

Budget Smoothing for Internet Ad Auctions: A Game Theoretic Approach

DENIS CHARLES, Microsoft Research, Redmond

DEEPARNAB CHAKRABARTY, Microsoft Research, India

MAX CHICKERING, Microsoft Research, Redmond

NIKHIL R. DEVANUR, Microsoft Research, Redmond

LEI WANG, Microsoft adCenter, Bellevue

In Internet ad auctions, search engines often throttle budget constrained advertisers so as to spread their spends across the specified time period. Such policies are known as *budget smoothing* policies. In this paper, we perform a principled, game-theoretic study of what the outcome of an ideal budget smoothing algorithm should be. In particular, we propose the notion of *regret-free* budget smoothing policies whose outcomes throttle each advertiser optimally, given the participation of the other advertisers. We show that regret-free budget smoothing policies always exist, and in the case of single slot auctions we can give a polynomial time smoothing algorithm. Inspired by the existence proof, we design a heuristic for budget smoothing which performs considerably better than existing benchmark heuristics.

Categories and Subject Descriptors: K.4.4 [Computers and Society]: Electronic Commerce

General Terms: Design, Algorithms, Economics

Additional Key Words and Phrases: Budget Smoothing, Ad Auctions, Online Algorithms, Equilibrium

1. INTRODUCTION

Advertising is the main source of revenue for search engines such as Google and Bing. For every search query, these search engines run an auction using the advertisers' keyword bids to determine which ads are shown alongside search results. Since the search traffic is unknown ahead of time, and since prices are determined by unknown and changing competitor bids, the amount of money that an advertiser can spend is uncertain and can vary significantly. Because of this, search engines allow advertisers to provide *budget constraints* that cap their spends over some period of time.

To respect budget constraints, search engines implement policies that “throttle out” advertisers from some of the auctions; these policies are commonly referred to as *budget smoothing* policies. A naïve budget smoothing policy, for example, is to throttle out an advertiser only when the budget runs out. With this policy, the market is thick at

Author's addresses: Denis Charles, Max Chickering and Nikhil R. Devanur, Microsoft Research, Redmond; Deeparnab Chakrabarty, Microsoft Research, Bangalore, India; Lei Wang, Microsoft adCenter, Bellevue.

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

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

the start of the day¹ and subsequently thins out as the day progresses; high-bidding advertisers get throttled out early, which can be unprofitable for both the advertisers and the search engine company.

Following are three desiderata of a good budget smoothing policy.

- . **High welfare:** The welfare of an advertiser is the value² derived over all auctions. The budget smoothing policy should attain high social welfare (i.e., the sum of all advertisers' welfare) because delivering more value can lead advertisers to spend more money, thereby increasing the long-term revenue of the search-engine company. While the total welfare of all the advertisers is a natural objective, care must be taken that each individual advertiser's welfare is kept reasonably high.
- . **No large short-term revenue loss:** As noted above, we expect to increase long-term revenue by increasing advertiser welfare, but we would also like the budget smoothing policy to not significantly decrease the short-term revenue.
- . **Even smoothing:** Advertisers have come to expect budget smoothing policies to spread budgets evenly over the course of the day; although it is not always clear why this would be desirable, it sometimes can give the advertiser a more representative sample of the day's auctions.

The problem of budget smoothing has been the source of many theoretical investigations, most notably in the form of the so-called *Adwords* problem introduced by [Mehta et al. 2005]. This problem is a simplified version of the budget smoothing problem that ignores many of the real-world constraints including the presence of multiple slots and the need to run a GSP auction to determine the prices. To address this, the Adwords problem has been generalized, and one example is the class of resource allocation problems considered by [Feldman et al. 2010; Devanur et al. 2011]. The budget smoothing problem does indeed fall in this class: for every search query, the algorithm needs to pick a *slate* of advertisers, i.e., the subset of advertisers that participate in the auction. The particular slate chosen determines the price paid by each advertiser in that auction. Given an objective of (e.g.) social welfare, the papers provide algorithms that achieve a near optimal guarantee under reasonable stochastic assumptions.

Unfortunately, an optimal solution to such a social-welfare optimization problem may not be good for the other desiderata discussed earlier. Suppose, for example, there are 1000 copies of the same auction, with two advertisers A and B. Further suppose that advertiser A bids 2, advertiser B bids 1, there is a reserve price of 0.1, the budget of advertiser A is 100 and the budget of advertiser B is 10000. The welfare maximizing solution throttles advertiser B out of all the auctions, letting advertiser A win them all for the reserve price of 0.1. This solution is terrible in terms of advertiser B's welfare, and it is worse in terms of revenue when compared to the naive policy or other baseline policies that we consider later.

One could replace the objective of social welfare with revenue, but because we care about keeping welfare high, this is not necessarily a desirable choice. One could try a linear combination of welfare and revenue, and so on, but it is not clear if there is some objective function that is appropriate for the problem. This leads to important

¹ We refer to a day as the duration for which the budget constraint is enforced. In practice, the budget constraint could be for the duration of a day, a week, a month, etc.

²In our experiments, we use the advertiser bids as approximations to their per-auction values; GSP is not a "truthful auction". Sometimes the true values can be deduced from the bids in GSP, in which case it is possible to use these true values instead.

conceptual questions. What is an ideal solution to the budget smoothing problem, even in hindsight when all the information about advertisers and auctions is known? Can we formally characterize an ideal solution in terms of a few key properties?

The main contribution of this paper is the introduction of a novel solution concept for the budget smoothing problem. This solution concept is game theoretic in nature rather than a solution to an optimization problem.

The outcome of any smoothing algorithm defines a *participation profile*, that is, the slate of participating advertisers for each auction. We introduce the notion of a *regret-free* profile. When the number of slots equals³ the number of advertisers, the regret-free profile is equivalent to a Nash equilibrium of a particular game. When the number of slots is fewer, the regret-free profile is a refinement of a Nash equilibrium that makes the equilibrium more robust. This refinement is similar to, but different from, related concepts such as trembling-hand perfect equilibrium. Consider the game with advertisers as players where the strategy of an advertiser is the choice of a subset of auctions to participate in. In a Nash equilibrium of this game, for any advertiser, given the participation of every other advertiser, switching to a different set of auctions does not improve his own utility.

When the number of slots is fewer than the number of advertisers, an advertiser who participates in an auction may not win any slot. In such a case, the Nash constraint says nothing about whether the advertiser should participate in that auction or not, since it does not affect his own utility. However his participation may affect the price paid by the other advertisers. This freedom may lead to “delicate” equilibria where some advertisers are made to participate in auctions that they would not participate in if there were an extra dummy slot with very low, but positive, click through rate. We define our regret-free notion so that such participation profiles are disallowed, thereby keeping only advertisers who pose “credible threats”. Therefore, our regret-free notion is a refinement of Nash equilibria to make it more robust.

