

Selective Call Out and Real Time Bidding

Tanmoy Chakraborty * Eyal Even-Dar † Sudipto Guha ‡ Yishay Mansour §
S. Muthukrishnan ¶

Abstract

Ads on the Internet are increasingly sold via ad exchanges such as RightMedia, AdECN and Doubleclick Ad Exchange. These exchanges allow real-time bidding, that is, each time the publisher contacts the exchange, the exchange “calls out” to solicit bids from ad networks. This aspect of soliciting bids introduces a novel aspect, in contrast to existing literature.

This suggests developing a joint optimization framework which optimizes over the allocation and well as solicitation. We model this selective call out as an online recurrent Bayesian decision framework with bandwidth type constraints. We obtain natural algorithms with bounded performance guarantees for several natural optimization criteria. We show that these results hold under different call out constraint models, and different arrival processes. Interestingly, the paper shows that under MHR assumptions, the expected revenue of generalized second price auction with reserve is constant factor of the expected welfare. Also the analysis herein allow us prove adaptivity gap type results for the adwords problem.

1 Introduction

A dominant form of advertising on the Internet involves *display* ads; these are images, videos and other ad forms that are shown on a web page when viewers navigate to it. Each such showing is called an *impression*. Increasingly, display ads are being sold through exchanges such as RightMedia, AdECN and DoubleClick Ad Exchange. On the arrival of an impression, the exchange solicits bids and runs an auction on that particular impression. This allows *real time bidding* where ad networks can determine their bids for each impression individually in real time (for an example, see [24]), and more importantly where the creative (advertisement) can be potentially produced on-the-fly to achieve better targeting [22].

This potential targeting comes hand in hand with several challenges. The Exchange and the networks face a mismatch in infrastructure and capacities and objectives. From an infrastructure

*Part of this work was done while visiting Google Research. Department of Computer and Information Science. University of Pennsylvania, Philadelphia, PA. Email tanmoy@cis.upenn.edu

†Google Research, 76 Ninth Ave, New York, NY. Email: evendar@google.com

‡Part of this work was done while visiting Google Research. Department of Computer and Information Science. University of Pennsylvania, Philadelphia, PA. Email sudipto@cis.upenn.edu

§Google Israel and The Blavatnik School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel, Email mansour.yishay@gmail.com

¶Google Research, 76 Ninth Ave, New York, NY. Email: muthu@google.com

standpoint, the volume of impressions that come to the exchange is very large compared to a smaller ad network limited in servers, bandwidths, geographic location preferences. This implies a bound on the number of auctions the network can participate in effectively. A network would prefer to be solicited only on impressions which are of interest to it, and in practice use a descriptive language to specify features of impressions (say, only impressions from NY). However this is an offline feature and runs counter to the attractiveness of real time bidding. Therefore the exchange has to “call out” to the networks selectively, simultaneously trying to balance the objective of soliciting as many networks as possible and increasing total value, as well as not creating congestion or situations where solicitations are not answered.

This leads to a host of interesting questions in *developing a joint optimization framework that optimizes over the allocation objective as well as the decisions to solicit the bids*. Specifically, a participant would be solicited for only a predetermined fraction of impressions. Moreover, these solicitations, referred to as “call outs” henceforth, need to be performed in a smooth manner and avoid burstiness. The impressions need to be managed in an online manner, which suggests the use of online algorithms. The burstiness properties suggest using a queueing model. And the overall goal is to optimize objectives such as (expected) welfare, revenue. While each of these issues have been considered in isolation, the overall challenge is to develop a joint framework, which in turn raises interesting questions about the interactions between different parts of the framework.

The call out framework is formally modeled in Section 1.1. The online allocation aspect, with a view that the call out constraints acts as budgets, is reminiscent of the online ad allocation framework for search ads, or the Adwords problem [19, 5, 7], and its stochastic variants [9, 27]. However the call out framework is significantly different, which we discuss below.

The Adwords problem is posed in the deterministic setting where the expected revenue is treated as a known deterministic reward of allocating an impression j to an advertiser i . The call out framework has no deterministic analogue; the rationale of the exchange is that the bids are not known. If the bids were known (or internal) then we would only call out the winners (assuming multiple slots) of the auction, which is the path taken by the Adwords problem. The call-out framework is similar to Bayesian mechanism design [21]. This has some fairly broad conceptual implications.

First, is the notion of “adaptivity gap”, where a policy is allowed to react to realization of the random variables. The analysis of adaptivity gap is the central question in the exchange setting. This is also relevant in the context of search ads and the adwords setting where the revenue is achieved on a click which is a random event. The adwords model uses the deterministic expectation but it is reasonable to allow an algorithm to adapt to this event (consider low click through rates and large bids, such that a payout affects the budget substantially). To the best of our knowledge, no analysis of adaptivity gap exists for the adwords problem but such a result will follow from our analysis. In the call out setting, when optimizing for welfare or total value, the assignment occurs *after* the bids are obtained, which has considerable gap in comparison to the assignment that assigns *before* obtaining the realizations (reduces to the expectations).

Second, many objective functions such as generalized second price with reserve (henceforth GSP-Reserve), for one or multiple slots, have a very different behavior in the Bayesian and deterministic settings. The gap between assignment *after* and *before* the realizations is more stark in this context – consider running Myerson’s (or similar) mechanism on the expected bids instead of the distributions.

Note that in GSP-Reserve we announce an uniform reserve price, *before* the bids are solicited as in [21]. For *known* deterministic bids, reserve prices can be made equal to the bid, and are not useful. Strong lower bounds hold for GSP without reserves [1].

Third, the notion of a comparison class in case of call out optimization framework requires more care. In the setting of these large exchanges, a comparison class with full foreknowledge of all information (in particular, the realization of the bids) is unrealistic. Moreover, the realizations of the bids depends on the networks which are called out, and two different strategies that call out to two different subsets will have completely different information. Thus to compare two algorithms, it appears that we should compare their expected outcome – *but each algorithm is allowed to be adaptive*. Thus a combination of stochastic and online models are in order in this setting.

This combination of stochastic and online models relates the call out framework to the stochastic variants of the Adwords problem [9, 27]. But while the similarity implies that Lagrangian decoupling techniques for separable convex optimization pioneered by Rockafellar [25] apply, *the different possible objectives of the call-out framework are not convex*. In fact, in the case of optimizing revenue in GSP (with reserve) or in posted price mechanisms, the objective is not submodular for all prices (as in welfare maximization). Submodular maximization with linear constraints has been studied, and while good approximation algorithms exist [6, 17], they are inherently offline – the key aspect of call out optimization is that the decision has to be made in an online fashion. The same is true for sequential posted price mechanisms analyzed in [8, 4] (albeit with more general matroid setting), the posted prices in the call out setting need to be announced in parallel (and the eventual allocation is sequential). Other than formulating the call out framework, significant contribution of this paper is to demonstrate that relaxations of natural objective functions can be made separable (as in [25]), yet with bounded loss in performance ratio. Subsequent to the formulation, standard techniques of online stochastic optimization can be applied.

1.1 Selective Call Out: The Model

Let n be the number of ad networks $1, 2 \dots n$. We assume that impressions arrive from a fixed (unknown) distribution over a finite set U_I , and that there exists a finite set of bid values U_B , where $L = \max\{u \mid u \in U_B\}$. In the following, ad networks will be indexed by $i \in \{1, 2 \dots n\}$, impressions by $j \in U_I$ and bid values by $k \in U_B$. The problem setting involves several steps:

1. An impression (or keyword) j , assumed drawn from a distribution \mathcal{D} , comes to the exchange. There may be *multiple slots* associated with a single impression, corresponds to text ads being blocked together, different locations in the page, which are often characterized by different *discount rate*. Let there be M slots, with discount rates $1 \geq \varrho_1 \geq \varrho_2 \dots \geq \varrho_M \geq 0$. If a bidder bids v , then it is assumed that the bid for the ℓ^{th} slot is $v\varrho_\ell$. The case of $M = 1$ is common and correspond to a basic pay-per-impression mechanism with discount rate 1. All subsequent discussions apply to this common case as well.
2. Given an impression j , the bid of ad network i for impression j is drawn from a fixed distribution \mathbf{V}_{ij} such that the bid is v with probability p_{ijv} . Note that the bids of different networks are likely to be correlated based on the perceived value of the impression, however, conditioned on j , the specific dynamics of different bidders can be construed to be independent.

We assume that the exchange has learned or can predict these p_{ijv} given the impression j . This is an assumption similar to estimation processes used by search engines to predict the value of different advertisements (albeit under static standing bids).

3. The exchange decides on the subset S_j of networks to call out, subject to the *Call-Out constraints*, which roughly bounds the rate at which the exchange can send impressions to solicit bid from an ad network. This decision is executed before seeing the next impression.

To define a specific problem in the above framework, we need to specify (i) an objective function (ii) a model for call out constraints, and (iii) the comparison class. The **goal** is to design a call-out policy that satisfies (ii) and is near optimal in the objective function (i), when compared to other algorithms in class (iii). We discuss the instantiations of (i),(ii), and (iii) in the following.

(i) Objective Functions. We consider three different objective functions. (a) *Total Value*: The sum of the maximum bids in S_j over the arriving impressions j . (b) *GSP-Reserve* (defined above) and (c) *Revenue under posted price mechanism (take-it-or-leave-it prices)*. All the quantities are in expectation. Total value corresponds to the welfare. The GSP with an uniform reserve price is a common mechanism used in these settings. Posted Prices (different networks may get different prices) are also used in this context. Note that the mechanism is *parallel* posted price because the prices are posted before any bids are obtained, and not sequential as in [8, 4].

(ii) Call Out Constraint Models. The simplest model for the call out constraints is a model where one impression arrives at the exchange at each time step and if the total number of arrivals is m , ad network i can be solicited at most times for some known $\rho_i \leq 1$. We will refer to this as *time average model under uniform arrival* – which describes the constraints at the outgoing and incoming sides of the exchange respectively. Most of the paper will focus on this model – *primarily because we can show that other common models reduce to this variant*. The non-initiated reader can skip the description of these models and proceed to (iii).

