads can be of any nature and our framework does not limit the designer’s ability for deciding the allocation and payments of them. This can be thought of as modeling a situation where we are designing a system for the TF bidders, and therefore care about their incentives. However, there will be some “legacy” bidders whose behavior does not reflect the new system (and perhaps cannot, because inherently changing their bid will result in them being reclassified as TF). Since the right way to treat these legacy bidders will depend on their exact nature, we do not put any constraints on what they are charged (and indeed neglect modeling this entirely in the proving the theoretical results of Section 4).

In section 5, we examine what happens when we look to maintain parity with existing GSP prices for non-TF bidders, and where the non-TF bidders behave as GSP bidders in equilibrium.

We denote an assignment of ads to slots by the permutation $\Pi = (\pi_1, \dots, \pi_n)$ where ad π_i is assigned to position i . The (expected) payment of ad π_i is the cost per click for being in position i multiplied by the CTR of position i . Throughout the paper we work with expected payments $p_i, \forall i \in [n]$ as opposed to cost per click. Further, we assume that all ads have quasilinear utilities, *i.e.*, if ad π_i is assigned to position i and pays p_i then its utility will be $u_i(\Pi, p) = f_i * v(\pi_i) - p_i$, where $v(\pi_i)$ is the valuation of ad π_i . Throughout the paper, we use $b(i)$ for the bid of ad i and if a TF ad, we also use $v(i) = b(i)$ to emphasize the fact that the bid and valuation are the same.

In this work we assume that the auctioneer can distinguish between TF ads and non-TF ads, *i.e.*, the nature of each ad is known. Since we are aware that this is a strong assumption in practice, in Section 6 we elaborate on this point and show some empirical results where the auctioneer is uncertain about the ads nature. Additionally, our model assumes away the growing richness that is part of the ad ecosystem that is part of the motivation for a switch. However, at this point basic text ads still represent the bulk of ad impressions, so a solution that works well for them would be useful to aid a near-term transition in anticipation of this future richness.

4. TOP INTERFERENCE FREE PAYMENT FRAMEWORK

In this section we introduce a framework for designing mechanisms in ad auction-like settings where it is easier to provide a payment rule than to give an allocation rule. In other words, we do not know exactly what we want the final allocation of ads to be, but we do know what we want the payment for ad π_i assigned to position i in the overall assignment Π to be. More formally, we explore the space of payment rules which are a set of n functions $\mathcal{P} = \{p^{(i)}\}_{i \in [n]}$, where function

$$p^{(i)} : \overbrace{\mathcal{T} \times \dots \times \mathcal{T}}^{n-i} \mapsto \mathbb{R}^+.$$

That is, if the ad in slot i is TF its payment will be $p^{(i)}(\pi_{i+1}, \dots, \pi_n)$. (Recall that we allow the payments of non-TF bidders to be arbitrary.) Note that $p^{(i)}$ is the expected payment of the ad assigned to position i and $p^{(i)}/f_i$ is its cost per click.

This formulation implicitly restricts the set of payment rules we consider. Without loss of generality, the payment of a truthful ad does not depend on its own bid. However, with loss of generality, we also assume that payment does not depend on the bids of ads assigned to slots above it. This is a natural restriction in a setting without externalities, and indeed one that is satisfied by both the GSP and VCG payment rules.

Our framework specifies two further intuitive properties which we require the payment rule satisfy. We show that these two properties are necessary in the sense that any anonymous mechanism whose payment rule for an ad depends only on ads below it can be implemented using our framework. By anonymous, we mean that permuting the input to the mechanism simply permutes the output (up to tie breaking among ads with identical bids). In order to specify these properties we need to restrict the domain of the payment rules to exclude nonsensical inputs where the TF bidders are mis-ordered, e.g., where TF ads are not assigned to slots monotonically with respect to their bid.

Let $\Pi = (\pi_1, \dots, \pi_n)$ denote the assignment of ads to positions. We use the notation $\Pi^{(k)} = (\pi_k, \dots, \pi_n)$ to show the partial assignment of the last $n - k + 1$ ads to positions k to n . In the following we define what partial assignments are valid and thus form the domain of payment rule $\{p^{(i)}\}_{i \in [n]}$.

DEFINITION 4.1 (VALID ORDERING). *A partial assignment of $n - k + 1$ ads $\Pi^{(k)}$ is valid if and only if for any $i, j \in \{k, \dots, n\}$ such that π_i and π_j are TF ads and $i < j$, we have $v(\pi_i) \geq v(\pi_j)$.*

Recall that Myerson's characterization of truthful mechanisms asks a designer to give a monotone allocation rule and the payments are then uniquely derived from the area above the curve. A monotone allocation rule, when seen from the payment perspective, implies monotone marginal increases of the payments (see Figure 1).




Fig. 1: The area above the allocation curve of a winner is his payment. The two arrows show the marginal increase of the payment at different points. If the allocation curve is monotone the marginal increases are also monotone.

DEFINITION 4.2 (MARGINAL OPERATOR $\nabla^{(i,j)}$). For two positions $i, j \in [n]$ where $i < j$, the marginal increase of payment rule \mathcal{P} for a valid assignment Π is

$$\nabla^{(i,j)}\mathcal{P}(\Pi) = \frac{p^{(i)}(\Pi^{(i+1)}) - p^{(j)}(\Pi^{(j+1)})}{f_i - f_j}.$$

Now we are ready to specify the first property which the payment rule \mathcal{P} should satisfy.

DEFINITION 4.3 (EXACT MARGINAL INCREASE (EMI)). The payment rule \mathcal{P} satisfies EMI, if for any valid assignment Π and position $i \in [n - 1]$, if π_{i+1} is a TF ad then

$$\nabla^{(i,i+1)}\mathcal{P}(\Pi) = v(\pi_{i+1}).$$

The intuition behind the EMI requirement (Definition 4.3) is that, since TF ads are shown in the order of their bid, the minimum bid TF ad π_i needed to get shown above the TF ad π_{i+1} is exactly $v(\pi_{i+1})$. Thus, the marginal payment he should make for being in slot i as opposed to $i + 1$ is exactly this minimum bid.