The best-response in the game we consider has a simple structure which is summarized by the following two observations:

- For a given auction and advertiser, the value-to-price ratio for winning a slot only depends on the valuation of the given advertiser and the next highest value, but is independent of *which* slot is won. We call this ratio as the *ROI*⁴.
- The optimum solution to a fractional knapsack problem has a simple structure: given prices and values for a set of (divisible) items, we identify the (fractional) subset of items that maximizes the total value subject to a budget constraint on the prices by picking the items in decreasing order of ROI, until the sum of the prices of the selected items exhausts the budget.

For an advertiser i , the participation of all other advertisers fixes the ROI for i in every auction. A best response mimics the optimal solution to the fractional knapsack problem; i.e., we pick the auctions for i in decreasing order of ROI until the budget of i is exhausted⁵. For auctions in which the advertiser does not win any slot, his ROI is zero. Our robustness concept calls for an extension of ROI that includes extra dummy slots; we call this extension the *potential ROI* or *pROI* for short. This is the ratio of the

³Or exceeds, because excess slots have no impact.

⁴Return on Investment.

⁵A fractional solution to the knapsack can be thought of as a randomized strategy. We only impose the budget constraint in expectation. In practice the price of a single auction is usually much smaller than the budget, and therefore an overflow in the last auction is not a concern.

valuation of the given advertiser to the next highest value (or the reserve price), even for advertisers that don't win any slot. The *robust* best response mimics the optimal solution to the fractional knapsack problem with pROIs instead of ROIs. A participation profile is regret-free if the participation set of each advertiser is simultaneously a robust best response.

The concept of a regret-free outcome is also very similar to a market equilibrium in the Fisher market model. With a GSP auction, the price faced by an advertiser is governed by the next-highest advertiser and cannot be set independent of the allocation. Instead, suppose that each slot in each auction had a posted-price. Then a regret-free outcome would essentially be equivalent to an equilibrium in a Fisher market, with the slots as items and advertisers as agents. Nonetheless this difference about whether the price is exogenous (posted-price) or endogenous (GSP price) is significant, and it's not clear if techniques from computing market equilibria are useful in our setting.

In addition to introducing the concept of regret-free participation profiles as a solution concept for the budget smoothing problem, we provide the following results:

- We show that a regret-free randomized participation profile always exists. The proof relies on Brouwer's fixed point theorem, and is surprisingly much simpler than other existence proofs of equilibria such as Nash equilibria and market equilibria.
- For the single slot case, we give a polynomial time algorithm to find a (weaker notion of a) regret-free participation profile. This algorithm is a variant of the deferred acceptance algorithm of [Gale and Shapley 1962] for finding stable matchings.
- We give a heuristic for the online version of the problem where the auctions are not known beforehand. The heuristic is inspired by the algorithm of [Devanur et al. 2011] for general resource allocation optimization problems in a stochastic setting. The heuristic maintains a threshold ROI for each advertiser that must be met in order for the advertiser to participate. Each threshold is continuously updated based on how fast the corresponding budget is being spent, so that a steady state set of thresholds corresponds to a regret-free participation profile.
- We perform extensive experimentation of this heuristic on real-world data. We provide evidence that the heuristic is an attractive option for a budget smoothing policy.

Experimental evaluation

We compare the performance of our heuristic against a benchmark *random throttling* heuristic that independently throttles each advertiser such that the rate of spend is spread uniformly across time. We run our experiments on samples of a few million auctions drawn from a set of 200 micromarkets⁶ over a 10-day period. We scale the budget in different ways to get different data sets.

We find that our heuristic consistently yields higher advertiser value while maintaining similar short-term revenue. We also consider a *regret measure* that measures the difference in the sum of per-auction ROI values obtained from both heuristics compared to the best participation in hindsight. We find that our algorithm's regret is much smaller than that of random throttling; that is, the online heuristic is doing what it is supposed to do. Finally, we find that the exhaust time (fraction of budget period at which an advertiser's budget is exhausted) is close to 1. Therefore the budget is indeed smoothed out as desired.

⁶A *micromarket* is a dense, well separated cluster of advertisers and auctions within the Bing advertiser/auction graph.

We also consider the effects of our heuristic versus random throttling on individual advertisers. We find that most advertisers get higher value with our heuristic, with $< 10\%$ of them having a decrease. Furthermore, $> 50\%$ of the advertisers had at least a 5% increase in value. The difference in spend between the two heuristics is typically small: 97% of advertisers had both spends within 5% of their budgets. Finally, of the top 10 advertisers by budget, none had a lower spend using our heuristic. Three of them has spend higher by $> 20\%$ and 5 of them had value higher by $> 5\%$.

Conclusions: Our heuristic results in an overall positive boost to social welfare while remaining neutral with respect to revenue. The effect on individual advertisers is similar to the aggregate effect: mostly positive in value and mostly neutral in spend.

Related Work

As noted earlier, the first theoretical investigation of budget smoothing was performed by [Mehta et al. 2005] as the Adwords problem. This paper was one of the first in the regime of Internet ad auctions, and since then the field has seen a great outpouring of research. The winner determination and pricing algorithm in [Mehta et al. 2005], however, was not GSP. [Abrams et al. 2007] investigated multiple slot auctions with GSP pricing and exhibited a linear program that maximizes search engine revenue when the advertiser bids and budgets are known. Subsequently, [Goel et al. 2010] gave a constant competitive *online* algorithm; they also show that the [Mehta et al. 2005] algorithm, and the greedy algorithm, performs poorly in this setting. None of these papers, however, consider the advertiser point of view that we take in this paper.

Another very rich line of work was the study of the Adwords problem in the stochastic setting. This was first studied by [Mahdian et al. 2012] in the partial information setting, and by [Devanur and Hayes 2009] in the random permutation setting. This was followed by a series of papers, [Feldman et al. 2009, 2010; Devanur et al. 2011; Haeupler et al. 2011; Mirrokni et al. 2012; Devanur et al. 2012], both for the Adwords problem and more general packing resource allocation frameworks in the stochastic setting. In fact, the heuristic in the experimental section of the paper is inspired by the algorithm in [Devanur et al. 2011].

There has also been work on bid optimization, where given the GSP auction and the budget constraints, advertisers try to optimize their bids for various keywords. Most approaches, such as [Feldman et al. 2007; Chakrabarty et al. 2008; Amin et al. 2012], cast this as a (multiple / multiple choice) knapsack problem and use offline/online algorithms to find efficient bids. On the other hand, in our work, we consider the bids to be immutable and rather change the participation profiles to increase welfare.

Finally, although GSP (with minor variants) is the auction of choice in practice today, it is known that the auction is not truthful, see [Varian 2006; Edelman et al. 2007; Agarwal et al. 2006]. In light of this, [Lame and Tardos 2010] recently studied the price of anarchy of a certain natural game induced by GSP, and showed that it is bounded by a constant. However, these papers consider just one auction in isolation as opposed to considering multiple auctions that are tied together by budget constraints.

Organization: We state our model, definitions and results formally in Section 2. The theoretical and experimental results are presented in Sections 3 and 4 respectively. Concluding remarks are in Section 5.