On the incoming side of the exchange, standard practice is to assume bursty (Poisson) arrivals. We consider this generalization. On the outgoing side, the simple model allow the possibility that the call-outs to a network are made on contiguous subset of impressions. This misses the original goal that the ad network would receive the impressions at a “smooth” rate. A common model used for behavior is the *token bucket model* [26]. A token bucket has two parameters, bucket size σ and token generation rate ρ . The tokens represent sending rights, and the bucket size is the maximum number of tokens we can store. The tokens are generated at a rate of ρ per unit time, but the number of tokens never exceeds σ . In order to send, one needs to use a token, and if there are no tokens, one can not send. The output stream of a (σ, ρ) token bucket can be handled by a buffer of size σ and a time average rate of ρ – the buffer is initially full. Unlimited buffer size corresponds to the time average model.

(iii) Comparison Class. Given the call out constraints, we define the class of admissible policies.

Definition 1 *An admissible call out policy specifies (possibly with randomization), for each arriving impression j , the subset S_j of ad networks to call out, while satisfying all call out constraints over the entire sequence of impressions. The policy bases its decision on the prior information about the bid distributions, and has no knowledge of the actual bid values. In the case of GSP-Reserve*

mechanism, the call out policy also decides the reserve price for each impression. In the case of a posted price mechanism, the call out policy also decides the posted/take-it-or-leave-it prices.

Our comparison class is the set of all admissible call out policies which know the bid distributions for every impression, but do not know the actual realization of the bids. The performance of a policy is measured as the expected (over the bid distributions and the impression arrivals) objective value obtained per arriving impression, when impressions are drawn from \mathcal{D} .

1.2 Our Results, Roadmap and Other related Work

We provide three algorithms LP-VAL, LP-GSP, and LP-POST for the three objectives discussed for the time average uniform arrival model. We then prove that any approximation for any additive objective function translates naturally to an approximation for the other constraint models. Recall, L is the largest possible bid. *A policy is α -approximate if it achieves at least α times the performance of the optimal policy for the corresponding objective.* The algorithms will have a natural two-phase approach where we use the t initial impressions as a sample as exploration and subsequently use/exploit this algorithm (see Section 2 for more discussion). We show that:

Theorem 1 *Suppose the optimal policy has expected total value at least $\delta > 0$. For any $\epsilon > 0$, LP-VAL with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta \epsilon})$ impressions gives a $(1 - \frac{1}{e} - \epsilon)$ -approximate policy.*

Theorem 2 *Suppose the optimal GSP-Reserve policy has expected revenue at least $\delta > 0$. For any $\epsilon > 0$, LP-GSP with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta \epsilon})$ impressions gives a $O(1)$ -approximate policy, if all bid distributions satisfy the monotone hazard rate (MHR) property. Moreover, the call outs of the policy derived from LP-GSP are identical to those of the policy derived from LP-VAL.*

In particular, we show that when every bidder is solicited – GSP-Reserve achieves a revenue that is $O(1)$ factor of optimal welfare, when all bid distributions satisfy the MHR property. This is a common distributional assumption in economic theory, and is satisfied by many distributions [3]. This result is in the same spirit as (but immediately incomparable to) the result in [4], which relates the optimum sequential posted price revenue to the optimal welfare under the same assumptions. We are unaware of such results about GSP-Reserve.

Theorem 3 *Suppose the optimal posted price policy has expected revenue at least $\delta > 0$. For any $\epsilon > 0$, LP-POST with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta \epsilon})$ impressions gives a $(1 - \frac{1}{e} - \epsilon)$ -approximate policy.*

We do **not** need the MHR assumption unlike the result in [4] for sequential posted prices, since the comparison classes are different (and the prices are posted in parallel). We next show that:

Theorem 4 *For a given distribution on impressions \mathcal{D} , suppose we have an α -approximate policy for an objective which is additive given the allocations and the realizations (and is at least $\delta > 0$) in the time average uniform arrival call out model. Let $\sigma_i > \sigma \forall i$. Then we can convert the policy to a $(\alpha - \frac{1}{\sigma-1})$ -approximate policy in the token bucket model. The result extends to Poisson arrivals.*

Roadmap: We summarize the results on online stochastic convex optimization in Section 2. We subsequently discuss the the total value problem in Section 3. We discuss the GSP-Reserve problem in Section 4. The posted price problem is discussed in Section 5. The token bucket model and other arrival assumptions are discussed in Section 6. To draw a contrast to the stochastic results, the adversarial order setting is discussed in Appendix A.

Other Related Work: A combination of stochastic and online components appear in many different settings [14, 15, 2, 13] which are not immediately relevant to the call-out problem. We note that the bandwidth-like constraints (where the constraint is on a parameter different than the obtained value, as is the case for call-outs) has not studied in the bandit setting (see [16, 23]) because the horizon is constrained. Finally, bidding and inventory optimization problems studied in the context of ad exchanges [11, 10, 20], are not related to the call out optimization problems.

2 Preliminaries

Consider a maximizing a “separable” linear program (LP) \mathcal{L} defined on Q *global* constraints with right hand side b_i , such that the Lagrangian relaxation produced by the transferring these constraints to the objective function decouples into a collection of independent non-negative smaller LPs \mathcal{L}_j over n' variables and *local* constraints. This implies that the objective function of \mathcal{L} is a weighted linear combination of \mathcal{L}_j . The uniqueness of the optimum solutions for \mathcal{L}_j implies that \mathcal{L} reduces to finding the Lagrangian multipliers. The unique solution is achieved by adding “small perturbations”, see Rockafellar [25]. However, this approach only provides a certificate of optimality and a solution, once we are given the Lagrangians. The approach does not give us an algorithm to find the Lagrangian multipliers themselves.

Devanur and Hayes [9] showed that if the smaller LPs could be sampled with the same probability as their contribution to the objective of \mathcal{L} , and the derivatives of the Lagrangian can be bounded, then the Lagrangians derived from a small number of samples (suitably scaled) can be used to solve the overall LP. The weighted sampling reduces to the prefix of the input if the \mathcal{L}_j s arrive in random order (see [12]). This was extended to convex programs in [27]. The number of sample bound requires several (easy) Lipschitz type properties:

1. The optimum value of \mathcal{L} is at least $\delta > 0$ and the optimum solution of \mathcal{L}_j is at most R .
2. For each setting of the Lagrangians, every \mathcal{L}_j has a unique optimum solution.
3. Reducing b_i by a factor of $1 - \epsilon$ does not reduce \mathcal{L} by more than a factor of $(1 - \epsilon)$.
4. \mathcal{L} does not change by more than a constant times $1 + \epsilon$ if we alter the value of the optimum Lagrangian multipliers by a factor of $1 + \epsilon$.

Theorem 5 *Sample $t = \tilde{O}(\frac{n'QR}{\delta\epsilon})$ of the smaller linear programs and consider the linear program \mathcal{L}' which corresponds to the union of these smaller linear programs and suitably scaled global constraints. If we use the optimum Lagrangian multipliers corresponding to the global constraints of \mathcal{L}' to solve the decoupled instances of \mathcal{L}_j as they are available (in an online fashion) then we produce a $1 + \epsilon$ approximation to the optimum solution of \mathcal{L} .*

In the setting of Adwords, the smaller LPs correspond to the arrival of an impression, and the associated assignment. Thus the stochastic framework of the Adwords problem is obviously of relevance to the call out optimization framework. However the focus shifts on solving the smaller LPs which encode the call out decision. A nice outcome of the approach is a simple two phase algorithm; an *Exploration phase* where the samples are drawn, and an *Exploitation phase* where the Lagrangian multipliers are used. If Q, n', R are small, then the exploration phase can be (relatively) short and this yields a natural algorithm. Thus the goal of the rest of the paper would be to formulate separable convex relaxations and achieve the mentioned properties.

3 The Total Value Problem

In this section, we prove Theorem 1, and describe LP-VAL. Let q_j denote the probability that impression j arrives. We shall add infinitesimal random perturbations to p_{ijv} which shall not affect the performance of any policy but ensure p_{ijv} are in general positions, that is, any combination of them will almost surely create a non-singular matrix.

The LP Relaxation:

1. Let x_{ij} be the (conditional) probability that advertiser i was called out on impression j .
2. Let $y_{ijv\ell}$ be the probability that advertiser i bid the value t and was assigned the slot ℓ (also conditioned on j). The constraints are named as $A(x, y)$ and $B(x, y)$ as shown.

$$LP1 = \max \sum_j q_j \sum_v \sum_i \sum_\ell v_{\ell} y_{ijv\ell} \quad \text{s.t.} \quad \left. \begin{array}{l} \sum_j q_j x_{ij} \leq \rho_i \\ \sum_i \sum_v y_{ijv\ell} \leq 1 \\ x_{ij} \leq 1 \\ \sum_\ell y_{ijv\ell} \leq p_{ijv} x_{ij} \\ x_{ij}, y_{ijv\ell} \geq 0 \end{array} \right\} B(x, y) \left. \vphantom{\sum_j q_j x_{ij}} \right\} A(x, y)$$

Decoupling: Let λ_i^* be the optimum Lagrangian variable for the constraint $\sum_j q_j x_{ij} \leq \rho_i$. $LP1$ then decouples to smaller LPs, $LP2(j, \vec{\lambda}_i^*)$ subject to the constraints $A(x, y)$.

$$LP1 = LP1(\vec{\lambda}_i^*) = \sum_i \lambda_i^* \rho_i + \sum_j q_j LP2(j, \vec{\lambda}_i^*) \quad \text{where} \quad LP2(j, \vec{\lambda}_i^*) = \max \left(\sum_v \sum_i \sum_\ell v_{\ell} y_{ijv\ell} - \sum_i \lambda_i^* x_{ij} \right)$$

Solving $LP2(j, \vec{\lambda}_i^*)$. We begin by considering the dual. Let $\tau_{j\ell}$ be the dual of the constraint $\sum_i \sum_v y_{ijv\ell} \leq 1$. Let ξ_{ijv} correspond to the dual of the constraint $\sum_\ell y_{ijv\ell} \leq p_{ijv} x_{ij}$. Let ζ_{ij} correspond to the dual of $x_{ij} \leq 1$.

$$DualLP2(j, \vec{\lambda}_i^*) = \min \sum_i \zeta_{ij} + \sum_\ell \tau_{j\ell} \quad \text{s.t.} \quad \begin{array}{l} \tau_{j\ell} + \xi_{ijv} \geq v_{\ell} \\ \zeta_{ij} - \sum_v \xi_{ijv} p_{ijv} \geq -\lambda_i^* \\ \tau_{j\ell}, \xi_{ijv}, \zeta_{ij} \geq 0 \end{array}$$

Lemma 3.1 Let $\tau_{j\ell}^*$ be the optimum dual variables for LP2. Then (i) For all ℓ there exists i and $v \geq \tau_{j\ell}^*/\varrho_\ell$ such that $\xi_{ijv}^* = v\varrho_\ell - \tau_{j\ell}^*$. (ii) $\tau_{j\ell}^*/\varrho_\ell$ is non-increasing in ℓ .