Note that in our setting there are non-TF ads that can be placed between TF ads. Therefore, we need to generalize EMI in order to make sure that the payments for TF ads remain incentive compatible.

DEFINITION 4.4 (MINIMUM MARGINAL INCREASE (MMI)). The payment rule \mathcal{P} satisfies MMI if for any valid assignment Π , position $i \in [n]$ such that π_i is a TF ad, and position $j \in \{1, \dots, i - 1\}$ we have

$$\nabla^{(j,i)}\mathcal{P}(\Pi) \geq v(\pi_i).$$

Now we give our algorithm to derive the final allocation given payment rule \mathcal{P} which satisfies EMI and MMI. Our algorithm is very simple and intuitive. It starts filling from position n all the way up to position 1. The TF ads get assigned to positions in the increasing order of their valuations. Let us assume that the current TF ad to be assigned to a position is π . Our algorithm tries to fill all the remaining positions by non-TF ads which are not yet assigned, choosing the ads sequentially such that the payment of the next position is minimized. Then, our algorithm puts ad π in the position i for which its profit is maximized. The algorithm then takes the next TF ad and restarts from position $i - 1$. The formal description of our Allocation Algorithm (AA) is given in Algorithm 1.

Algorithm 1: Allocation Algorithm (AA)

input : n ads $\{1, \dots, n\}$ and payment rule $\{p^{(i)}\}_{i \in n}$.
output: assignment (π_1, \dots, π_n) of ads to positions and their payments.

- 1 $T \leftarrow \text{Extract-TF-ads}(\{1, \dots, n\});$
- 2 $N \leftarrow \text{Extract-NonTF-ads}(\{1, \dots, n\});$
- 3 $\ell \leftarrow n;$
- 4 **while** $|T| > 0$ **do**
- 5 Let $\pi \in T$ be a TF ad with minimum value ;
- 6 Remove π from T ;
- 7 Fill-With-NonTF-ads ;
- 8 $i \leftarrow \arg \max_{j \in \{\ell - |N|, \ell - |N| + 1, \dots, \ell\}} f_j \cdot v(\pi) - p^{(j)}(\pi_{j+1}, \dots, \pi_n);$
- 9 $\pi_i \leftarrow \pi;$
- 10 $N \leftarrow N - \{\pi_{i+1}, \dots, \pi_\ell\};$
- 11 $\ell \leftarrow i - 1;$
- 12 **end**
- 13 Fill-With-NonTF-ads ;
- 14 Set the payment of π_i to be $p^{(i)}(\pi_{i+1}, \dots, \pi_n);$

/* Sub-procedure Fill-With-NonTF-ads provisionally assigns all the remaining non-TF ads to next available positions. At each step it selects a non-TF ad which makes the next payment as small as possible. */

- 15 Fill-With-NonTF-ads:
- 16 **begin**
- 17 $N' \leftarrow N;$
- 18 **for** $i \leftarrow \ell$ **downto** $\ell - |N| - 1$ **do**
- 19 $\pi_i \leftarrow \arg \min_{\pi \in N'} p^{(i-1)}(\pi, \pi_{i+1}, \dots, \pi_n);$
- 20 $N' \leftarrow N' - \{\pi_i\}$
- 21 **end**
- 22 **end**

Description of the algorithm. Set T contains all the TF ads which are not yet assigned. Similarly set N contains all the non-TF ads which are not yet assigned. In Line 3 we initialize the value of ℓ which keeps the index of current position to be filled. In Line 5 we select a TF ad with minimum value in order to assign it to a position. In Line 7 we provisionally fill the next $|N|$ positions with non-TF ads. In Lines 8 and 9 we find and assign a position with the best profit for TF ad π . In Line 10 we remove all the non-TF ads which are assigned permanently (appear after the position i) from N . In Line 13 we fill the remaining positions by the rest of non-TF ads. Finally at Line 14 we set the payments of allocated ads according to $p^{(i)}$.

Note that at Lines 5 and 19, we might have multiple valid choices, in which case we break the ties by the choosing the ad with the smallest index. The only other instance where a tie can happen is at Line 8, when we select the largest feasible j .

The following theorem shows that mechanisms derived from our framework are incentive compatible for TF ads.

THEOREM 4.5. *Given a set of payment functions $\mathcal{P} = \{p^{(i)}\}_{i \in [n]}$ that satisfy *EMI* and *MMI*, the mechanism derived from applying AA to \mathcal{P} is incentive compatible for TF ads.*

PROOF. Let \mathcal{M} be the resulting mechanism after applying AA to set of payment functions. Observe that \mathcal{M} assigns TF ads to positions in the increasing order of their value, *i.e.*, the larger the value of a TF ad is, the higher position he receives. This follows from Line 5 of AA.

In order to prove incentive compatibility of mechanism \mathcal{M} , we show that an arbitrary TF ad gets the best utility when he bids his true valuation. Assume that θ is an arbitrary type profile, \mathcal{M} outputs assignment $\Pi = (\pi_1, \dots, \pi_n)$ for θ , and π_k is a TF ad. We show that utility of π_k does not increase if he bids v' considering three cases: (1) he is considered in the same iteration of Algorithm 1, (2) he is considered in a later iteration, and (3) he is considered in an earlier one.

Case (1): Since he is considered in the same iteration, all that changes is that Line 8 optimizes with respect to v' rather than v , giving him a weakly worse position. Thus, he does not benefit.

Case (2): Since he is considered in a later iteration, some other TF with value $v'' \geq v$ is considered in his original iteration and assigned to slot k'' . By MMI, $\nabla^{(k', k'')} \mathcal{P}(\Pi') \geq v'' \geq v$. Thus, his marginal payment for all the clicks he gets beyond what he would get in slot k'' is at least his value, and he is no better off than he would have been originally taking slot k'' , which is a contradiction.