2. FORMAL STATEMENT OF MODEL AND RESULTS

We operate in the ad-auction setting with N advertisers and M auctions. Advertiser $i \in [N]$ specifies a budget B_i for participating in these auctions, and a bid b_{ij} for the j th auction, with j ranging from 1 to M . We denote the relevance of the ad by advertiser i for the auction j as p_{ij} , a parameter between 0 and 1. We let $v_{ij} := b_{ij}p_{ij}$ denote the *rank-score* of advertiser i in auction j . We assume each auction has K ad-slots in each auction. We associate a slot click-through-rate (slot CTR) $\theta_\ell \in [0, 1]$ with the ℓ th slot, and for simplicity we assume θ_ℓ 's are non-increasing in ℓ . Finally, each auction j is associated with a reserve price r_j .

Each auction j is a generalized second price auction (GSP) on a *subset* $S_j \subseteq [N]$ of the advertisers that *participate* in the auction. This set could be smaller than $[N]$ due to the budget smoothing policy (or throttling) of the search engine. The crux of this paper is to arrive at this subset of participating advertisers in a systematic fashion, as discussed in the introduction. In the subsequent paragraphs we put this into a formal framework.

Given a participation set S_j , the GSP auction sorts the advertisers in S_j in the order of decreasing rank-score. Suppose this ordering is

$$v_{1j} > v_{2j} > \dots > v_{\ell j} > r_j > v_{(\ell+1)j} > \dots > v_{tj}$$

where $t := |S_j|$, and we have renamed the advertisers in S_j for notational convenience. If $\ell \leq K$, then the set $\{1, \dots, \ell\}$ are the *winners* of this auction; otherwise, the set of winners is $\{1, \dots, K\}$. The cost per click (CPC) of a winner i is precisely the minimum bid required to win the current position. Therefore,

$$\text{CPC}(i) := \frac{\max(v_{(i+1)j}, r_j)}{p_{ij}}$$

We assume that the ad of advertiser i shown on slot ℓ in the auction j gets a click with probability equal to the product of the relevance and the slot CTR. Thus the effective CTR is given by $\text{CTR}(i, j, \ell) := p_{ij} \cdot \theta_\ell$. Therefore the expected spend of advertiser i on winning the ℓ th slot of auction j is

$$\text{spend}(i, j) := \frac{\max(v_{(i+1)j}, r_j)}{p_{ij}} \cdot (p_{ij}\theta_\ell) = \theta_\ell \max(v_{(i+1)j}, r_j)$$

We assume that the bids made by the advertisers are a proxy for the utility they derive from getting a single click. Therefore, the utility obtained by advertiser i in winning ℓ th slot of auction j is

$$\text{util}(i, j) = b_{ij} \cdot (p_{ij}\theta_\ell) = v_{ij}\theta_\ell$$

We now define one of our central driving concepts: the return on investment (ROI) of advertiser i in auction j . It is defined, quite simply, as the ratio of $\text{util}(i, j)$ and $\text{spend}(i, j)$ and is the bang-per-buck the advertiser obtains by winning a slot in auction j .

$$\text{ROI}(i, S_j) := \text{util}(i, j) / \text{spend}(i, j) = \frac{v_{ij}}{\max(v_{(i+1)j}, r_j)}.$$

$\text{ROI}(i, S_j)$ is zero if i does not win any slot. As discussed in the introduction, we consider a robust version of ROI, which we call pROI, defined as

$$\text{pROI}(i, S_j) := \frac{v_{ij}}{\max(v_{(i+1)j}, r_j)},$$

even if i does not win any slot. ROI and pROI differ only for such non-winning advertisers; for them, ROI is zero and pROI is non-zero.

Regret-free Participation Profiles. A budget smoothing algorithm chooses a participation set S_j for all auctions $j \in [M]$. Once this set is chosen, GSP is run as described earlier. We call the set $\{S_1, \dots, S_M\}$ the *participation profile* generated by the smoothing algorithm.

We say that an advertiser i has been *throttled* from auction j if $i \notin S_j$. We let $\text{spend}(i)$ denote the total spend of the advertiser, thus $\text{spend}(i) := \sum_{j \in [M]} \text{spend}(i, j)$. Finally, we define $\text{minROI}(i)$ as his minimum pROI among those auctions in which he participates. We stress that the minimum is over the *potential* ROIs and not the actual ROIs.

$$\text{minROI}(i) := \min_{j: i \in S_j} \text{pROI}(i, S_j).$$

We call a participation profile (S_1, \dots, S_M) **regret-free** if for each advertiser, the auctions in which he participates is a robust best response to the participation profile of all other advertisers. This translates into the following three conditions: (1) no advertiser overspends, (2) no advertiser with budget remaining is ever throttled, and (3) if an advertiser is ever throttled from an auction (and thus necessarily has spent his budget), then his pROI in this auction is no larger than his minROI.

Regret-free participation profiles may not always exist due to “knapsack issues”. Consider two auctions with one slot where one bidder bids \$5 on both with a total budget of \$7, while the second highest bidder in both cases bids \$4 with a very large budget. If the highest bidder participates in both auctions, feasibility is violated, otherwise maximal participation is violated.

To fix this, we move to *randomized* participation profiles. For each auction j instead of a single participation set S_j , we now have a distribution \mathcal{S}_j on participation sets. Each advertiser’s spend is now a random variable, and the relevant quantity is $\text{Exp}[\text{spend}(i)]$ where the expectation is over the randomization in the participation sets. We extend the definition of $\text{minROI}(i)$ in the strongest possible way:

$$\text{minROI}(i) := \min_{S_j \in \text{supp}(\mathcal{S}_j): i \in S_j} \text{pROI}(i, S_j)$$

where $\text{supp}(\mathcal{S}_j)$ is the support of \mathcal{S}_j .

Definition 2.1 (Main Definition). A randomized participation profile $(\mathcal{S}_1, \dots, \mathcal{S}_M)$ is regret-free if

- (1) *Feasibility:* For each advertiser i , $\text{Exp}[\text{spend}(i)] \leq B_i$.
- (2) *Maximal Participation:* If $\text{Exp}[\text{spend}(i)] < B_i$, then for all j , $\forall S \in \text{supp}(\mathcal{S}_j)$, $i \in S$.
- (3) *Optimal Throttling:* If $i \notin S_j$ for some $S_j \in \text{supp}(\mathcal{S}_j)$, then $\text{pROI}(i, S_j \cup i) \leq \text{minROI}(i)$.

To go back to the example mentioned above, if advertiser 1 is throttled from auction 2 with 60% probability, then the expected spend is \$7, and the optimal throttling condition remains satisfied due to symmetry.

Remark 2.2. Even though we allow arbitrary randomization, at an equilibrium the randomization takes a very simple form. The only reason to randomize is to handle the knapsack issue: for a given advertiser, when considering the auctions in the decreasing order of ROI, there is exactly one which takes the spend over the budget. In order to have the spend be exactly equal to the budget, the advertiser participates in this “straddling auction” with an appropriate probability. This holds even if there are multiple auctions with the same ROI; they can be considered in an arbitrary order to find

the one that straddles the budget. This randomization is only needed for theoretical niceness. In practice, the price in any auction is small compared to the budget therefore for the knapsack problem that captures an advertiser’s best response, the integral solution given by the greedy algorithm is very close to the fractional optimum. The fact that randomized strategies are not crucial is also reflected in our heuristic which only uses pure strategies.