Proof: For every ℓ there must be some i, v such that we have $\tau_{j\ell}^* + \xi_{ijv}^* = v\varrho_\ell$. Otherwise we can keep decreasing $\tau_{j\ell}^*$, keeping all other variables the same and contradict the optimality of the dual solution. Now $\xi_{ijv}^* \geq 0$ and the condition on t follows. The condition corresponds to the set of points (v, ξ_{ijv}^*) in the two dimensional (x, y) plane being above the lines $\{y = \varrho_\ell x - \tau_{j\ell}^*\}$.

For the second part, consider $\tau_{j\ell}^*$ and the i, v such that we have $\tau_{j\ell}^* + \xi_{ijv}^* = v\varrho_\ell$. Define t to be the support of ℓ . Let $v \geq \tau_{j\ell}^*/\varrho_\ell$ be the largest such support of ℓ . Consider $\tau_{j(\ell-1)}^*$. We have $(\tau_{j(\ell-1)}^* + \xi_{ijv}^*)/\varrho_{\ell-1} \geq v = (\tau_{j\ell}^* + \xi_{ijv}^*)/\varrho_\ell$. But $\varrho_{\ell-1} \geq \varrho_\ell$ and thus $\tau_{j\ell}^*/\varrho_\ell$ are non-increasing in ℓ . Moreover, if $\varrho_\ell = \varrho_{\ell-1}$ then $\tau_{j\ell}^* = \tau_{j(\ell-1)}^*$ (starting from the support of $\ell - 1$). \square

Decoupling LP2($j, \vec{\lambda}_i^*$) itself. Consider LP2($j, \vec{\lambda}_i^*$) with the Lagrangians $\tau_{j\ell}^*$. The problem decomposes under the constraints $B(x, y)$, to

$$LP2(j, \vec{\lambda}_i^*) = \sum_{\ell} \tau_{j\ell}^* + \sum_i LP3(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) \quad \text{where} \quad LP3(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) = \max \left(\sum_{\ell} \sum_v [v\varrho_\ell - \tau_{j\ell}^*] y_{ijv\ell} - \lambda_i^* x_{ij} \right)$$

Lemma 3.2 Define $\ell(v) = \arg \max_{\ell'} \{ \varrho_{\ell'} v - \tau_{j\ell'}^* \mid \varrho_{\ell'} v > \tau_{j\ell'}^* \}$ and $\ell(t) = M + 1$ if the set is empty. Set $y_{ijv\ell}^* = p_{ijv}$ if $\ell = \ell(v)$ and 0 otherwise. If $\sum_v \sum_{\ell} v\varrho_{\ell} y_{ijv\ell}^* \geq \lambda_i^*$ we set $x_{ij} = 1$ and $y_{ijv\ell} = y_{ijv\ell}^*$. Otherwise we set $x_{ij} = y_{ijv\ell} = 0$.

Proof: LP3($j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i$) is optimized at $x_{ij} = 1$ or $x_{ij} = 0$. This is because if $0 < x_{ij} < 1$ and $\sum_{\ell} \sum_v [v\varrho_{\ell} - \tau_{j\ell}^*] y_{ijv\ell} - \lambda_i^* x_{ij} > 0$ then we can multiply all the variables by $1/x_{ij}$ and have a better solution. If the latter condition is not true then $x_{ij} = 0$ is an equivalent solution. If $x_{ij} = 1$ the optimal setting for $y_{ijv\ell}$ is $y_{ijv\ell}^*$. (Note that $y_{ijv\ell}^*$ is uniquely determined for a fixed t .) Thus the overall optimization follows from comparing the $x_{ij} = 1$ and $x_{ij} = 0$ case. \square

Interpretation and the Call Out Algorithm: Given $\{\tau_{j\ell}^*\}_{\ell=1}^M$, the distribution $\{p_{ijv}\}$ for i , is divided into at most $M + 1$ pieces (some of the pieces can be a single point) given by the upper envelope (the constraint max) of the lines $\{\varrho_{\ell} x - \tau_{j\ell}^*\}_{\ell=1}^M$ and the line $y = 0$, in the x - y coordinate plane. Intuitively, seeing the value $x = t$, if the upper envelope corresponds to the equation $\varrho_{\ell} x - \tau_{j\ell}^*$ then we are “interested” in the slot ℓ . If the weighted (by ϱ_{ℓ}) sum of interests, given by $\sum_v v\varrho_{\ell} y_{ijv\ell}^*$ exceeds λ_i^* , then it is beneficial to call out i . We call out based on this condition and allocate the slots in decreasing order of bids.

Analysis: The LP2 solution satisfies: $LP1 = \sum_{i,v} v p_{ijv} \varrho_{\ell(v)}$ and $\sum_{i,v:\ell(v)=\ell} p_{ijv} = 1$. For each slot ℓ , let $w_i(\ell) = \sum_{v:\ell(v)=\ell} v p_{ijv}$ and $u_i(\ell) = \sum_{v:\ell(v)=\ell} p_{ijv}$. Order the i in non-increasing order of $w_i(\ell)/u_i(\ell)$ inside the slot. If we call out to i and get t , then for the sake of analysis we will consider its contribution to slot $\ell(t)$ only. Moreover, we stop the contribution to a slot ℓ if any any of the i return a value t with $\ell(t) = \ell$. The best M ordered bids outperform the analyzed contribution in every scenario. Therefore it suffices to bound the contribution of the analysis.

Lemma 3.3 Suppose we are given a set of independent variables Y_i such that $\Pr[Y_i \neq 0] = u_i$ and $E[Y_i] = w_i$. Consider the random variable Y corresponding to the process which orders the variables $\{Y_i\}$ in non-increasing order of w_i/u_i , and stops as soon as the first non-zero value is seen. Then $E[Y] = \sum_i \prod_{i' < i} (1 - u_{i'}) w_i \geq \sum_i w_i (1 - e^{-1})$.

Proof: Let $F(\{(w_i, u_i)\}) = \sum_i \prod_{i' < i} (1 - u_{i'}) w_i$. Let $\lambda = \sum_i w_i / \sum_i u_i$. Given the sequence $\{(w_i, u_i)\}$ where w_i/u_i are non-increasing, if there exists an i such that $w_i/u_i \neq w_{i+1}/u_{i+1}$, then define a new sequence $\{(w'_i, u_i)\}$ as follows:

$$w'_{i'} = \begin{cases} w_{i'} & \text{if } i' \neq i, i+1 \\ w_i - \Delta & \text{if } i' = i \\ w_{i+1} + \Delta & \text{if } i' = i+1 \end{cases} \quad \text{where} \quad \Delta = \frac{\frac{w_i}{u_i} - \frac{w_{i+1}}{u_{i+1}}}{\frac{1}{u_i} + \frac{1}{u_{i+1}}}$$

Note that $\sum_i w_i = \sum_i w'_i$ and w'_i/u_i remains non-increasing. Now $F(\{(w_i, u_i)\}) - F(\{(w'_i, u_i)\}) = \prod_{i' < i} (1 - u_{i'}) \Delta - \prod_{i' < i+1} (1 - u_{i'}) \Delta = \prod_{i' < i} (1 - u_{i'}) u_i \Delta > 0$. Thus, we can repeatedly perform the above steps till we get a sequence such that w'_i/u_i remains the same for all i and $\sum_i w_i = \sum_i w'_i$. Clearly $w'_i = \lambda u_i$ in this case. The function F will continue to decrease, and

$$F(\{(w_i, u_i)\}) \geq F(\{(w'_i, u_i)\}) = \sum_i \prod_{i' < i} (1 - u_{i'}) \lambda u_i = \lambda \left(1 - \prod_i (1 - u_i) \right) \geq \lambda (1 - e^{-\sum_i u_i})$$

But $\frac{1}{x}(1 - e^{-x})$ is decreasing over $[0, 1]$ and the worst case is $x = 1$. \square

In slot ℓ (renumbering the advertisers in the order of $w_i(\ell)/u_i(\ell)$) we get an expected reward of $w_i(\ell)$ if we reach i . But the events are independent in a particular slot. Thus the expected reward in a slot is bounded by (using independence and Claim 3.3) to be $(1 - e^{-1})$ times $\sum_i w_i(\ell)$. We now apply linearity of expectation across the slots – observe that the events across the slots are quite correlated. The expected reward is at least $(1 - e^{-1})$ times $\sum_\ell \sum_i w_i(\ell) = LP1$. Theorem 1 follows from Lemmas 3.2 and 3.3 and the application of Theorem 5.

4 Generalized Second Price with reserve (GSP-Reserve)

The call outs for this problem would be exactly the same as the algorithm in Section 3. We will however adjust the reserve prices. The reserve price will be the same for all the advertisers being called out on that impression. In fact either we will run a single slot auction with a reserve price, or simply GSP for the M slots. The decision will depend on the LP solution found for this specific impression (and the contributions of different parts of the LP). Recall that the bid distribution \mathbf{V}_{ij} of advertiser i on impression j is assumed to satisfy the MHR property. We use the following:

Lemma 4.1 (Lemma 3.3 in [4]) For any random variable V following an MHR distribution, let $v^* = \arg \min_v \{v | v \Pr[V \geq v] \geq \frac{1}{2} \sum_{v' \geq v} v' \Pr[V = v']\}$. Then $\Pr[V \geq v^*] \geq e^{-2}$.