Case (3): Without loss of generality, let π_k be the bidder in the lowest slot (according to π) who can benefit from lowering his bid. Let k'' be the highest slot below k such that $\pi_{k''}$ (with value v'') is TF. By the taxation principle, there is a price that $\pi_{k''}$ faced for every slot at or below k'' , and at those prices he preferred k'' . π_k could have faced those same prices by bidding $v'' - \epsilon$ for sufficiently small ϵ , and as $v \geq v''$ he too prefers slot k'' among all those options. By EMI, $\nabla^{(k''-1, k'')} \mathcal{P}(\Pi) = v'' \leq v$. Thus, he weakly prefers taking slot $k'' - 1$ to taking slot k'' . Since slot $k'' - 1$ was one of his options, he weakly prefers slot k to it, a contradiction. \square

Having shown that every mechanism derived from our framework is truthful, it is natural to characterize the class of mechanisms that are implementable with our framework. We show that this class is characterized by three natural axioms and one technical one.

First note that payment functions $\{p^{(i)}\}_{i \in [n]}$ only use the bid and nature of the ads and do not use the identity (index) of the ads to determine payments². This means that mechanisms derived from our framework satisfy anonymity, defined formally below.

DEFINITION 4.6. [*Anonymous Mechanism (AM)*] A mechanism $(\mathcal{M} = (x, p))$ with allocation function x and payment function p is anonymous if the following holds. Let θ and θ' be two type profiles that are permutations of each other (*i.e.* the set of natures and bids are the same but the identities of ads are permuted) and have no ties. Say $\theta = \sigma(\theta')$, for some fixed permutation σ . Then we have $x(\theta) = \sigma(x(\theta'))$ and $p(\theta) = \sigma(p(\theta'))$. For type profiles with ties we require permutations to permute the output, except that the payments and allocations of tied bidders can be exchanged arbitrarily.

Secondly, note that the payment of the ad assigned to position i is specified by looking only at the ads that are assigned to positions below i . Therefore, mechanisms in our framework also satisfy the following property.

DEFINITION 4.7 (TOP INTERFERENCE FREE (TIF)). A mechanism $\mathcal{M} = (x, p)$ satisfies TIF, if when an ad changes its type and gets a better position then the allocation

²This would appear to rule out current systems that incorporate other information such as click probability and other quality measures into a rank score. However, our results still apply in this more general setting as long as the rank scores are treated anonymously.

of ads assigned to lower positions remains unaltered. More formally, let $x(\theta)$ be the allocation given by x on type profile $\theta = \{\theta_1, \dots, \theta_n\}$ and $x(\theta')$ the allocation given by x on type profile θ' where $\theta'_h = \theta_h, \forall h \in [1, \dots, k-1, k+1, \dots, n]$ and $\theta'_k \neq \theta_k$. Assume that ad k is in position i with allocation $x(\theta)$ and in position j with allocation $x(\theta')$ such that $j < i$. Mechanism \mathcal{M} satisfies TIF if the ads assigned to positions $i+1$ to n are the same in both allocations $x(\theta)$ and $x(\theta')$.

In the following theorem, we prove that our framework can implement all mechanisms that are incentive compatible, anonymous, and top interference free, as well as satisfying an additional technical axiom (one which seems to be satisfied for reasonable mechanisms). Hence, requiring EMI and MMI does not restrict the designer in ways that current standard designs such as VCG and GSP do not.

THEOREM 4.8. *A mechanism is derived from our framework if and only if it satisfies IC, AM, TIF, and 2T.*

See Appendix A for a definition of the technical axiom (2T), a proof of the theorem, and a discussion of 2T.

5. PRICING FUNCTIONS

In the preceding section we designed a general framework. In this section, we apply it to the desired special case of transitioning from a GSP auction to a VCG auction. To do so, we must decide how to price GSP bidders, since our framework is silent about how they should be charged. We make perhaps the simplest decision, to charge them the same amount a truthful bidder would be charged in the same slot, and show that this has several desirable properties.

We begin by discussing what happens when users can evolve and switch types: that is, they are either TF or GSP bidders, and in each case act as utility maximizers. When the hybrid mechanism is first put into use, all bidders are GSP bidders, but as time goes on some bidders will change to being truthful. Given perfect rationality, this means i 's bid changes from $b(i)$ to $v(i)$. We can then show that if GSP bidders begin from the *lowest revenue* Symmetric Nash Equilibrium (SNE) then the revenue and allocations are unaltered, provided a particular hybrid pricing function is employed, irrespective of the order in which users transition from GSP to TF bids. (As SNE rank bidders in decreasing order of bid, in this section we assume that ad i is in slot i .) Formally,

PROPOSITION 5.1. *For any GSP algorithm where GSP bidders bid as in the lowest revenue SNE and truthful bidders bid truthfully, the revenue, allocation and prices paid will be independent of the number and identity of TF and GSP bidders if and only if the payment function satisfies MMI and EMI and further satisfies*

$$p^{(i-1)}(\Pi^{(i)}) = \begin{cases} p^{(i)}(\Pi^{(i+1)}) + v(i)(f_{i-1} - f_i) & \text{if } \theta_i = (TF, v(i)) \\ b(i)f_{i-1} & \text{if } \theta_i = (GSP, b(i)) \end{cases} \quad (1)$$

PROOF. Both directions of the proof follow almost directly from the definitions of a lowest SNE, EMI and MMI. In order for payments of TF and GSP bidders to be identical, the payment functions $p^{(i)}(\Pi^{(i+1)})$ must be the same regardless of whether i is TF or GSP, and independent of the mix of bidder types in $\Pi^{(i+1)}$. Consequently no GSP or TF bidder wants to change bid or position, since by the definition of an SNE

$$\left(v(i) - p^{(i)}(\Pi^{(i+1)})\right) f_i \geq \left(v(i) - p^{(j)}(\Pi^{(j+1)})\right) f_j \quad \text{for all } i, j. \quad (2)$$

By standard arguments about the lowest SNE (see, e.g., [Varian 2007; Roberts et al. 2013]) we in fact have that for all i ,

$$b(i) = b(i+1) + v(i)(f_{i-1} - f_i). \quad (3)$$

Thus, by induction, the two conditions of (1) are in fact equal at the lowest SNE. This gives that the form is necessary and sufficient for prices to coincide, as pricing must be equivalent to the case $\theta_j = (GSP, b(j))$ for all $j \geq i$. As this outcome is equivalent to the outcome of a truthful auction, it follows that EMI and MMI are satisfied as well. \square

In the statement of the proposition, we say that it applies to any GSP algorithm. By this we mean that GSP is, strictly speaking, just a payment rule. It can be applied to a variety of rank score allocation rules. As long as the one chosen admits an SNE, the proposition applies. The following corollary results from applying a sufficient condition [Roberts et al. 2013] for this.