Remark 2.3. Recall that $\text{spend}(i, j)$ itself is an expectation, taken over the click probabilities. Therefore the budget constraint we require is only on the expected spend whereas in reality the budget constraint is on the actual spend. However, the small bid to budget ratio comes to the rescue once again. We assume that the clicks on different auctions are independent events which implies that the actual spend is highly concentrated around the expected spend, as can be seen by any of the concentration inequalities such as the Chernoff-Hoeffding bounds. Therefore the constraint on the expected spend is a good enough approximation.

Our main theorem is the following.

THEOREM 2.4. *For any bids, budgets, auctions and slots, there always exists a regret-free randomized participation profile.*

The proof of the above theorem is via Brouwer’s fixed point theorem and is non-constructive. Nevertheless, the proof technique, which we describe in detail in §3.1, leads to a heuristic that performs well in our experiments; details can be found in §4. Furthermore, unlike the situation in general equilibrium theory where existence of equilibria does not necessarily imply finite time algorithms, we observe that existence of regret-free randomized participation profiles implies finite time algorithms.

COROLLARY 2.5. *The distribution in any regret-free randomized participation profile has probabilities that are rational numbers and can be evaluated in finite time.*

A theoretical question left open by our work is whether there exist polynomial (in N, M, K) time algorithms to find regret-free randomized participation profiles. In fact, we do not know such algorithms even in the case when M and K are constants.

Our second result states that in the case of a single slot, we can find *deterministic* participation profiles that are ‘almost’ regret-free. In particular, we relax the requirement about the losing advertisers being a credible threat. Also, the feasibility criteria is violated by at most one bid per advertiser. We define $\text{minROI}'(i)$ to be the minimum ROI among all the auctions that i wins.

$$\text{minROI}'(i) := \min_{j: \text{spend}(i, j) > 0} \text{ROI}(i, S_j).$$

THEOREM 2.6. *In the case of single slots, there exists a polynomial time algorithm which given an auction instance finds a deterministic participation profile (S_1, \dots, S_M) such that*

- (1) **Feasibility:** *For each advertiser i , $\text{spend}(i) \leq B_i + v_{max}^i$, where $v_{max}^i = \max_j v_{ij}$*
- (2) **Maximal Participation:** *If $\text{spend}(i) < B_i$, then $i \in S_j$ for all j .*
- (3) **Optimal Throttling:** *If $i \notin S_j$, then $\text{ROI}(i, S_j \cup i) \leq \text{minROI}'(i)$.*

3. THEORETICAL RESULTS

3.1. Regret-Free Randomized Participation Profiles always exist.

We start by proving Theorem 2.4. At a very high level, we maintain a parameter, which we call a threshold ROI, for each advertiser. These parameters determine in which auctions the advertisers participate, ensuring the optimal throttling condition. The goal is then to find the parameters which also maintain feasibility and maximum participation; the existence of this is proved via a fixed point theorem.

Proof of Theorem 2.4: We associate with each advertiser i a parameter $\alpha_i \geq 1$ which denotes the minROI of this advertiser. In particular, we guarantee that in any of the auctions that i participates, his pROI is at least α_i . Furthermore, in any of the auctions i doesn't participate, we ensure that his pROI if he had participated is at most α_i .

Definition 3.1. Given $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_N)$, and an auction j , we say that a participation set S_j is α -respecting if for all advertisers i ,

If $i \in S_j$, $\text{pROI}(i, S_j) \geq \alpha_i$.
 If $i \notin S_j$, then $\text{pROI}(i, S_j \cup i) < \alpha_i$.

Next we show a procedure which takes α and an auction j and returns a participation set respecting α . Fix an auction j with reserve price r_j . Renumber the advertisers so that $v_{1j} > v_{2j} > \dots > v_{Nj}$. The procedure builds the set S_j by considering advertisers in decreasing order, rejecting an advertiser i if $v_{ij}/v_{i'j} < \alpha_i$, where i' is the most recent advertiser added to S_j . If there is no such i' , then v_{ij} is compared to the reserve price r_j . The detailed pseudocode is given below.

```

1: Input:  $v_{1j} \geq v_{2j} \geq \dots \geq v_{Nj}$ 
            $(\alpha_1, \alpha_2, \dots, \alpha_N)$ ;  $\alpha_i \geq 1$ .
           reserve price  $r_j$ 
2: Output:  $S_j$ : an  $\alpha$ -respecting participation set for auction  $j$ .
3: current.price =  $r_j$ ;  $S_j = \emptyset$ ;
4: for  $i = N \rightarrow 1$  do
5:   if  $\frac{v_i(j)}{\text{current.price}} \geq \alpha_i$  then
6:      $S_j = S_j \cup i$ .
7:     current.price =  $v_{ij}$ 
8:   end if
9: end for
    
```

Fig. 1. Procedure ReturnParticipationSet (j)

CLAIM 1. *The participation set returned by procedure ReturnParticipationSet is an α -respecting set.*

PROOF. Whenever an advertiser i is being considered, current.price is either the rank-score of an advertiser i' who is present in S_j and no one strictly between i and i' are present, or is the reserve price. Therefore, if i were to be put in S_j , his pROI would be $v_{ij}/\text{current.price}$. Since no advertiser between i and i' will be in S_j , and since an advertiser in S_j is never deleted, we get that if $i \in S_j$ at the end, $\text{pROI}(i, S_j) \geq \alpha_i$. Otherwise, $\text{pROI}(i, S_j \cup i) < \alpha_i$. That is, the final S_j is α -respecting. \square

Given α , we can find a participation profile which is α -respecting and therefore satisfies the optimal throttling condition. Also note that if $\alpha_i = 1$ for any advertiser, then that advertiser participates in all auctions. The resulting profile, however, need not be feasible. Therefore, if we can find α_i 's such that the resulting participation profile is feasible *and* $\alpha_i = 1$ for all advertisers i whose spends are smaller than their budgets, then we have obtained a regret-free participation profile. In the next section we show how to modify the above procedure to return randomized participation profiles, and prove using Brouwer's fixed point theorem that there exists α 's which lead to regret-free randomized participation profiles.

Existence of Regret-Free Participation Profiles. For each advertiser i , there are a finite number of obtainable pROI values. These are precisely $\{v_{ij}/r_j, v_{ij}/v_{i'j} : j \in [M], v_{i'j} < v_{ij}\}$. Let the cardinality of this set be L_i ; note that $L_i \leq NM$. Arrange these L_i rational numbers in increasing order and let them be $1 = f_0 < f_1 < f_2 < \dots < f_{L_i}$. (Actually, these numbers should be indexed with a superscript i which we omit for brevity). Let F be an upper bound on f_{L_i} , for all i .