The next lemma is a restatement of Lemma 3.2 and the subsequent analysis.

Lemma 4.2 Given an impression $j(t)$, and define $v_1(t) = \max_{\ell: \varrho_1 \neq \varrho_\ell} (\tau_{j1}^* - \tau_{j\ell}^*) / (\varrho_1 - \varrho_\ell)$. The call out to a set $S(t) \neq \emptyset$, ensures that $\sum_i \sum_{v \geq v_1(t)} p_{ijv} \geq 1$.

Definition 2 Given an impression $j(t)$, and the call out decision to a set $S(t) \neq \emptyset$ at time t , let $v_m^*(i, t) = \min\{v | 2v \Pr[V_{ij} \geq v] \geq \sum_{v' \geq v} v' p_{ijv'}\}$ and $\Psi(t) = \{i | i \in S(t) \text{ and } v_1(t) \leq v_m^*(i, t)\}$. Note that using Lemmas 4.1, and 4.2, we have $|\Psi(t)| \leq \lfloor e^2 \rfloor = 7$ since $i \in \Psi(t)$ contributes a probability mass of at least e^{-2} .

Lemma 4.3 Given an impression $j(t)$, and the call out decision to a set $S(t) \neq \emptyset$ at time t , we can set a single threshold $v^*(t) \geq v_1(t)$ such that if we set a reserve price $v^*(t)$ for a single slot then the revenue (ignoring the multiplicative discount factor ϱ_1) is at least $\frac{1}{4(7e^2+1)} \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v p_{ijv}$.

Proof: Let $\sum_{i \in S(t)} \sum_{v \geq v_1(t)} v p_{ijv} = Z$. We have two cases, (i) $\sum_{i \in \Psi(t)} \sum_{v \geq v_1(t)} v p_{ijv} \geq 7e^2 Z / (7e^2 + 1)$ or (ii) otherwise. In case (i), pick the $i \in \Psi(t)$ such that $\sum_{v \geq v_1(t)} v p_{ijv}$ is maximized, which is at least $e^2 Z / (7e^2 + 1)$ since $|\Psi(t)| \leq 7$. Let V_{ij} be the random variable that corresponds to the bid of advertiser i on impression j . Now since $v_m^*(i, t) \geq v_1(t)$ we have that

$$\sum_{v \geq v_m^*(i, t)} v p_{ijv} \geq \Pr[V_{ij} \geq v_m^*(i, t) | V_{ij} \geq v_1(t)] \sum_{v \geq v_1(t)} v p_{ijv} \geq \Pr[V_{ij} \geq v_m^*(i, t)] \sum_{v \geq v_1(t)} v p_{ijv} \leq \frac{1}{e^2} \sum_{v \geq v_1(t)} v p_{ijv}$$

which is at least $Z / (7e^2 + 1)$. Now, if we set $v^*(t) = v_m^*(i, t)$ then just from i we have $\sum_{i \in S(t)} \sum_{v \geq v^*(t)} p_{ijv} \geq \frac{1}{2} \sum_{v \geq v_m^*(i, t)} v p_{ijv}$ and therefore in this case the lemma is true.

In case (ii), we have $\sum_{i \in S(t) \setminus \Psi(t)} \sum_{v \geq v_1(t)} v p_{ijv} \geq Z / (7e^2 + 1)$. But for each $i \in S(t) \setminus \Psi(t)$ we have $v_1(t) \sum_{v \geq v_1(t)} p_{ijv} \geq \frac{1}{2} \sum_{v \geq v_1(t)} v p_{ijv}$ and as a consequence, $v_1(t) \sum_{i \in S(t) \setminus \Psi(t)} \sum_{v \geq v_1(t)} p_{ijv}$ is at least $Z / (2(7e^2 + 1))$. Consider setting $v^*(t) = v_1(t)$. Let $p = \sum_{i \in S(t) \setminus \Psi(t)} \sum_{v \geq v_1(t)} p_{ijv}$. Since $p \leq 1$ (from definition of $v_1(t)$, see Lemma 4.2) the probability of sale is at least $(1 - \frac{1}{e})p$ which is bounded below by $p/2$. The Lemma follows in this case as well. \square

Lemma 4.4 Given an impression $j(t)$, and the call out decision to a set $S(t) \neq \emptyset$ at time t , consider (i) If $\sum_{i \in S(t)} \sum_v v \varrho_\ell(v) y_{ijv\ell}^* \geq 3 \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v \varrho_1 y_{ijv1}^*$ then call-out to $S(t)$ and run regular GSP. (ii) Otherwise call-out to $S(t)$ and run a single slot auction with the threshold $v^*(t)$ given by Lemma 4.3. This algorithm gives a revenue which is $\Omega(1)$ factor of the LP bound on efficiency which is given by $\sum_v \sum_{i \in S(t)} \sum_v v \varrho_\ell y_{ijv\ell}^*$. Note that the call out decisions are based on optimizing the total value/ efficiency of the slots, and thus are feasible.

Proof: Let the non-increasing ordered list of values that are returned for a time step be $a_1(t)$. Suppose we are in case (i). Then the revenue of GSP is at least $\sum_{r=1}^M \varrho_r a_{r+1}(t)$. Now since ϱ_r are decreasing, and $a_r(t)$ are non-increasing,

$$\sum_{r=1}^M \varrho_r a_{r+1}(t) \geq \sum_{r=1}^M \varrho_r a_r(t) - \varrho_1 a_1(t) \implies \mathbf{E} \left[\sum_{r=1}^M \varrho_r a_{r+1}(t) \right] \geq \mathbf{E} \left[\sum_{r=1}^M \varrho_r a_r(t) \right] - \varrho_1 \mathbf{E} [a_1(t)] \quad (1)$$

We know from the rounding in Section 3 that $\mathbf{E} \left[\sum_{r=1}^M \varrho_r a_r(t) \right] \geq (1 - \frac{1}{e}) \sum_{i \in S(t)} \sum_v v \varrho_{\ell} y_{ijv\ell(v)}^*$.

We observe that $\mathbf{E} [a_1(t)] \leq \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v y_{ijv1}^*$. This is easily seen if we write an LP for the maximum value seen (this LP is for analysis only). Let x_{iv} be the probability that $i \in S(t)$ is the maximum with value v . Then (we drop the index j for convenience):

$$\mathbf{E} [a_1(t)] \leq LPMAX = \max \sum_i \sum_v v x_{iv} \quad \text{s.t.} \quad \begin{array}{l} \sum_i \sum_v x_{iv} \leq 1 \\ x_{iv} \leq p_{iv} \\ x_{iv} \geq 0 \end{array}$$

The optimum solution of $LPMAX$ is $x_{iv}^* = p_{iv}$ for $v > \tau$ and $x_{iv}^* \leq p_{iv}$ for one i and $v = \tau$. Here τ is the optimum dual variable for the constraint $\sum_{i \in S(t)} \sum_v x_{iv} \leq 1$. Note that $\sum_{i \in S(t)} \sum_v x_{iv}^* = 1$.

For $v > v_1(t)$ we have $y_{ijv1}^* = p_{ijv}$ and $v < v_1(t)$ we have $y_{ijv1}^* = 0$. Moreover $\sum_{i \in S(t)} \sum_{v \geq v_1(t)} y_{ijv1}^* = 1$. Likewise for $v > \tau$ we have $x_{iv}^* = p_{iv}$ and $v < \tau$ we have $x_{iv}^* = 0$ and $\sum_{i \in S(t)} \sum_{v \geq \tau} x_{iv}^* = 1$.

Suppose that $\tau < v_1(t)$. We arrive at a contradiction because $\sum_{i \in S(t)} \sum_{v \geq \tau} x_{iv}^* > \sum_{i \in S(t)} \sum_{v \geq v_1(t)} y_{ijv1}^* = 1$ which implies that we are exceeding the probability mass of 1 for the maximum. On the other hand if $\tau > v_1(t)$, then we again have a contradiction that $\sum_{i \in S(t)} \sum_{v \geq v_1(t)} y_{ijv1}^* > \sum_{i \in S(t)} \sum_{v \geq \tau} x_{iv}^* = 1$ which implies $\{y_{ijv1}^*\}$ were not feasible.

As a consequence, $\tau = v_1(t)$ and for $v > \tau = v_1(t)$ we have $x_{iv}^* = y_{ijv1}^* = p_{ijv} = p_{iv}$. For $v < \tau = v_1(t)$ we have $x_{iv}^* = y_{ijv1}^* = 0$. Therefore $\mathbf{E} [a_1(t)] \leq LPMAX = \sum_{i \in S(t)} \sum_v v x_{iv}^* = \sum_{i \in S(t)} \sum_v v y_{ijv1}^*$ as claimed. Applying this claim to Equation 1, and the fact that $\sum_{i \in S(t)} \sum_v v y_{ijv1}^* \leq \frac{1}{3} \sum_{i \in S(t)} \sum_v v \varrho_{\ell} y_{ijv\ell(v)}^*$, we get

$$\mathbf{E} \left[\sum_{r=1}^M \varrho_r a_{r+1}(t) \right] \geq \left(1 - \frac{1}{e} - \frac{1}{3} \right) \sum_{i \in S(t)} \sum_v v \varrho_{\ell} y_{ijv\ell(v)}^*$$

Thus in this case the expected revenue is $\Omega(1)$ of the LP bound on the efficiency.

Suppose we are in case (ii). By Lemma 4.3 we are guaranteed an expected revenue of $\Omega(1)$ times

$$\varrho_1 \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v p_{ijv} \geq \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v \varrho_1 y_{ijv1}^* \geq \frac{1}{3} \sum_{i \in S(t)} \sum_v v \varrho_{\ell(v)} y_{ijv\ell(v)}^*$$

In this case also the expected revenue is $\Omega(1)$ of the LP bound; the lemma follows. \square

We are ready to prove Theorem 2.