COROLLARY 5.2. *The result holds true for any GSP ranking function that uses a rank score of the form*

$$y(b, i) = (g(i)b - h(i))^+, \quad (4)$$

where g and h are arbitrary non-negative values that can depend on i .

The necessary and sufficient conditions only hold when we start from a *lowest* SNE. For example if we are in another SNE, then moving just one bidder i from GSP to TF will not change the position or prices paid by i or those below i , but potentially changes prices (and hence positions) of bidder(s) above i (since equality in (3) need no longer hold for $i-1$). Hence, we want to construct price functions that satisfy EMI and MMI when other equilibria hold, and for more general non-truthful prices. Specifically, we shall consider two natural examples, where the pricing functions for ad i are the same for truthful and non-truthful i . First,

$$\mathbf{A} : p^{(i-1)}(\Pi^{(i)}) = \max \left(p^{(i)}(\Pi^{(i+1)}) + v_{max}(\Pi^{(i)})(f_{i-1} - f_i), b_{max}(\Pi^{(i)})f_{i-1} \right) \quad (5)$$

where

$$v_{max}(\Pi^{(i)}) \stackrel{def}{=} \max \{v(\theta_j) : j \geq i \cap \theta_j = (TF, v(j))\} \quad (6)$$

$$b_{max}(\Pi^{(i)}) \stackrel{def}{=} \max \{b(\theta_j) : j \geq i \cap \theta_j = (GSP, b(j))\} \quad (7)$$

are the largest TF valuation and GSP bid at or below i , respectively, and

$$\mathbf{B} : p^{(i-1)}(\Pi^{(i)}) = \max \left(p^{(\arg \max v(\Pi^{(i)}))}(\Pi^{(\arg \max v(\Pi^{(i)})+1)}) + v_{max}(\Pi^{(i)})(f_{i-1} - f_{\arg \max v(\Pi^{(i)})}), b_{max}(\Pi^{(i)})f_{i-1} \right) \quad (8)$$

where

$$\arg \max v(\Pi^{(i)}) \stackrel{def}{=} \arg \max_j \{v(j) : j \geq i \cap \theta_j = (TF, v(j))\} \quad (9)$$

is the identity of the largest TF ad at or below i .

Either of these hybrid auctions is consistent with Proposition 5.1, and so in fact they are identical in this case. It is easy to see, however, that they do differ in other scenarios. In some sense, we can think of these as the two extremes of reasonable prices. In any “reasonable” extension of GSP, an advertiser ought to pay at least the bid of a GSP bidder below him, and \mathbf{B} is the lowest set of prices consistent with this, EMI, and MMI. At the other extreme, \mathbf{A} charges the highest prices that are consistent with no GSP bidder paying more than his bid.

OBSERVATION 1. *When setting prices according to (5), GSP bidders pay exactly their bid when an indifferent TF bidder is put below them.*

PROOF. Let b be the GSP bidder assigned to slot i and v be the value of the truthful bidder assigned to slot $i + 1$ and indifferent between that and slot i . Then

$$f_i(v - b) = f_{i+1}v - p^{(i+1)}(\Pi^{(i+1)}).$$

Rewriting shows that the GSP bidder pays exactly his bid. \square

Since the main concern with switching to VCG is the loss of revenue, we use **A** in our simulations. However, this does come at a cost in welfare relative to **B**, since it will tend to put higher TF bidders below lower GSP bidders more often.

6. SIMULATION RESULTS

We saw in Proposition 5.1 that if bidders always play the lowest SNE, we can perfectly identify which bidders have adjusted to the new truthful auction, and that adjustment consists of instantly switching to the advertiser’s true value, then there would be no effect on efficiency or revenue from switching to the hybrid auction. Of course, none of these assumptions are realistic. In this section, we discuss a variety of simulations that analyze the practical effects of the hybrid auction in more realistic scenarios.

6.1. Simulation Setup

We base our simulations on a *non-random* sample of Bing data on 3984 auctions. It is a filtered subset of a larger random sample that ensures the auctions are “interesting.” In particular, we wanted thick auctions (with at least 12 participants), and with other properties such that techniques for inferring true values from GSP bids could give reasonable answers. The metrics have been normalized. Nevertheless, we believe the sample is representative enough to allow a meaningful exploration of our approach.

We restrict each auction to the top 12 participants, and only actually run an auction for the top three slots. In order to run our simulations we need to have an estimate of true valuation of GSP ads. One estimate of true valuations is to assume that GSP ads have played the minimum symmetric Nash equilibrium and invert their bids to their valuations. In this case, we would essentially be baking in the first of the assumptions from Proposition 5.1, so unsurprisingly the transition would happen without any changes as the allocation and payment of ads remain identical at each point of time.

Instead, we use the stochastic formulation from Pin and Key [2011]. This approach derives the valuations under the hypothesis that each advertiser chooses her bid to maximize her expected net utility under the assumption that she faces a stationary bid distribution. In our calculations we assume that the CTRs are known, with the opposing bid distribution estimated from the opponents’ empirical bid distributions.

In our simulations we run four different mechanisms.

- **GSP.** The first mechanism is GSP run on the original set of bids when no updates have happened. This represents the current state of the world and serves as a benchmark to which the other approaches can be compared.
- **VCG-V.** The second mechanism is VCG run on the final set of true valuations when all the ads have updated their bids. This represents the ideal end state when all bidders have transitioned to being truthful. It also serves as a sanity check on the reasonableness of our value estimation (i.e. it should display similar performance to GSP).
- **HYBRID.** The third mechanism is the one derived from our framework, using pricing rule **A** described in (5).

- **VCG-B.** Finally the last mechanism is VCG run on the current set of bids when some ads have updated their bids and some have not. This is the obvious alternative strategy for transitioning: simply transition directly to VCG and wait for bidders to catch up.