Now given α_i in $[1, F]$, we identify q such that $f_q \leq \alpha_i < f_{q+1}$. Let $\alpha_{i\downarrow}$ denote f_q and $\alpha_{i\uparrow}$ denote f_{q+1} . (For convenience, let $f_{L_i+1} = F$ and $f_{L_i+2} = \infty$). The following procedure describes a randomized participation profile.

- 1: Input: $(\alpha_1, \alpha_2, \dots, \alpha_n); \alpha_i \geq 1$.
- 2: Output: Randomized Participation Profile (S_1, S_2, \dots, S_M) .
- 3: For each i , evaluate $\alpha_{i\downarrow}$ and $\alpha_{i\uparrow}$ as described above.
 Define $p_i := \frac{\alpha_{i\uparrow} - \alpha_i}{\alpha_{i\uparrow} - \alpha_{i\downarrow}}$.
 Set random variable α'_i to $\alpha_{i\downarrow}$ with probability p_i and $\alpha_{i\uparrow}$ with probability $(1 - p_i)$.
- 4: For each j run ReturnParticipationSet(j) with the random α'_i 's. This induces the required distribution.

Fig. 2. Procedure RandPartProf

CLAIM 2. *For any advertiser i , for any set $S_j \in \text{supp}(S_j)$, the following holds. If $i \in S_j$, then $\text{pROI}(i, S_j) \geq \alpha_{i\downarrow}$. If $i \notin S_j$, then $\text{pROI}(i, S_j \cup i) \leq \alpha_{i\downarrow}$.*

PROOF. Suppose $i \in S_j$. Since $\alpha'_i \geq \alpha_{i\downarrow}$, we get $\text{pROI}(i, S_j) \geq \alpha_{i\downarrow}$. From **Claim 1**, if $i \notin S_j$, then we get $\text{pROI}(i, S_j \cup i) < \alpha_{i\uparrow}$. This implies $\text{pROI}(i, S_j \cup i) \leq \alpha_{i\downarrow}$, since pROI takes discrete values. \square

This implies the following claim.

CLAIM 3. *The randomized participation profile (S_1, \dots, S_M) satisfies the optimal throttling condition of [Definition 2.1](#).*

CLAIM 4. *For all advertisers i , $\text{Exp}[\text{spend}(i)]$ is a continuous function of α .*

PROOF. $\text{Exp}[\text{spend}(i)] = \sum \text{Pr}[\alpha'] \cdot \text{spend}(i, \alpha')$, where α' be an instantiation in the random choices of the advertisers, and $\text{spend}(i, \alpha')$ is the total spend of advertiser i in the deterministic participation profile obtained in running ReturnParticipationSet(j) for every auction j . The summation is over all possible α' 's. Therefore, $\text{Exp}[\text{spend}(i)]$ is a piecewise linear continuous function. \square

We now define a mapping $\phi : [1, F]^N \mapsto [1, F]^N$ as follows. Given α , calculate $\text{Exp}[\text{spend}(i)]$ in the procedure `RandPartProf` with input α .

$$\phi(\alpha)_i := \Pi(\alpha_i + \eta(\text{Exp}[\text{spend}(i)] - B_i)) \quad (1)$$

where $0 < \eta < 1$ is any scalar, and Π is the projection operator defined as $\Pi(x) = x$ if $x \in [1, F]$, $\Pi(x) = F$ if $x > F$ and $\Pi(x) = 1$ if $x < 1$.

From [Claim 4](#), we get that ϕ is a continuous function defined over a compact space. Therefore, Brouwer's fixed point theorem guarantees the existence of $\alpha^* := (\alpha_1^*, \dots, \alpha_N^*)$ such that $\phi(\alpha^*) = \alpha^*$. The following lemma proves [Theorem 2.4](#).

LEMMA 3.2. *RandPartProf with input α^* returns a regret-free randomized participation profile.*

PROOF. [Claim 3](#) implies the optimal throttling condition. We now prove the feasibility and maximal participation condition. Suppose there's an advertiser i with $\text{Exp}[\text{spend}(i)] < B_i$. We claim that $\alpha_i^* = 1$; this is because $\alpha_i^{*\prime} = \Pi(\alpha_i^* - \delta)$ for some $\delta > 0$. These two are the same iff $\alpha_i^* = 1$. The fact that $\alpha_i^* = 1$ means that bidder i participates in all auctions. Therefore maximal participation is guaranteed.

Suppose there's an advertiser i with $\text{Exp}[\text{spend}(i)] > B_i$. Once again, by a similar argument above we'll have $\alpha_i^* = F$. However, at $\alpha_i^* = F$, $\text{Exp}[\text{spend}(i)] = 0$. Since budgets are positive, we get a contradiction. \square

■ ([Theorem 2.4](#))

Proof of Corollary 2.5: This follows from guessing the support of the regret-free profile. For each auction j , let S_j be our guess of the support. Given (S_1, \dots, S_M) , there is a finite procedure to check the optimal throttling condition (note that this depends only on the support, and not on the probabilities). Furthermore, the support also tells us which advertisers must satisfy $\text{Exp}[\text{spend}(i)] = B_i$. Once we have these, the existence of probabilities is a simple linear program. The existence theorem guarantees that for one guess (out of the 2^{MN^K} possible guesses) will lead to a feasible solution. The linear program has rational entries, and thus the solution is rational. ■ ([Corollary 2.5](#))

3.2. Polynomial Time Algorithm for Single Slot

In this section, we prove [Theorem 2.6](#), which we restate below. Recall that for this case, we focus on the weaker notion of regret-freeness where we only care about actual ROI and not potential ROI. In particular, recall that $\text{minROI}'(i)$ is the minimum ROI among all the auctions that i wins: $\text{minROI}'(i) := \min_{j: \text{spend}(i,j) > 0} \text{ROI}(i, S_j)$.

THEOREM 2.6. *In the case of single slots, there exists a polynomial time algorithm which given an auction instance finds a deterministic participation profile (S_1, \dots, S_M) such that*

- (1) Feasibility: For each advertiser i , $\text{spend}(i) \leq B_i + v_{max}^i$, where $v_{max}^i = \max_j v_{ij}$
- (2) Maximal Participation: If $\text{spend}(i) < B_i$, then $i \in S_j$ for all j .
- (3) Optimal Throttling: If $i \notin S_j$, then $\text{ROI}(i, S_j \cup i) \leq \text{minROI}'(i)$.

The proof differs from the one in the previous section. In fact, our algorithm is similar to the deferred acceptance algorithm for stable matching. The algorithm is deterministic; the flip side is that it works only for the single slot case.

Proof of Theorem 2.6: For each auction j , we associate the order σ_j on the advertisers in decreasing order of rank-scores. That is, $v_{\sigma_j(1)j} \geq v_{\sigma_j(2)j} \geq \dots \geq v_{\sigma_j(N)j}$. In fact,

it suffices to restrict attention to $n \leq N$ such that $v_{\sigma_j(n)j} > r_j$. For simplicity, we will henceforth assume σ_j is the identity permutation.