Proof:(Of Theorem 2). Let APP be the policy that approximately maximizes the efficiency. Let OPT_G be the optimum GSP-Reserve policy. Let OPT_B be the optimum policy which maximizes the total value.

Given a policy P let $\text{GSP}(P)$ denote the expected revenue of the policy if the charged as GSP, $\text{BESTKW}(P)$ denote the expected (weighted) efficiency. Then for any policy $\text{GSP}(P) \leq \text{BESTKW}(P)$. Let $R(\text{APP})$ be the revenue of the policy in Lemma 4.4. Therefore, for some absolute constant $\alpha \geq 1$,

$$\text{GSP}(\text{OPT}_G) \leq \text{BESTKW}(\text{OPT}_G) \leq \text{BESTKW}(\text{OPT}_B) \leq LP1 \leq \alpha R(\text{APP})$$

The theorem follows (again appealing to Theorem 5). \square

5 The Posted Price Revenue Problem

In this section we prove Theorem 3. The flow of ideas would be similar to that of Section 3, however we would require a different LP.

1. Let x_{ijv} denote the probability that advertiser i was called out on impression j with price t .
2. Let $y_{ijv\ell}$ be the probability that advertiser i was offered price t and assigned to slot ℓ . The revenue generated in this event is $\varrho_\ell t$.

Let $\tilde{p}_{ijv} = \sum_{t \geq \varrho_\ell} p_{ijvt}$ denote the probability that on impression j , advertiser i has a valuation at least t . The LP is:

$$LP4 = \max \sum_j q_j \sum_v \sum_i \sum_\ell v \varrho_\ell y_{ijv\ell} \quad \text{s.t.} \quad \left. \begin{array}{l} \sum_j q_j \sum_v x_{ijv} \leq \rho_i \\ \sum_i \sum_v y_{ijv\ell} \leq 1 \\ \sum_v x_{ijv} \leq 1 \\ \sum_\ell y_{ijv\ell} \leq \tilde{p}_{ijv} x_{ijv} \\ x_{ijv}, y_{ijv\ell} \geq 0 \end{array} \right\} B(x, y) \left. \vphantom{\sum_j q_j \sum_v x_{ijv}} \right\} A(x, y)$$

Decoupling: Let λ_i^* be the optimum dual variable for the constraint $\sum_j q_j \sum_v x_{ijv} \leq \rho_i$. The result of the decoupling is (subject to $A(x, y)$):

$$LP4 = LP4(\vec{\lambda}_i^*) = \sum_i \lambda_i^* \rho_i + \sum_j q_j LP5(j, \vec{\lambda}_i^*) \quad \text{where} \quad LP5(j, \vec{\lambda}_i^*) = \max \left(\sum_v \sum_i \sum_\ell v \varrho_\ell y_{ijv\ell} - \sum_i \sum_v \lambda_i^* x_{ijv} \right)$$

Solving LP5($j, \vec{\lambda}_i^*$). Let $\tau_{j\ell}^*$ be the dual of the constraint $\sum_i \sum_v y_{ijv\ell} \leq 1$. And again we have a decoupling (subject to $B(x, y)$):

$$LP5(j, \vec{\lambda}_i^*) = \sum_\ell \tau_{j\ell}^* + \sum_i LP6(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) \quad \text{where} \quad LP6(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) = \max \sum_v \sum_\ell (t \varrho_\ell - \tau_{j\ell}^*) y_{ijv\ell} - \lambda_i^* x_{ijv}$$

The next lemma follows from inspection (assuming x_{ijv} being fixed).

Lemma 5.1 *In the optimum solution to $LP6(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i)$ we must have $\sum_\ell y_{ijv\ell} = \tilde{p}_{ijv} x_{ijv}$ and moreover we should have*

$$y_{ijv\ell}^* = \begin{cases} \tilde{p}_{ijv} x_{ijv} & \text{if } \ell = \arg \max_{\ell'} \{v \varrho_{\ell'} - \tau_{j\ell'}^* | t \varrho_{\ell'} > \tau_{j\ell'}^*\} \\ 0 & \text{otherwise} \end{cases}$$

Definition 3 Note that this fixes a **slot function** $f_{ij}(v) = \arg \max_{\ell} \{v \varrho_{\ell} - \tau_{j\ell}^* | v \varrho_{\ell} > \tau_{j\ell}^*\}$ for each t . Again note that this corresponds to the upper envelope of the lines $y = x \varrho_{\ell} - \tau_{j\ell}^*$ and $y = 0$. Thus $f_{ij}(v)$ can be represented as a piecewise linear function of v . We can view this as “ i is offered price v with the understanding that if i accepts then i we will attempt to give the slot $f_{ij}(v)$ to i ”. Note slot $M + 1$ corresponds to $\varrho_{M+1} = 0$.

Therefore we have: $LP6(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) = LP7(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i)$ where

$$LP7(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i) = \max \sum_v \left[(t \varrho_{f_{ij}(v)} - \tau_{j f_{ij}(v)}^*) \tilde{p}_{ijv} - \lambda_i^* \right] x_{ijv} \quad \text{s.t.} \quad \begin{cases} \sum_v x_{ijv} \leq 1 \\ x_{ijv} \geq 0 \end{cases}$$

Lemma 5.2 In the optimum solution to $LP7(j, \vec{\lambda}_i^*, \vec{\tau}_{j\ell}^*, i)$ we have

$$x_{ijv'}^* = \begin{cases} 1 & \text{if } v' = \arg \max_v (v \varrho_{f_{ij}(v)} - \tau_{j f_{ij}(v)}^*) \tilde{p}_{ijv} | (t \varrho_{f_{ij}(v)} - \tau_{j f_{ij}(v)}^*) \tilde{p}_{ijv} > \lambda_i^* \\ 0 & \text{otherwise} \end{cases}$$

Note that we can compute v' easily because $f_{ij}(v)$ can be represented as a relatively simple function of v , based on Definition 3.

Thus we are offering an unique price $v' = \arg \max_v \{(v \varrho_{f_{ij}(v)} - \tau_{j f_{ij}(v)}^*) \tilde{p}_{ijv} | (v \varrho_{f_{ij}(v)} - \tau_{j f_{ij}(v)}^*) \tilde{p}_{ijv} > \lambda_i^*\}$ to the advertiser i (with the intuitive idea that the advertiser would be considered for slot $f_{ij}(v')$). We denote $v' = \infty$ if the condition does not hold for any value of t possible for i . The probability that the advertiser accepts is $\tilde{p}_{ijv'}$ and this is also $y_{ijv'\ell}^*$ where $\ell' = f_{ij}(v')$. Therefore given j , to advertiser i we offer a price of $v'(i, j)$ which is a function i, j with a slot $\ell(i, j) = f_{ij}(v'(i, j))$ in mind.

The Interpretation and the Rounding: We note that for every slot ℓ we have $\sum_{i:\ell=\ell(i,j)} y_{ijv'\ell}^* \leq 1$. And the expected reward is $\sum_{i:\ell=\ell(i,j)} \varrho_{\ell} t'(i, j) y_{ijv'\ell}^*$.

For each slot we order the advertiser in non-increasing order of $v'(i, j)$, and we can perform the same analysis as Lemma 3.3. And for each slot we expect a revenue of $(1 - \frac{1}{e})$ fraction. Since we have one slot in mind for each advertiser we can sum up across the slots and expect a $(1 - \frac{1}{e})$ fraction of the LP revenue. **Note:** that in a given scenario, after all the bids are accepted, we can make the best allocations and only increase the revenue in the process. Thus Theorem 3 follows.

6 Handling Bursts: Token Buckets and Poisson Arrivals

We shall now present the proof of Theorem 4. Recall that the token bucket model starts with a full buffer – and thus unlimited buffer size corresponds to the time average model.

The Uniform Arrival case: We first consider the uniform arrival case, where an impression arrives every unit time step. Consider any algorithm \mathcal{A} which makes call outs according to the time average model. We define \mathcal{A}' to behave exactly as \mathcal{A} except that in the case the token bucket is empty for advertiser i , no call out is made to i in \mathcal{A}' .

Let the expected value received by \mathcal{A} in sending impressions to ad network i be R_i , so that the total expected value obtained by our algorithm is $R = \sum_{i=1}^n R_i$. Let the corresponding values obtained by our algorithm in the token bucket model be R'_i and R' respectively, and so $R' = \sum_{i=1}^n R'_i$. We assume that in the beginning, each token bucket is full of tokens. This is safe to assume when the process runs for sufficiently long time $T \gg \sigma_i/\rho_i$, since one can simply ignore the impressions arriving in the first σ_i/ρ_i units of time, and lose only a small fraction of the objective value in expectation.

Lemma 6.1 $R'_i \geq (1 - \frac{1}{\sigma_i-1})R_i$.

Proof: Let us focus on a single token bucket, that of ad network i . We simply need to calculate the expected fraction of times that an attempted sending by our algorithm to ad network i succeeds. We shall show that this fraction is at least $(1 - \frac{1}{\sigma_i-1})$. Given a set of arriving impressions, every order of their arrival is equally likely, so every impression is equally likely to have a failed attempt at sending, and so the value obtained is at least $(1 - \frac{1}{\sigma_i-1})R_i$.

For every arriving impression, the algorithm decides whether to attempt sending it to ad network i . We lower bound the number of successful sending by the following process: send whenever the algorithm attempts sending till the token bucket is empty, then neglect all impressions till the bucket becomes full (dormant period), then again send whenever the algorithm attempts to do so, and so on. This can be observed to be a lower bound because for every impression that the algorithm attempts but fails to send while the said process succeeds, there must be a unique previous successful sending by the algorithm during the dormant period of the process.