6.2. Perfectly Rational Bidders

The first set of simulations we run assumes that bidders are perfectly rational and that they know that any bid change will result in them being classified as truthful. Such a bidder would directly update his bid to his true valuation without bidding any other intermediate amounts since the mechanism is incentive compatible. In this simulation we assume that at each time step one randomly selected ad decides to change its bid from the GSP value to the true value.

Figure 2 shows the normalized average revenue, welfare, and click yield³ for different mechanisms during the transition. The estimated revenue from ultimately running VCG (i.e. VCG-V) is close to GSP, which is consistent with the reasonableness of our value estimation procedure. Immediately switching to VCG (curve VCG-B) results in a significant revenue drop, which is steadily recovered as more advertisers update their bids. In contrast, there is a more modest revenue drop under the hybrid mechanism (since bidders are not always following the lowest SNE). In particular, revenue always dominates directly switching, substantially so in the initial time steps. In both the estimated welfare and click yield there are no significant differences between VCG-B and Hybrid auction. The observation that welfare and click yield do not differ much in the Hybrid auction and in the VCG-B strengthens the importance of the revenue improvement that the former has over the latter because it is not coming at the cost of other important metrics. Note also that, in the worst case, the drop in welfare is less than 1.5% and the drop in click yield is less than 0.3% from the optimal case. Thus, we focus on revenue in our subsequent simulations.

6.3. Cautious Bidders

The second set of simulations relaxes the idea that bidders are willing to immediately jump to their true value, no matter how large a bid increase this implies. Instead we parameterize them with a triple (p, q, i) . At each time step, bidders decide randomly whether to update their bid, doing so with probability p . If their consumer surplus decreased in the last step (i.e. because the bid changes of others changed their slot or increased their price) they update with a higher probability q . This allows us to model advertisers who are attentive only when needed. Finally, when they update they increase their bid by a percentage i until they reach their true value. Bidders are treated as truthful as soon as they change their bid.

Different parameterizations lead to somewhat different pictures, but all share the same general trends as in Figure 2. Note however, that relative to that figure their x-axis has been compressed, since it now takes significantly more than 12 rounds for all bidders to fully adapt. Figure 3, with parameters $(0.3, 0.6, 0.1)$ shows that these cautious updates hurt the performance of the hybrid relative to a direct switch to VCG. There is still a benefit for the first 8 rounds, but then essentially all bidders are classified as truthful, so performance is the same as if we had switched directly. This points to the need for more subtle methods of determining how to treat bidders than just on the basis of whether or not they have updated their bids since the process began. Figure 4 shows that the benefits persist longer if we have bidders who are lax about updating (unless something bad happens) with parameters $(0.1, 0.9, 0.1)$. Larger values

³Here click yield includes both bad clicks, the ones users do not stay in the clicked website for enough time, and good clicks.

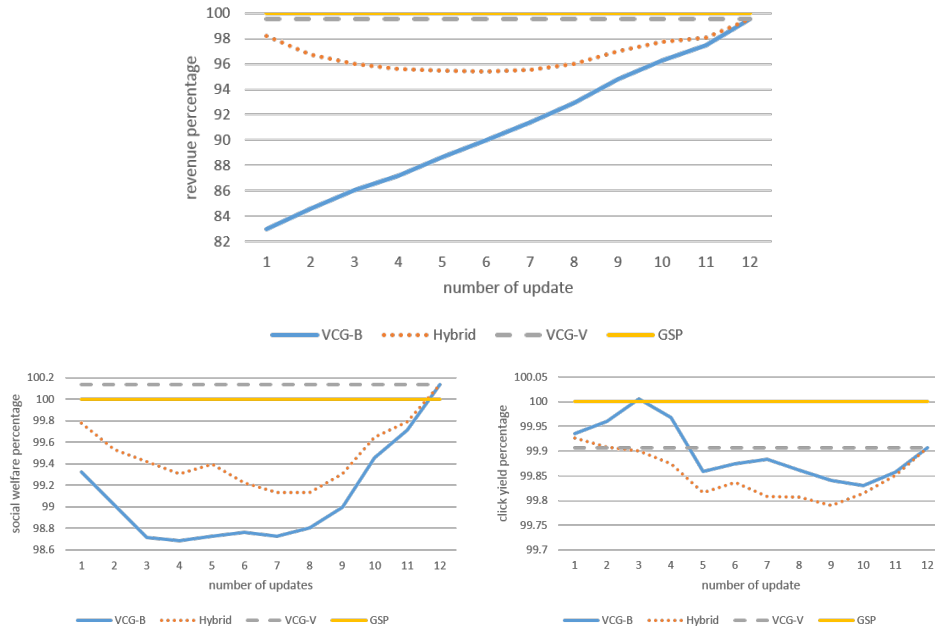


Fig. 2: Perfectly Rational Bidders

of i (not shown) lead to more of the benefits of the hybrid approach being maintained, since the period when a bidder is not GSP but not yet truthful is shortened.




Fig. 3: Cautious Bidders




Fig. 4: Cautious Bidders, Lax Updates

6.4. Identification Strategies

In the preceding simulations, bidders are homogeneous. Thus, in some sense trying to identify which bidders are truthful is meaningless: bidders differ only in the number of times they have increased their bid. In this simulation, we consider instead a heterogeneous population of bidders. Some update more frequently and in smaller increments, representing advertisers who use automated tools to optimize their campaigns, while others update only occasionally (and in practice there are many advertisers who go

long periods of time between updates). Intuitively, we can likely tolerate simply treating the former as truthful immediately since they will rapidly adapt to the new setting, while we can update the latter as soon as they first change their bid, since that first change represents a large step towards their true value. Advertisers in between are more problematic.

Here we model each bidder's behavior by an update factor x where $x \in [0..1]$ is a real number. The bidder with update factor x updates her bid with probability of x and increases her bid by factor $1 + \frac{b}{x}$ where base increment percentage b is the parameter of the simulation. Note that the larger x is the less is the increment factor. At the start we assign the update factors to bidders by selecting a random number from $[0..1]$.