In what follows, we represent a participation set S_j as a pair (i, i') , $i < i'$, which represents the set $\{i, i', i' + 1, \dots, N\}$. For auction j , we associate the following ‘preference list’ on pairs.

$$\Pi_j := ((1, 2), (2, 3), (1, 3), (3, 4), (2, 4), (1, 4), \dots, (n-1, n), (n-2, n), \dots, (1, n))$$

Our algorithm proceeds in iterations with each iteration having a proposal phase and a disposal phase as in the deferred acceptance algorithm; in our case, auctions propose and advertisers dispose. Each advertiser i maintains a list of tentative auctions he participates in. The list Λ_i for advertiser i , has entries of the form $\{(j, (i, i'))\}$ where $v_{ij} > v_{i'j}$. That is, it participates and wins auction j , and the second-highest advertiser is advertiser i' . We will maintain that any auction j appears in at most one Λ_i . If auction j appears in exactly one $(j, (i, i')) \in \Lambda_i$ for some i, i' , we call it tentatively allocated to (i, i') . Otherwise, it is called unallocated. Initially all Λ_i are empty, and all auctions are unallocated. We will maintain the running counter $\text{spend}(i)$ to denote $\sum_{(j_k, (i, i_k)) \in \Lambda_i} v_{i_k j_k}$, the tentative spend of advertiser i .

In the proposal phase, every unallocated auction j ‘proposes’ to the first pair (i, i') in its preference list Π_j . This causes $(j, (i, i'))$ to be added to Λ_i , and j is tentatively allocated to (i, i') . Note that the invariant on Λ_i ’s is maintained. We delete (i, i') from Π_j . At this point, each auction is now tentatively allocated.

Subsequently, in the disposal phase, each advertiser i sorts the tuples $(j, (i, i'))$ in Λ_i in decreasing order of $v_{ij}/v_{i'j}$. He then deletes the tuples at the end, until deleting any more would underspend his budget. More precisely, suppose Λ_i is sorted as $((j_1, (i, i_1)), (j_2, (i, i_2)), \dots, (j_k, (i, i_k)))$, then we delete all tuples with index $> r$ if

$$\sum_{1 \leq \ell < r} v_{i_\ell j_\ell} < B_i \quad \text{and} \quad \sum_{1 \leq \ell \leq r} v_{i_\ell j_\ell} \geq B_i \quad (2)$$

All auctions whose associated tuples are deleted, that is $\{j_{r+1}, \dots, j_k\}$, are rendered unallocated, and we move back to the proposing round.

The algorithm terminates if either there are no unallocated auctions subsequent to the disposal phase, or Π_j is empty for every unallocated auction.

CLAIM 5. *The algorithm terminates in $O(MN^2)$ iterations.*

PROOF. After every proposal phase the size of $\sum_{j \in [M]} |\Pi_j|$ decreases by at least one. \square

At the end, the final participation profile (S_1, \dots, S_M) is determined as follows. If j is tentatively allocated to pair (i, i') , we have $S_j := \{i, i', i' + 1, \dots, N\}$. Else, $S_j = \emptyset$. The following claim implies the feasibility condition. In fact, the feasibility condition is maintained after every iteration.

CLAIM 6. *After the disposal phase of any iteration, $\text{spend}(i) \leq B_i + v_{max}^i$.*

PROOF. This follows from (2). \square

CLAIM 7. *If at the disposal phase of some iteration t , an advertiser i disposes some auction, then henceforth $\text{spend}(i) \geq B_i$.*

PROOF. Whenever an advertiser disposes at some iteration t , right after that step we have $\text{spend}(i) \geq B_i$ from (2). Furthermore, Λ_i remains unchanged till i has to dispose again. \square

CLAIM 8. *If $\text{spend}(i) < B_i$, then $i \in S_j$ for all j .*

PROOF. From Claim 7, we need to bother only about advertisers who never dispose. Consider such an advertiser i and consider an auction j which has been allocated (i_1, i_2) . Again, for simplicity, assume σ_j is the identity. We claim that $v_{ij} < v_{i_2j}$ which will imply $i \in S_j$. Suppose $v_{ij} > v_{i_1j}$. Then (i, i_1) lies before (i_1, i_2) in Π_j , and therefore since i must have disposed this. Similarly, if $v_{i_1j} > v_{ij} > v_{i_2j}$, then (i, i_2) lies before (i_1, i_2) in Π_j , and thus i must have disposed before. \square

CLAIM 9. $\min_{(j_k, (i, i_k)) \in \Lambda_i} \frac{v_{ij_k}}{v_{i_kj_k}}$ *is monotonically nondecreasing over iterations.*

PROOF. Whenever Λ_i is modified, the auctions with least ROI are deleted. \square

CLAIM 10. *If $i \notin S_j$, then $\text{ROI}(i, S_j \cup i) \leq \min \text{ROI}'(i)$.*

PROOF. Let $(j_k, (i, i_k))$ be the auction achieving $\min \text{ROI}'(i)$. From Claim 8, we know that $\text{spend}(i) \geq B_i$. Suppose auction j has been allocated $(j, (i_1, i_2))$ and $S_j \neq \emptyset$. If $v_{ij} < v_{i_1j}$, then $\text{ROI}(i, S_j \cup i) = 0$ since it doesn't win. (Note that $\text{pROI}(i, S_j \cup i)$ is non-zero.) If $v_{ij} > v_{i_1j}$, then (i, i_1) lies before (i_1, i_2) in Π_j . Therefore, i must have disposed (i, i_1) implying $v_{ij}/v_{i_1j} \leq \min \text{ROI}'(i)$ at that time. From the previous claim, we see that remains true at the end as well. If $S_j = \emptyset$, then either $(i, N) \notin \Pi_j$, that is, $v_{ij} < r_j$ in which case $\text{ROI}(i, S_j \cup i) = 0$; or, i disposed (i, N) implying $v_{ij}/r_j \leq \min \text{ROI}'(i)$ at that time, and which remains true till the end. \square

Claim 5, Claim 6, Claim 8, and Claim 10 prove the theorem. ■ (Theorem 2.6)

4. EXPERIMENTAL RESULTS

4.1. Heuristics

We first describe our benchmark budget smoothing algorithm, `RandomThrottling`, which is inspired by heuristics commonly used in practice. We then describe our heuristic, `WaterLevel`, inspired by the proof of Theorem 2.4.

`RandomThrottling` operates on the premise that if at any time an advertiser i has spent faster than his average rate, then he is throttled with a certain probability depending on the over-spend. For each advertiser i , we maintain a participation rate π_i that is continuously updated; it increases if the advertiser is underspending and decreases if he is overspending. We define the precise (multiplicative) update below after first introducing our new heuristic.

The new heuristic, which we call the `WaterLevel` algorithm, is inspired by the proof of Theorem 2.4. The heuristic starts off with an initial guess of $\alpha_i = 1$. Recall (from §3.1) that this implies every advertiser participates in every auction to begin with. As the day progresses, we modify the α_i 's of various advertisers depending on the rate of spend. After m of the M auctions, we expect each advertiser to have spent m/M fraction of his budget. The heuristic compares the fraction of actual budget spent after m auctions and bumps up α_i if the ratio is higher than m/M , and bumps down α_i if it is lower.