At time step t , let X_t denote the amount of tokens in the bucket. For all t , $0 \leq X_t \leq \sigma_i$, and $X_0 = \sigma_i$. Let τ be a stopping time, defined as the first time when $X_t < 1$. It is easy to see that $\mathbf{E}[\tau] < \infty$, since there is a non-zero probability of a burst of impressions that are sent to ad network i . Let $Z_t = \sigma_i - X_t$, so $Z_0 = 0$. Note that $\sigma_i \geq Z_\tau > \sigma_i - 1$. Let Y_t be the number of impressions sent to ad network i before or at time step t . Also, at every time step, let $\alpha_i \leq \rho_i$ denote the probability that the algorithm sends an arriving impression to ad network i , i.e., $\alpha_i = \sum_j q_j x_{i,j}$. We prove the following subclaim:

Claim 1 $\mathbf{E}[Y_\tau] \geq (\sigma_i - 1)^2 \alpha_i / \rho_i$.

Proof: First, we observe that $Y_t - Y_{t-1}$ is 1 with probability α_i , and zero otherwise. So $(Y_t - \alpha_i t)$ is a martingale. Since $\mathbf{E}[\tau] < \infty$, so applying Doob's optional stopping theorem on the martingale, we get that $\mathbf{E}[Y_\tau - \alpha_i \tau] = 0$, so $\mathbf{E}[Y_\tau] = \alpha_i \mathbf{E}[\tau]$. We shall now show that $\mathbf{E}[\tau] \geq (\sigma_i - 1)^2 / \rho_i$.

Let \mathbb{F}_t , $t \geq 0$ denote the filtration for the sequence $\{Z_t\}$, that is, \mathbb{F}_t is the information about all the values $Z_0, Z_1 \dots Z_t$. We now identify a process A_t such that $Z_t^2 - A_t$ is a martingale, and $A_0 = 0$. Clearly, such a process exists (by Doob decomposition), and $A_{t+1} = \mathbf{E}[Z_{t+1}^2 | \mathbb{F}_t] - Z_t^2$. At each time step, Z_t decreases by ρ_i (lower bounded by zero), and increases by 1 (upper bounded by σ_i) with probability α_i . Thus, when $\sigma_i - 1 \geq Z_t \geq \rho_i$,

$$\begin{aligned}
\mathbf{E} [Z_{t+1}^2 | \mathbb{F}_t] &= (1 - \alpha_i)(Z_t - \rho_i)^2 + \alpha_i(Z_t + 1 - \rho_i)^2 \\
&= (1 - \alpha_i)(Z_t^2 - 2Z_t\rho_i + \rho_i^2) + \\
&\quad \alpha_i(Z_t^2 + 2Z_t(1 - \rho_i) + (1 - \rho_i)^2) \\
&= Z_t^2 + 2(\alpha_i - \rho_i)Z_t + \alpha_i(1 - \rho_i)^2 + (1 - \alpha_i)\rho_i^2 \\
&\leq Z_t^2 + \rho_i^2 + \alpha_i(1 - 2\rho_i) \leq Z_t^2 + \rho_i
\end{aligned}$$

The inequalities follow from the facts that $\alpha_i \leq \rho_i \leq 1$ and $Z_t \geq 0$. Moreover, when $\rho_i \geq Z_t \geq 0$, Z_{t+1} is zero with probability $(1 - \alpha_i)$, and at most $1 - \rho_i + Z_t \leq 1$ with probability α_i . Thus $\mathbf{E} [Z_{t+1}^2 | \mathbb{F}_t] \leq \alpha_i \leq \rho_i$. This implies that $A_{t+1} - A_t < \rho_i$ and so $A_t < \rho_t$, for all $t < \tau$.

Now we apply Doob's optional stopping theorem to obtain that $\mathbf{E} [Z_\tau^2 - A_\tau] = 0$. Since $Z_\tau > \sigma_i - 1$, we have $(\sigma_i - 1)^2 - \rho_i \mathbf{E} [\tau] \leq 0$. This proves the claim. \square

The process will succeed in all attempts to send impressions to ad network i till time τ . After this, the process is dormant and neglects all impressions for the next Z_τ/ρ_i steps, so that the bucket becomes full, and then the same argument can be repeated. During the dormant period, the expected number of attempted sending by the original algorithm is at most $Z_\tau\alpha_i/\rho_i$. So the fraction of sending attempts that are successful is at least

$$\frac{\frac{(\sigma_i - 1)^2 \alpha_i}{\rho_i}}{\frac{(\sigma_i - 1)^2 \alpha_i}{\rho_i} + \frac{\sigma_i \alpha_i}{\rho_i}} > 1 - \frac{1}{\sigma_i - 1} .$$

This completes the proof of Lemma 6.1. \square

Lemma 6.1 implies that if $\sigma_i > \sigma \forall i$, then $R' \geq (1 - \frac{1}{\sigma-1})R$.

Poisson Arrival Process. Now we consider the token bucket model under the assumption that the time difference between the arrival of two consecutive impressions is drawn from an exponential distribution (with mean 1, for normalization). This implies that the arrival process is a unit rate *Poisson process*. The tokens fill up *continuously*, at a rate of ρ_i tokens per unit time. Extending our analysis from discrete to continuous martingales, in proof of Lemma 6.1, we get,

Lemma 6.2 $R'_i \geq (1 - \frac{1}{\sigma_i - 1})R_i$ in the Poisson arrival process as well.

Proof: X_t, Y_t and $Z_t = \sigma_i - X_t$ are now continuous processes. The Doob optional stopping theorem holds for continuous martingales, and we can find a process A_t such that $Z_t^2 - A_t$ is a continuous martingale. Let τ be again the smallest t such that $X_t < 1$. It again boils down to showing that $\mathbf{E} [\tau] \geq (\sigma_i - 1)^2/\rho_i$. To show this, we need to show that $A_t < \rho_i$ for all $t < \tau$. Between time t and $t + dt$, where dt is an infinitesimal increment, there is a probability dt that an impression arrives (since the arrival time is a Poisson process), and thus probability $\alpha_i dt$ that the algorithm attempts to send an impression to ad network i , while $\rho_i dt$ tokens are generated. Thus, when $0 < Z_t < \sigma_i - 1$, we have (neglecting higher powers of dt):

$$\begin{aligned}
\mathbf{E} [Z_{t+dt}^2 | \mathbb{F}_t] &= (1 - \alpha_i dt)(Z_t - \rho_i dt)^2 + \alpha_i dt(Z_t + 1 - \rho_i dt)^2 \\
&= Z_t^2 + \alpha_i dt - 2Z_t(\rho_i - \alpha_i)dt \leq Z_t^2 + \rho_i dt
\end{aligned}$$

So we infer that $A_{t+dt} - A_t \leq \rho_i dt$, and so $A_t < \rho_i t \forall t$. Finally, applying Doob’s optional stopping theorem for continuous martingales, we get that $\mathbf{E}[Z_\tau^2 - A_\tau] = 0$, which implies that $(\sigma_i - 1)^2 - \rho_i \mathbf{E}[\tau] < 0$. \square

This completes the proof of Theorem 4.

7 Experimental Study

In this section we experimentally explore aspects of our model and performance of algorithms, focusing solely on *sales* as our measure of performance in our simulations, and sticking to the basic mechanism where ad networks report their bids. In each online query, there is a single impression with discount factor 1, and a minimum price set by the publisher. The impression is considered to be sold if at least one of the ad networks that are called out returns a bid higher than this minimum price. The objective is to maximize the number of impressions sold. Sales can be viewed as a special case of the efficiency problem, where all bids are 0 or 1 (they are willing to buy the item at a minimum price quoted by the publisher, or not), and so the corresponding algorithm is applicable. The following questions are important.

1. How much does estimating the bid distributions help? We assumed that the Exchange will estimate the *survival probability* p_{ijk} ’s for ad network i , impression j and bid k (that is, probability that the bid is at least k), via machine learning and data mining techniques. Can we require less out of these techniques? Motivated by sponsored search systems, it is tempting to only estimate the expected bid for each i, j , rather than the entire distribution of p_{ijk} . Further, for the sales metric, it suffices to only estimate the probability of a bid above k for i, j , while for the value metric, we use the entire distribution of p_{ijk} . We compare performance of algorithms that use different amounts of information about p_{ijk} ’s.
2. *How does the error in estimation affect the performance of algorithms?* Methods that estimate p_{ijk} will have errors. We study the influence of such errors on our algorithms.
3. *What is the benefit of the optimization over simple, natural schemes?* We evaluate our LP-based solution and compare it to other simpler schemes.

Before proceeding to the simulation setup, we list the algorithms we consider. The linear programming based algorithm uses probability thresholds to decide on call outs. However alternative (and simpler) implementations typically would choose subsets based on some criterion. Thus we have a natural partition of set and threshold based algorithms.

Set Based Algorithms: All our candidate algorithms sort the ad networks according to some criterion, which may depend on the arriving impression as well as history of the process. The algorithms then attempt to call out to the top ad networks in this ordering, which succeeds if those ad networks have bandwidth remaining. It remains to specify how many ad networks are picked for a particular impression. Set based algorithms are simply those that pick the top k ad

networks for every impression. The first two algorithms are baseline algorithms that are motivated by scheduling algorithms, but they ignore the probability estimations. The next two algorithms use the probability estimation to guide the call outs (and are based on different but widely occurring ways of characterizing high probability subsets).

1. **Random:** chooses a random ordering.
2. **MaxRemBand:** orders by (decreasing) remaining bandwidth.
3. **MaxProb:** orders by (decreasing) survival probability.
4. **MaxExp:** orders by (decreasing) expectation of a bid. While expected bid value can be clearly shown to be inadequate on specially designed examples, we evaluate the value of this information on our more generally generated data sets.

Threshold Algorithms: The above algorithms all have a parameter k , the size of the set of ad networks to which call out is attempted for each impression (subject to bandwidth availability). Each of the algorithms in this class induce an ordering among ad networks and choose a prefix of this ordering such that the survival probabilities $\sum_i p_i x_i$ of the chosen ad networks add up to a *threshold*. If the sum exceeds the threshold, then the algorithm probabilistically decides whether to attempt call out to the last ad network in the chosen ordering. The probability of this event is set such that the sum of survival probabilities times the probability of getting called out sums up to *exactly* the threshold. Note that our LP-based algorithm is close to a threshold algorithm with the threshold set at 1, except that the Lagrangians z_i^* that were learned also acts as a *cut-off*. We found this algorithm to be conservative in spending the call outs and decided to set this to a variable threshold. The class of threshold algorithms is justified by the Lemma A.1. Below are all the threshold algorithms that we consider (again, we only make the call outs if the actual bandwidth is available).