Our identification strategy is as follows. We update a bidder to a TF bidder with probability $IncPer/BasePer$ where $IncPer$ is the total percentage of increases that the bidder has made since the start of the transition and $BasePer$ is a parameter of the simulation. Moreover we update a bidder to a TF bidder with probability $IncNum/BaseNum$ where $IncNum$ is the total number of increases that the bidder has made since the start of the transition and $BaseNum$ is a parameter of the simulation. Therefore, we describe our simulations by a triple $(b, BasePer, BaseNum)$. Note that the smaller $BasePer$ or $BaseNum$ are, the more likely is to tag an ad as a TF bidder.

Figure 5 shows the performance of the hybrid auction in three different simulations: $(0.2, 0.1, 5)$, $(0.2, 0.2, 10)$, and $(0.2, 0.3, 30)$. In all the simulations the base increment percentage b is 0.2. In simulation $(0.2, 0.1, 5)$ values $BasePer$ and $BaseNum$ are the smallest among others, hence, we tend to tag a bidder as a TF ad more likely. As one expects, this reduces the overall revenue as it is more likely that a bidder who is not fully adapted (does not bid his true valuation) get tagged as a TF bidder. In simulation $(0.2, 0.3, 30)$ values $BasePer$ and $BaseNum$ are the largest among others, hence, we are less likely to tag a bidder as a TF ad and hence get more revenue. Although these simulations suggest that we should less likely tag bidders as TF, but we note that in order to give incentive to bidders to update at all, we cannot decrease this probability too much. Thus, perhaps the biggest message is that there is a tension between preserving revenue and giving bidders an incentive to update their bids.

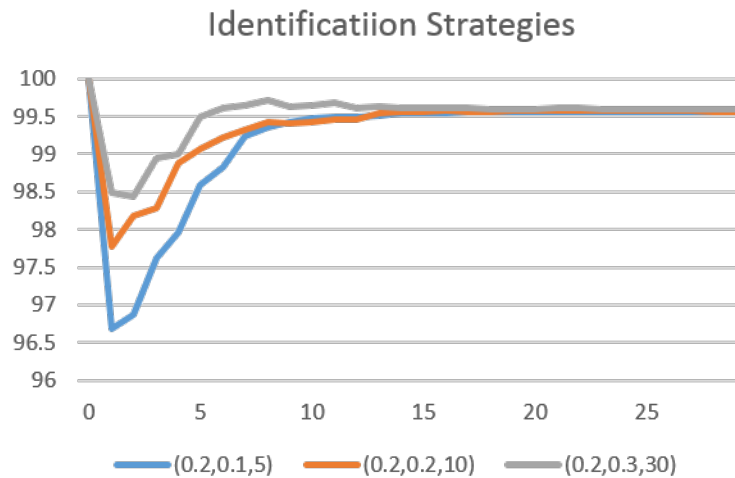


Fig. 5: Identification Strategies

7. CONCLUDING REMARKS

We have described a framework for a hybrid mechanism design where bidders may be either truthful or not, and compete for goods where there is an intrinsic ordering by worth amongst the goods. If we can identify each group in this dichotomy, then we have shown that the fundamental building block is a payment function, which maps the bids for lower value goods to the payment for a good, and which needs to satisfy two properties. The two properties, EMI and MMI relate solely to the bids of truthful agents, and place constraints on the discrete derivative of the payment function; the payment for non-TF bidders may be arbitrary.

We have given details of a bottom-up allocation procedure, which when used with an EMI and MMI payment function gives an Incentive Compatible (IC) mechanism, and hence gives incentive for bidders to change type to truthful. If in addition, the mechanism is Top Interference Free, TIF, then this characterization is both necessary and sufficient for IC anonymous mechanisms which satisfy an additional technical axiom. Any mechanism derived from a bottom-up procedure, such as standard GSP or VCG, are all examples of TIF mechanisms.

The motivation for our work is sponsored search, where ad-slots are auctioned off. But as the previous paragraph suggests, our framework does not have to be tied to this setting. Our particular starting point was wanting to provide a pathway for migrating from one non-truthful mechanism (GSP) to a truthful mechanism whilst mitigating the revenue loss that would occur if there was a switch to a truthful mechanism but bidders did not immediately update (increase) their bids to their true valuation. We have shown empirically that revenue loss can indeed be mitigated if users do switch between being non-truthful with a GSP bid or truthful and bid their true valuation, but that the situation is more nuanced if users update partially (i.e. are not fully truthful).

The latter illustrates the difficulties involved in ascertaining a users type, and points to the need for further research into the best way to classify users, and the trade-offs involved in mis-classification. The example pricing functions we gave appear to work well for vanilla sponsored-search auctions, and this was our intended target. For use in richer settings (e.g. where there are richer ad variants), the pricing functions can be straightforwardly adapted for bottom-up or TIF settings, but would need additional work to adapt them to non-TIF settings.

REFERENCES

- AGGARWAL, G., FELDMAN, J., AND MUTHUKRISHNAN, S. 2007. Bidding to the top: Vcg and equilibria of position-based auctions. In *Approximation and Online Algorithms*. Springer, 15–28.
- AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*. ACM, 1–7.
- AGGARWAL, G., MUTHUKRISHNAN, S., PÁL, D., AND PÁL, M. 2009. General auction mechanism for search advertising. In *Proceedings of the 18th international conference on World wide web (WWW)*. ACM, 241–250.
- ARCHER, A. AND KLEINBERG, R. 2008. Truthful germs are contagious: a local to global characterization of truthfulness. In *Proceedings of the 9th ACM Conference on Electronic Commerce*. 21–30.
- ASHLAGI, I., BRAVERMAN, M., HASSIDIM, A., AND MONDERER, D. 2010. Monotonicity and implementability. *Econometrica* 78, 5, 1749–1772.
- ASHLAGI, I., MONDERER, D., AND TENNENHOLTZ, M. 2007. Mediators in position auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 279–287.