More precisely, the heuristic maintains a parameter η which is called the step-size, and a parameter γ which is called the *front loading parameter*. After each auction j ,

the spend $\text{spend}(i, j)$ is evaluated, and α_i is modified to $\alpha_i \cdot \exp\left(\eta \cdot \left(\frac{\text{spend}(i, j)}{B_i} - \frac{\gamma}{M}\right)\right)$. With the parameter γ set to 1, the advertiser is expected to exhaust his budget at precisely the last moment. A value of $\gamma > 1$ allows the advertiser to try to exhaust the budget slightly earlier. This update function is inspired by an algorithm in [Devanur et al. 2011]. The complete heuristic is described in Fig. 3.

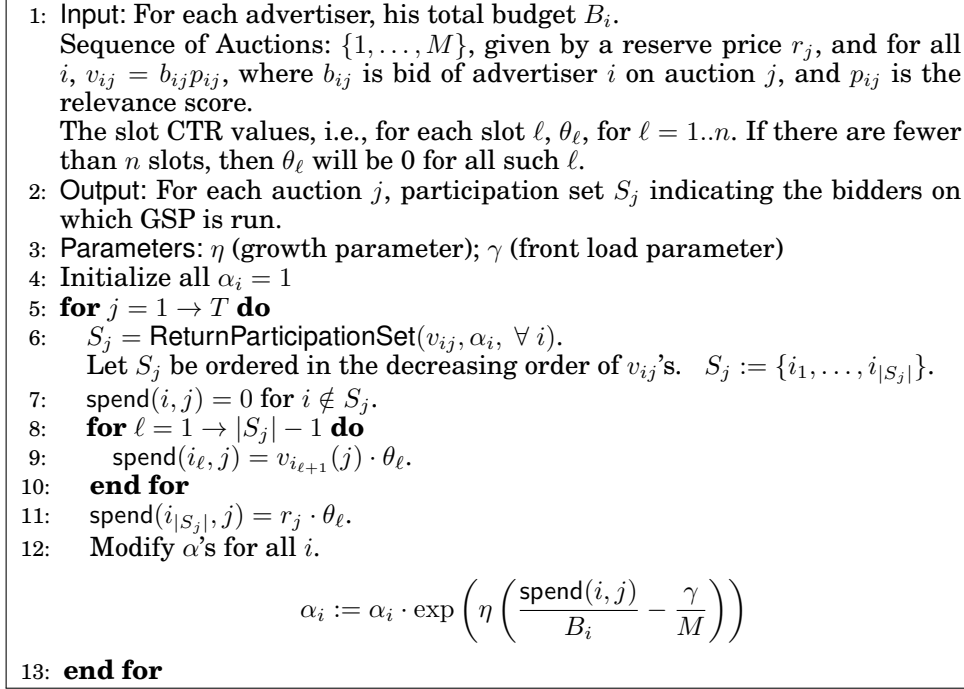


Fig. 3. The WaterLevel heuristic.

The update rule for RandomThrottling is similar, except for a change in sign, since the participation rate decreases with a higher spend.

$$\pi_i := \pi_i \cdot \exp\left(\eta \left(\frac{-\text{spend}(i, j)}{B_i} + \frac{\gamma}{M}\right)\right)$$

4.2. Results

The goal of our experiments is to compare the performances of the two algorithms on instances that are as close to the real auctions as possible. Due to the enormity of the number of advertisers and auctions in the Bing marketplace, an ideal solution is to find a small piece of the market that is isolated from the rest. Towards this end we used a clustering algorithm on the advertiser-auction graph and found a set of 200 clusters/micro-markets whose graph conductance is less than 10%.⁷ From these micro-markets we obtained two datasets consisting of ≈ 3.70 million and ≈ 10 million auctions corresponding to different lengths of time periods. We refer to these two datasets

⁷ The conductance of a subset of vertices of a graph is the ratio of the number of edges crossing the subset to the number of vertices in the subset (or the complement of the subset, whichever is smaller).

as *Medium* and *Large*. For these auctions we recorded all the information required to be able to re-run the auction. In particular, we recorded reserve prices, number of ad slots, bids and probability of click for each ad. Furthermore, we noted these details for all the *eligible* ads for the auctions prior to any filtering. Thus the input data for the experiments was independent⁸ of the throttling decisions made by the budget smoothing policy in effect. In our simulations we assumed that the click probabilities were exact; in fact, we treat click probabilities as fractions so that the effect due to randomization of clicks is eliminated. This is necessary for offline experiments, but also avoids skewing of results due to errors in click probabilities. We created advertiser budgets using the observed spends in each sample; we were not able to use real budget constraints because they are not given at a granularity that matches the time period, the traffic sample or the auctions in the micro-markets.

We also created different instances by scaling the budgets differently, to study the effect of budgets becoming more constraining. We created two budget levels, which we call *Generous* and *Constraining*. The Constraining budgets are half as much as the Generous budgets. Finally, for each instance we ran both algorithms with γ set to 1 and 1.2. The runs with $\gamma = 1.2$ are referred to as *Frontloaded* and the ones with $\gamma = 1$ as *Non-frontloaded*.

In addition to measuring the value generated and the revenue obtained by the algorithms, we also measure a notion of regret which we define next.

Definition 4.1. Let A be a budget smoothing algorithm which chooses slate S_j for auction j . We define the regret of a given advertiser i w.r.t the outcome of this algorithm. For auction j , let q_j be the price this advertiser faces assuming the rest of the advertisers participate according to the throttling decisions of the algorithm. Let

$$OPT(i) := \max\left\{\sum_j \frac{v_{ij}x_j}{q_j} : 0 \leq x_j \leq 1, \sum_j x_j q_j \leq B_i\right\} \text{ and}$$

$$Real(i) := \sum_{j: spend(i,j) > 0} \frac{v_{ij}}{q_j}.$$

The *Regret* for advertiser i is defined as the quantity $Regret(i) := OPT(i) - Real(i)$, and the regret of the algorithm A is defined as the quantity

$$Regret(A) := \sum_i Regret(i).$$

Remark 4.2. If we assume that algorithm A respects budget constraints of every advertiser, then $Real(i) \leq OPT(i)$ for each advertiser and equality holds if and only if the algorithm allows the advertiser to participate in the top ROI auctions from the advertiser's standpoint. If we allow fractional allocation of auctions the regret metric for a regret-free equilibrium allocation is zero.

We report our findings from these experiments next.

4.2.1. Aggregates. We first compare the social welfare, total revenue and total regret of `WaterLevel` with that of `RandomThrottling`. The following table shows the percentage increase in the social welfare and total revenue of `WaterLevel` over that of

⁸The bids we observe may have been an adaptation to the current smoothing policy. However, we believe the experiments still show the general features of the algorithm since the baseline heuristic was also run on the same data.

RandomThrottling. For regret, it shows the total regret of WaterLevel as a percentage of the regret of RandomThrottling. We also report the following quantities:

- Percentage of advertisers for whom the value went down.
- Percentage of advertisers for whom the value rose by at least 5%.
- Percentage of advertisers for whom the spend changed by less than 5%.
- Percentage of advertisers in WaterLevel for whom the budget was exhausted at a time 10% earlier than expected (the expected time at which budget exhausts depends on the γ parameter.) This is the only statistic that depends on WaterLevel alone.