1. **Th-Random:** chooses a random ordering.
2. **Th-MaxRemBand:** orders on maximum remaining bandwidth.
3. **Th-Prob:** orders on (decreasing) survival probability.
4. **Th-LP:** This is the LP based algorithm discussed above.

7.1 The Simulation Setup

We simulate the above algorithms using synthetic data generated from specific natural distributions, and vary parameters in the data generation such as impression arrival rate and range of minimum prices of impressions.

Implementation: There are 32 ad networks, and the bandwidth of each ad network is implemented as a *token bucket* or a buffer. Tokens get generated in a bucket at a uniform rate, which reflects the bandwidth, and an attempt to call out to an ad network succeeds if and only if there is at least one

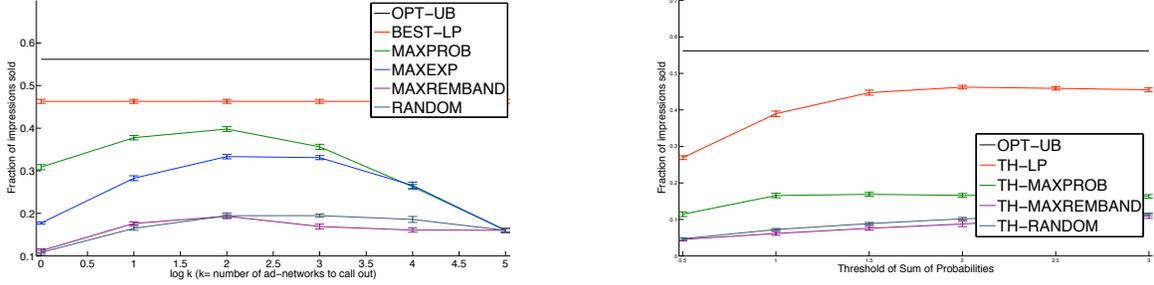


Figure 1: Set based and Threshold algorithms for Gaussian bid distributions

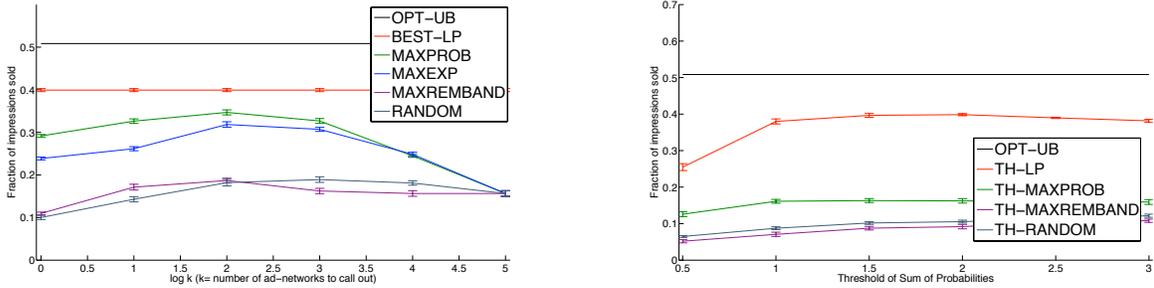


Figure 2: Set-based and Threshold algorithms for Pareto bid distributions

token, which is consumed by the call out, in the bucket corresponding to the ad network. Moreover, a token bucket has a limit on the number of tokens it can store, and if the bucket is full, tokens generated at that time are lost. This reflects the burst-size allowed in the communication. Unless specified otherwise we use a bucket size of 5 for most of our simulation results (see section 7.4). The ad networks have (continuous) token generation rates chosen uniformly at random between 5 and 50 per unit time. Impressions arrive according to a Poisson clock with a fixed rate. This rate is the expected time lapse between two consecutive impressions. Again, unless specified otherwise, we use a rate of 0.003 for the Poisson clock. The average token generation rate per impression, ρ_i , varies from 0.015 to 0.15 when the rate of the Poisson clock is 0.003.

Bid distribution: All bids are drawn from bounded distributions, which put all its probability mass in the range $[0, R]$ for a fixed positive value R . Gaussian is a very commonly seen distribution, with exponential decay, while Pareto is a heavy-tailed distribution which has polynomial decay (*power law*), and is also often observed in online ad scenarios. The means of these Gaussian or Pareto distributions are chosen uniformly at random from the range $[0, 0.5R]$. The Gaussian distributions are then given a standard deviation uniformly between 0 and 0.5 times the mean, while the degree of the polynomial pdf of Pareto distributions is uniformly chosen between 2 and 5. Subsequent to their choice, all distributions were truncated to $[0, R]$ – note that the truncation affects the expectation and survival probabilities.

Verticals and Minimum Prices: Different types of impressions should have different value to an ad network, so we have 10 different types of impressions, which we refer to as *verticals*. Each arriving impression is assigned a vertical uniformly at random. The bid distribution of each ad

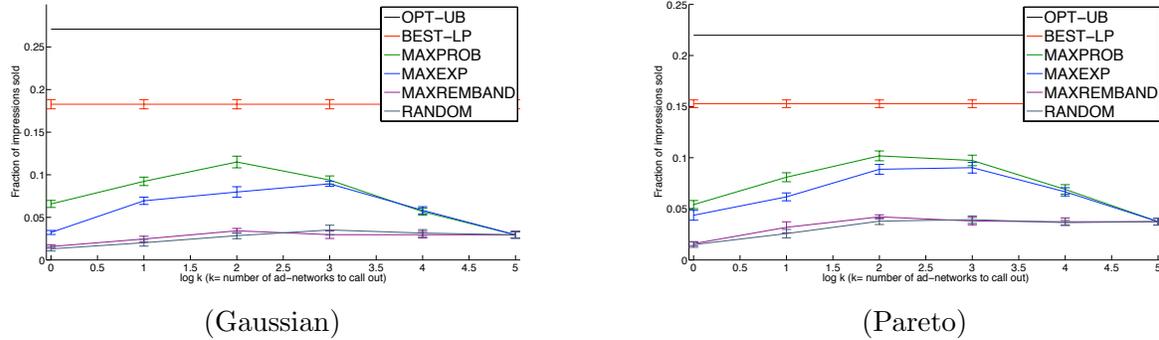


Figure 3: Set-based algorithms for bid distributions with higher minimum prices

network for an impressions depends on its vertical only. However, the survival probabilities can still vary for impressions from the same vertical, because their minimum prices are independently chosen. Thus, impressions are iid drawn from a fixed distribution. 2000 impressions arrive in each simulation, while our LP-based algorithm is given 500 impressions to learn from the same distribution of impressions.

The minimum prices of each impression come from the range $[0.2R, R]$ for most of the simulations, but we also verified our broad findings when the range of minimum prices is $[0.5R, R]$. Higher minimum prices reflect lower levels of maximum possible sales.

For our set-based algorithms, we try out all powers of 2 for the value of k , that is, 1, 2, 4, 8, 16 and 32. For our threshold algorithms, we try 0.5, 1.0, 1.5, 2.0, 2.5 and 3.0 as our threshold values. This completes the description of the simulation set up, and we are now ready to describe our findings.

7.2 Performance of various Strategies

We tried our algorithms for 2000 impressions in each simulation. While this stream may seem small, we checked the standard deviation of our results when different streams drawn from the same distribution were used, and found them to be sufficiently small. With 10 different such streams, the standard deviation was found to be about 2% or less for all the algorithms. The performance (along with error bars indicating the deviation) of the various strategies are plotted in Figures 1, 2 and 3. While the performance of threshold algorithms are plotted against the threshold parameter, that of the set based algorithms are plotted against k on a logarithmic scale, that is, $\log k$. BEST-LP indicates the performance of our LP-based algorithm for optimal choice of threshold, and is shown along with set based algorithms for comparison. OPT-UB indicates the upper bound on an optimal policy that is given by LP1 in the previous section. Note that the real offline optimum can be significantly smaller, but is hard to compute. The following conclusions emerge:

1. The LP-based algorithm surpasses the performance of the other algorithms by a large margin, about 20% more than the closest performer with optimal parameter choice. This gap grows to 85% when the reserve prices are raised, and the performance of all the algorithms (and also the optimum) falls. Thus, our LP-based algorithm is even more valuable in a world where

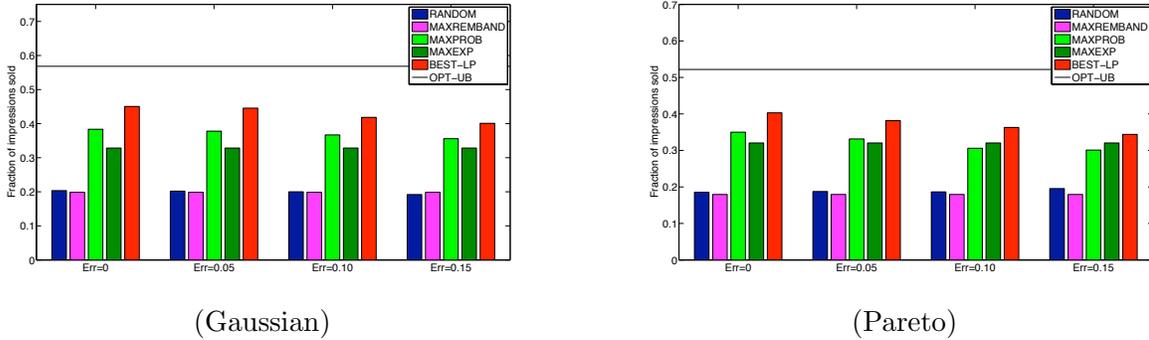


Figure 4: Set-based algorithms for different bid distributions with erroneous bid estimations

fewer items can be sold. Figure 3, show the effect of raising the minimum prices for set based algorithms.