- ATHEY, S. AND NEKIPELOV, D. 2010. A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop*.
- BACHRACH, Y., CEPPI, S., KASH, I. A., KEY, P., AND KUROKAWA, D. 2014. Optimising trade-offs among stakeholders in ad auctions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 75–92.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., AND KYROPOULOU, M. 2011. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 81–90.
- CARY, M., DAS, A., EDELMAN, B., GIOTIS, I., HEIMERL, K., KARLIN, A. R., MATHIEU, C., AND SCHWARZ, M. 2007. Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 262–271.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2005. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords.
- EDELMAN, B. AND SCHWARZ, M. 2010. Optimal auction design and equilibrium selection in sponsored search auctions. *The American Economic Review*, 597–602.
- FRONGILLO, R. M. AND AND, I. A. K. 2014. General truthfulness characterizations via convex analysis. In *10th Conference on Web and Internet Economics (WINE)*. To Appear.
- GOMES, R. AND SWEENEY, K. 2014. Bayes–nash equilibria of the generalized second-price auction. *Games and Economic Behavior* 86, 421–437.
- LEME, R. P. AND TARDOS, É. 2009. Sponsored search equilibria for conservative bidders. In *Fifth Workshop on Ad Auctions*.
- LUCIER, B., PAES LEME, R., AND TARDOS, E. 2012. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 361–370.
- MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 58–73.
- PIN, F. AND KEY, P. 2011. Stochastic variability in sponsored search auctions: observations and models. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 61–70.
- ROBERTS, B., GUNAWARDENA, D., KASH, I. A., AND KEY, P. 2013. Ranking and tradeoffs in sponsored search auctions. In *Proc. of the Fourteenth ACM Conf. on Electronic Commerce*. EC '13. 751–766.
- THOMPSON, D. R. AND LEYTON-BROWN, K. 2013. Revenue optimization in the generalized second-price auction. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. ACM, New York, NY, USA, 837–852.
- VARIAN, H. R. 2007. Position auctions. *international Journal of industrial Organization* 25, 6, 1163–1178.
- VARIAN, H. R. AND HARRIS, C. 2014. The VCG auction in theory and practice. *The American Economic Review* 104, 5, 442–445.
- VOROBAYCHIK, Y. 2009. Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 583–590.

A. PROOF OF THEOREM 4.8

Before we begin, we observe that our requirements ensure that the mechanism allocates truthful ads in increasing order of value.

OBSERVATION 2. *If \mathcal{M} satisfies IC, AM, and TIF then it allocates TF ads in increasing order.*

PROOF. Suppose for contradiction that $v_1 < v_2$ but ad (TF, v_1) got a better slot than ad (TF, v_2) . By IC, if we replace (TF, v_1) by another copy of (TF, v_2) , its slot can only improve. By TIF, this means that the slot of the original copy is unchanged. But by AM, if the original ads had been permuted the ad with value v_1 could raise its bid to v_2 and receive a worse slot, contradicting IC. \square

We begin with sufficiency. Let $\mathcal{M}' = (x', p')$ be a mechanism with allocation function x' and payment function p' which satisfies IC, AM, TIF, and 2T (see below). We use \mathcal{M}' to propose payment rule \mathcal{P} such that mechanism \mathcal{M} derived from applying \mathcal{P} to TIFP framework is equivalent to \mathcal{M}' . In order to deal with technicalities of ties, we assume that in the case of ties \mathcal{M}' use the same tie breaking rule as our framework.

Let $\Pi^{(k)} = (\pi_k, \dots, \pi_n)$ for $k \in \{3, \dots, n\}$ be the assignment of \mathcal{M}' for positions k to n . We say that ad i is less than j with regard to $\Pi^{(k)}$ (show by $i \prec_{\Pi^{(k)}} j$), if there exist a type profile θ such that

- The allocation $x'(\theta)$ is the same as $\Pi^{(k)}$ for positions k to n .
- $i, j \in \theta$.
- Ad i is assigned to position $k - 1$ and ad j is assigned to a position better than i ($x'(i) > x'(j)$).

Intuitively, $i \prec_{\Pi^{(k)}} j$ means that fixing $\Pi^{(k)}$ allocation x' prefers to assign i to position k over j . In the following lemma we prove that $\prec_{\Pi^{(k)}}$ is a total order over all the ads which can be assigned to position $k - 1$ fixing $\Pi^{(k)}$. It turns out that TIF is almost, but not quite strong enough to prove this lemma. The difficulty is that it has no “bite” when applied to the case of $k = 3$ (i.e. the final 2 slots). That is, ad in slot 2 has no (non-fixed) ads below it (so the definition is vacuous), while an ad in slot 1 that changes type and remains in slot 1 forces the other ad to stay in slot 2 (again making the definition vacuous). Thus, we need a property that ensures the mechanism is well-behaved in this case.

DEFINITION A.1 (TWO TRANSITIVE(2T)). \mathcal{M} is two transitive if for all choices of $\Pi^{(3)}$ the relation $\prec_{\Pi^{(3)}}$ is transitive.

LEMMA A.2. The relation $\prec_{\Pi^{(k)}}$ is antisymmetric (if $i \prec_{\Pi^{(k)}} j$ then $j \not\prec_{\Pi^{(k)}} i$), total, and transitive.

PROOF. We prove the lemma by contradiction. To show antisymmetry, let θ be a type profile for which $i \prec_{\Pi^{(k)}} j$ and θ' be a type profile for which $j \prec_{\Pi^{(k)}} i$. Consider a sequence of type profiles that are intermediate between θ' and θ in the sense that the transition from one profile to the next results from changing the type of a single ad from its value in θ' to its value in θ . Let $\theta' = \theta_1, \theta_2, \dots, \theta_{a-1}, \theta_a = \theta$ be the sequence of type profiles.

We show that this sequence maintains the following invariant: the ad in slot $k - 1$ is not i and has already changed its type. Clearly this is true for θ' , since j is in slot $k - 1$ in $x(\theta')$ and its type will never change. Suppose it is true for θ_b , and let ad e be the ad that changes type between θ_b and θ_{b+1} . By our invariant, e is not in slot $k - 1$ of $x(\theta_b)$. Thus, by TIF, if the ad in slot $k - 1$ changes from $x(\theta_b)$ to $x(\theta_{b+1})$ it must be that in $x(\theta_{b+1})$ the ad in slot $k - 1$ is in fact e , which is not i and has already changed its type. This shows that i is not in slot $k - 1$ of $x(\theta)$, contradicting our assumption.