We report this for each of the 8 runs, one for each data set medium and large, for each of the two budget levels generous and constrained, and for each γ parameter of 1 and 1.2.

| Conf. | VI | RI | Reg | VD | 5% VI | SC | ET |
|----------|-----|-----|-----|-----|-------|-----|-----|
| M, G, F | 12% | 1% | 56% | 18% | 66% | 93% | 3% |
| M, C, F | 22% | 2% | 39% | 12% | 80% | 95% | 7% |
| M, G, NF | 17% | 0% | 6% | 16% | 68% | 86% | 3% |
| M, C, NF | 29% | 2% | 4% | 11% | 81% | 91% | 8% |
| L, G, F | 13% | 1% | 62% | 15% | 68% | 96% | 7% |
| L, C, F | 24% | 2% | 41% | 8% | 87% | 98% | 10% |
| L, G, NF | 18% | -1% | 6% | 12% | 69% | 90% | 5% |
| L, C, NF | 31% | 2% | 4% | 5% | 89% | 96% | 9% |

Legend: M: Medium, L: Large, G: Generous, C: Constrained, F: Frontloaded, NF: Non-Frontloaded. VI: Value increase, RI: Revenue Increase, Reg: Ratio of Regret of WaterLevel to that of RandomThrottling, VD: advertisers with a decrease in value, 5% VI: advertisers with at least 5% increase in value, SC: Advertisers with less than 5% change in spend, ET: Advertisers with exhaust time 10% sooner than expected.

Observations.

- WaterLevel consistently gives a higher value than Random-throttling. The relative increase in value is higher in case the budgets are more constraining. In fact most advertisers don't see their value decreasing and a majority of the advertisers see at least a 5% boost to their value.
- The revenue of WaterLevel and RandomThrottling are essentially the same. In fact most advertisers do not see a big change in the spend.
- WaterLevel does a very good job of actually minimizing regret as intended, especially without frontloading.

4.2.2. *Change in value vs Change in revenue.* For each advertiser, we consider the percentage change in his value and the percentage change in his spend when going from RandomThrottling to WaterLevel. The results are shown for the (M, G, NF) run in the left chart of Figure 4; all other runs were similar.

Observations. The scatter plot forms an "L" shape with the origin at the intersection and we see that majority of advertisers get a higher value at the same spend. Quite a few of them spend a lot more for a small increase in value. We conjecture that the former are budget constrained and their budgets are better utilized; the latter are not

budget constrained and they spend more since other advertisers' increased efficiency results in higher prices to these advertisers.

4.2.3. *Regret.* We take a closer look at the distribution of regrets between `WaterLevel` and `RandomThrottling`. In the right chart of Figure 4 we plot the distribution of the regret for both algorithm for the (M, G, NF) run. For both algorithms, we have sorted the individual advertiser regret in descending order and have plotted a line graph of the regret values in this order. The graph has a logarithmic y-axis.




Fig. 4. Experimental findings for Medium Non-Frontloaded dataset

Observation. We observe that `WaterLevel` minimizes regret more effectively than does `RandomThrottling`.

5. CONCLUDING REMARKS

In this paper we introduced a novel solution concept for the budget smoothing problem. Our work raises many open problems, including the following prominent ones.

- Is there a polynomial time regret-free budget smoothing algorithm? Even if a single auction is repeated multiple times?
- What is the complexity of finding the social welfare maximizing regret-free participation profile?
- Is there an online algorithm that approximates the regret-free participation profile? Even in the case of a single slot?
- We have considered a model where the advertisers' bids remain the same and studied the effects of changing the budget smoothing policy. It is natural to expect that the advertisers will indeed react⁹ to a new policy and change their bids accordingly. Studying this stage theoretically is an interesting and important open problem.

The experimental results indicate that `WaterLevel` is a good candidate for budget smoothing. It benefits a majority of the advertisers while making a few worse-off. The overall value improvement is quite significant while the overall revenue impact is neutral. We are optimistic about this heuristic becoming practically useful.

⁹ Our experiments show that the effects of the new budget smoothing policy with the same bids is mostly an increase in advertiser value, which indicates that the reaction to the new policy will be mostly positive. In particular, we expect this increase in value to induce the advertisers to bring in more advertising dollars.

REFERENCES

- ABRAMS, Z., MENDELEVITCH, O., AND TOMLIN, J. 2007. Optimal delivery of sponsored search advertisements subject to budget constraints. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 272–278.
- AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *ACM Conference on Electronic Commerce*.
- AMIN, K., KEARNS, M., KEY, P., AND SCHWAIGHOFER, A. 2012. Budget optimization for sponsored search: Censored learning in mdps. In *UAI*.
- CHAKRABARTY, D., ZHOU, Y., AND LUKOSE, R. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *WINE*.
- DEVANUR, N., SIVAN, B., AND AZAR, Y. 2012. Asymptotically optimal algorithm for stochastic adwords. In *ACM Conference on Electronic Commerce*.
- DEVANUR, N. R. AND HAYES, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *ACM Conference on Electronic Commerce*, J. Chuang, L. Fortnow, and P. Pu, Eds. ACM, 71–78.
- DEVANUR, N. R., JAIN, K., SIVAN, B., AND WILKENS, C. A. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*. EC '11. ACM, New York, NY, USA, 29–38.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97, 1, 242–259.
- FELDMAN, J., HENZINGER, M., KORULA, N., MIRROKNI, V. S., AND STEIN, C. 2010. Online stochastic packing applied to display ad allocation. In *ESA*. 182–194.
- FELDMAN, J., MEHTA, A., MIRROKNI, V., AND MUTHUKRISHNAN, S. 2009. Online stochastic matching: Beating $1-1/e$. In *FOCS '09: Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, Washington, DC, USA, 117–126.
- FELDMAN, J., MUTHUKRISHNAN, M., PÁL, M., AND STEIN, C. 2007. Budget optimization in search-based advertising auctions. In *ACM Conference on Electronic Commerce*.
- GALE, D. AND SHAPLEY, L. 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–14.
- GOEL, A., MAHDIAN, M., NAZERZADEH, H., AND SABERI, A. 2010. Advertisement allocation for generalized second-pricing schemes. *Oper. Res. Lett.* 38, 6, 571–576.
- HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *WINE*. 170–181.
- LAME, R. P. AND TARDOS, E. 2010. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *FOCS*.
- MAHDIAN, M., NAZERZADEH, H., AND SABERI, A. 2012. Online optimization with uncertain information. *ACM Transactions on Algorithms* 8, 1.
- MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. 2005. Adwords and generalized on-line matching. In *FOCS 05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 264–273.
- MIRROKNI, V., GHARAN, S. O., AND ZADIMOGHADDAM, M. 2012. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *SODA*.
- VARIAN, H. 2006. Position auctions. *International Journal of Industrial Organization* 26, 6, 1163–1178.