2. The algorithms Random and MaxRemBand perform quite poorly compared to all other algorithms that use information about bid distributions. This shows that *bid estimation is useful*.
3. The performance of MaxExp was better than Random and MaxRemBand, but was worse than MaxProb. Thus for the distribution used in the simulation, information about expectation was useful, but less useful than the survival probability of the bids.
4. The other threshold based algorithms perform quite poorly. In the rest of the experimental evaluation, we will ignore these algorithms.

7.3 Robustness and Error estimation

Estimating survival probabilities can be a tough problem, and can be loosely compared to the problem of estimating click-through-rates. It is thus very likely that the estimates will be inaccurate, and it is very important to understand how useful such erroneous estimates are. We add noises to the survival probabilities that are normally distributed (round to 0 or 1 if the resulting value is below 0 or above 1, respectively), and have standard deviations ranging from 0.05 to 0.15. We find that the algorithms using this information lose some performance, but still perform consistently better than the algorithms without any information about bid distributions. We show the plots for both distributions, in Figure 4. For comparison, we choose to represent an algorithm by its peak performance, that is, its performance for optimal choice of parameters.

7.4 Sensitivity analysis

For all the plots described above, the size of the token buckets were fixed at 5 and the arrival rate was fixed to 0.003. We tried out bucket sizes of 2, 5, 15, and 45. As expected, the performance of each algorithm is greater when bucket size is larger; but there was no substantial differences in the conclusions we observed. In fact, the performance of the algorithms were within a standard deviation of a different bucket size (we omit showing them to avoid clutter) There was no difference

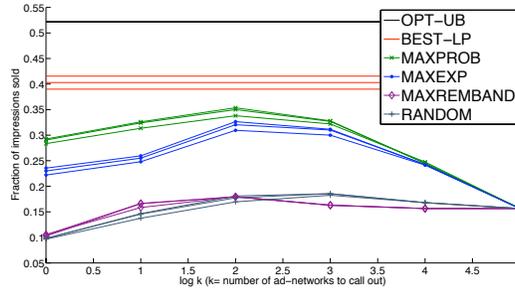


Figure 5: Set based algorithms for Pareto bid distributions with burst sizes 2, 5 and 15

in performance between bucket sizes 15 and 45. We superimpose the performance plots of bucket sizes 2, 5, 15 and show this plot for Pareto distributions in Figure 5. We also varied impression arrival rate to note its effect on our algorithms. We found, as expected, that performance improves if arrival rate slows down, but the relative performances of the algorithms remain unchanged. Figure 6 shows this for the peak performances of the algorithms, for Pareto bid distributions.

The plots for Gaussian distributions and threshold algorithms are similar in both cases.

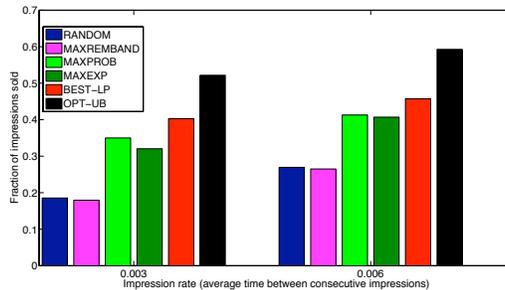


Figure 6: Peak performances of algorithms for different impression arrival rates for Pareto bid distributions

8 Conclusion & Future Work

We initiate a formal study of bandwidth and resource constraints faced by ad exchanges and ad networks as they move into the arena of real time bid solicitation. Our conceptual framework of learning bid distribution followed by online decision making should be useful for further study of real time bidding systems. Our online algorithms are fast and can be implemented to run in real time within this conceptual framework. The technical framework of our algorithm, with a short learning phase followed by a long online phase, can generally be applied to any problem that involves maximizing linear objectives, or objectives that can be approximated by a concave or linear objective, with linear constraints.

References

- [1] Y. Azar, B. E. Birnbaum, A. R. Karlin, and C. T. Nguyen. On revenue maximization in second-price ad auctions. *Proc. of ESA*, pages 155–166, 2009.
- [2] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. Knapsack secretary problem with applications. *Proceedings of Approx*, pages 16–28, 2007.
- [3] R. E. Barlow, A. W. Marshall, and F. Proschan. Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics*, 34(2):375–389, 1963.
- [4] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. *STOC*, also at <http://arxiv.org/abs/0907.4166>, 2010.
- [5] N. Buchbinder, K. Jain, and J. Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. *Proc. of ESA*, pages 253–264, 2007.
- [6] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint (extended abstract). *IPCO*, pages 182–196, 2007.
- [7] D. Chakrabarty and G. Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and gap. *Proceedings of FOCS*, pages 687–696, 2008.
- [8] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. *STOC*, also at <http://arxiv.org/abs/0907.2435>, 2010.
- [9] N. Devanur and T. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proc. 10th Annual ACM Conference on Electronic Commerce (EC)*, 2009.
- [10] A. Ghosh, P. McAfee, K. Papineni, and S. Vassilvitskii. Bidding for representative allocations for display advertising. In *Proc. of 5th Workshop on Internet and Economics(WINE)*, 2009.
- [11] A. Ghosh, B. I. P. Rubinstein, S. Vassilvitskii, and M. Zinkevich. Adaptive bidding for display advertising. In *WWW*, pages 251–260, 2009.
- [12] S. Guha and A. McGregor. Approximate quantiles and the order of the stream. *Proc. of PODS*, pages 273–279, 2006.
- [13] S. Guha and K. Munagala. Adaptive uncertainty resolution in bayesian combinatorial optimization problems. *ACM Transactions of Algorithms*, (to appear).
- [14] J. Kleinberg, Y. Rabani, and E. Tardos. Allocating bandwidth for bursty connections. *SIAM J. Of Computing*, 30(1):171–217, 2001.
- [15] R. Kleinberg. A multiple-choice secretary problem with applications to online auctions. *Proceedings of SODA*, pages 630–631, 2005.
- [16] R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *STOC*, pages 681–690, 2008.

- [17] A. Kulik, H. Shachnai, and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. *SODA*, pages 545–554, 2009.
- [18] A. Marchetti-Spaccamela and C. Vercellis. Stochastic on-line knapsack problems. *Mathematical Programming*, 68:73–104, 1995.
- [19] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.
- [20] S. Muthukrishnan. Adexchange: Research issues. In *Proc. of 5th Workshop on Internet and Economics(WINE)*, 2009.
- [21] R. B. Myerson. Optimal auction design. *Math. of OR*, 6(1):5873, 1981.
- [22] M. Nolet. <http://www.mikeonads.com/2009/08/30/rtb-part-i-what-is-it/>.
- [23] S. Pandey and C. Olston. Handling advertisements of unknown quality in search advertising. *Proceedings of NIPS*, pages 1065–1072, 2006.
- [24] P. release. <http://www.doubleclick.com/insight/blog/archives/doubleclick-advertising-exchange/announcing-the-new-doubleclick-ad-exchange.html>.
- [25] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [26] A. S. Tanenbaum. *Computer Networks, 3rd Edition*. Prentice-Hall, 1996.
- [27] E. Vee, S. Vassilvitskii, and J. Shanmugasundaram. Optimal online assignment with forecasts. In *EC*, 2010.

A The Adversarial Model

We show that no online algorithm, deterministic or randomized, can give any approximation to the offline optimum, even when the instance has only one ad network, and the bids are 0 or 1.

Theorem 6 *For any given $0 < \alpha \leq 1$, there exists no randomized algorithm which is α -approximate, if the length of the sequence is arbitrarily large.*

Proof: Maximizing the sales objective in the presence of only one ad network subsumes the online knapsack problem (the value of an item is the probability that the bid exceeds minimum price) with unit size of items, for which no algorithm can give better than $1/\Omega(\log m)$ approximation [18]. If m is allowed to be large, then no algorithm giving approximation independent of m can exist. \square

One key feature of the adversarially constructed sequences in the theorem above is that OPT can be arbitrarily small compared to the length of the sequence. This is unlikely to happen; moreover, if it were to happen, then there would not be much benefit out of constructing the real time bidding system. Thus we make the assumption that OPT is reasonably large. But even with that assumption, we don't fare particularly well. Following is a simple observation about sales.

Lemma A.1 *For a given impression j , suppose ad network i , with survival probability p_{ij} is called out with probability x_{ij} . Then the probability of selling the impression is $1 - \prod_{i=1}^n (1 - p_{ij}x_{ij})$. If $\sum_{i=1}^n p_{ij}x_{ij} = c$, then the probability of selling the impression is at most c and at least $1 - \frac{1}{e^c}$.*

Proof: The probability of selling the impression is $1 - \prod_{i=1}^n (1 - p_{ij}x_{ij})$. It is upper bounded by $\sum_i p_{ij}x_{ij} = c$. The probability is lower bounded by $1 - \prod_{i=1}^n (1 - c/n)^n$ where n is the number of ad-networks, and this is lower bounded by $1 - \frac{1}{e^c}$. \square

Theorem 7 *Suppose it is promised that $\text{OPT} \geq \delta m$, where $\delta > 0$ is known to the algorithm. Then, no algorithm can be better than $1/\Omega(\log(1/\delta))$ -approximate for the sales or total value objective. Moreover, a simple randomized algorithm is $1/O(\log(n/\delta))$ -approximate for the sales objective.*

Proof: The hardness instances for online knapsack problems also imply that if the value of all items are guaranteed to lie between L and U , $L < U$, then no algorithm can give better than $1/\Omega(\log(U/L))$ approximation [18], if m is allowed to be arbitrarily large. The promise $\text{OPT} \geq \delta m$ allows us to construct any instance where values of items range from δ to 1, thus yielding the lower bound. For the upper bound on sales, we use the following simple algorithm: Let us assume, wlog, that δ is a negative power of 2. Choose a random cut-off t from the set $H = \{\delta/2, \delta, 2\delta, 4\delta \dots 1\}$. Now send each impression to any $2/t$ ad networks who has survival probability between t and $2t$ and has bandwidth remaining. If there are less than $2/t$ such qualifying ad networks for an impression, then send it to all qualifying ad networks. This algorithm gives a $1/O(\log(n/\delta))$ approximation; the proof is standard. \square