To see that the relation is total, take some θ and i and j . Create θ' by replacing all ads other than i and j shown in a slot above k in $x(\theta)$ with a copy of either i or j . Thus, by an inductive argument that shows this does not change the allocation below slot $k - 1$, some copy of i or j must be allocated to slot $k - 1$.

Finally, transitivity follows via a similar construction. If $k > 3$, simply transform θ to θ' by replacing all ads above slot k with copies of one of the relevant i, j , or ℓ . If $k = 3$, transitivity is by assumption (i.e. 2T). \square

Now we are ready to specify how we build set of payment functions $\mathcal{P} = \{p^{(i)}\}_{i \in [n]}$ from \mathcal{M}' . Let $\Pi^{(k)}$ be a valid assignment of ads to positions k to n . By Lemma A.2 we have a total ordering of candidate ads for position $k - 1$, so set $p^{(k-1)}(\Pi^{(k)})$ to be the infimum of TF ads among those candidates.

Now we prove that the mechanism \mathcal{M} derived from our framework using payment rule \mathcal{P} always gives the same allocation and payments as \mathcal{M}' which finishes the proof of this sufficiency.

LEMMA A.3. *Mechanism $\mathcal{M} = (x, p)$ is the same as $\mathcal{M}' = (x', p')$.*

PROOF. We need only verify that \mathcal{M} always gives the same allocation as \mathcal{M}' as the fact that the payments are the same (at least up to a constant) then follows via revenue equivalence. (The constant can be matched by changing p to p'' to include this constant shift.)

Now we prove by contradiction that the allocation functions x and x' are the same. Let θ be a type profile for which $x(\theta) \neq x(\theta')$. Let $\Pi^{(k)} = (\pi_k, \dots, \pi_n)$ be the largest common suffix of $x(\theta)$ and $x(\theta')$ and assume for now that $p^{(k-1)}(\Pi^{(k)})$ is finite. Let e be the ad assigned to position $k - 1$ in $x(\theta)$ and e' be the ad assigned to position $k - 1$ in $x'(\theta)$. Note that

$$e' \prec_{\Pi^{(k)}} e \tag{10}$$

since $x'(\theta)$ assigned positions $k - 1$ to e' .

Because $x(\theta)$ assigns position $k - 1$ to e as opposed to e' this means that

$$p^{(k-2)}(e, \pi_k, \dots, \pi_n) < p^{(k-2)}(e', \pi_k, \dots, \pi_n)$$

(Recall Lines 5 and 19 of algorithm AA). This means that there exists a TF ad x with valuation $\nabla^{(k-2, k-1)} \mathcal{P}((e, \pi_k, \dots, \pi_n)) < v(x) < \nabla^{(k-2, k-1)} \mathcal{P}((e', \pi_k, \dots, \pi_n))$. Now if we replace the rest of ads with TF bidders with valuation $v(x)$ then they appear before e but after e' . This contradicts with Equation 10 and the fact that $\prec_{\Pi^{(k)}}$ is a total order for any $\Pi^{(k)}$.

Now we deal with the case where $p^{(k-1)}(\Pi^{(k)})$ is infinite. Intuitively, this is the case where only non-truthful ads can be shown before the suffix $\Pi^{(k)}$. Since prices are all infinite, we need a way for the algorithm to match the order that \mathcal{M}' chooses. We do this by allowing prices of the form (∞, a) , where a is the type of a non-TF ad. The total order $\prec_{\Pi^{(k)}}$ then gives a well defined notion of the lowest price as the one whose a is lowest according to that ordering. With this enlarged set of prices, the proof proceeds as before. \square

Finally, the necessary part is easy to prove. Let \mathcal{M} be a mechanism derived from using TIFP framework. The AM and TIF properties follow by the fact that in algorithm AA (Figure 1) when assigning the next TF ad, AA uses only its value and neither its index nor the value of higher TF ads. The IC property of \mathcal{M} is the result of Theorem 4.5. 2T follows from the greedy nature of the allocation.

A.1. Discussion of Two Transitivity (2T)

Two transitivity is a technical assumption. The intuition is that we require the mechanism to be well-behaved when considering the top two slots, which TIF is not strong enough to enforce. If all non-TF ads are of the same nature, a sufficient condition is

that \mathcal{M}' is monotone for non-TF ads (MN). That is, if a non-TF ad raises its bid, it gets a (weakly) better slot.

LEMMA A.4. *IF \mathcal{M} satisfies IC, AM, TIF, and MN and all non-TF ads have the same nature then it satisfies 2T*

PROOF. By AM and IC/MN, $\prec_{\Pi(3)}$ is transitive if all 3 ads are TF or non-TF respectively. Thus, WLOG let 2 be TF and 1 be non-TF (replace IC by MN below if only 1 is TF). There are 3 cases.

Case 1: $N \prec V_1 \prec V_2$. By IC, $N \prec V_2$ (otherwise a truthful ad V_1 could raise its bid and go from slot 1 to slot 2).

Case 2: $V_1 \prec N \prec V_2$. By AM+IC, $V_1 \prec V_2$ (otherwise a truthful ad V_2 could raise its bid and go from slot 1 to slot 2 when facing N).

Case 3: $V_1 \prec V_2 \prec N$. By IC, $V_1 \prec N$ (otherwise a truthful ad V_1 could raise its bid and go from slot 1 to slot 2). \square

Such a nice sufficient condition is not obvious if there is more than 1 nature of non-TF ad. Non-degenerate examples still appear to satisfy 2T, but we do not know of a less technical way to explain the way in which they are non-degenerate. To see why, consider an example with 2 slots. If the bidders have at least one truthful ad, this becomes a form of second price auction. However, when there are two non-truthful ads of different natures, nothing obviously constrains the rule for determining the order in which they are shown in a way that corresponds to enforcing transitivity. This same example shows why we require MN in Lemma A.4. Without it we would be equally at a loss in this situation.