

Preferred Deals in General Environments

Yuan Deng¹, Sébastien Lahaie² and Vahab Mirrokni²

¹Duke University

²Google Research

ericdy@cs.duke.edu, {slahaie, mirrokni}@google.com

Abstract

A preferred deal is a special contract for selling impressions of display ad inventory. By accepting a deal, a buyer agrees to buy a minimum amount of impressions at a fixed price per impression, and is granted priority access to the impressions before they are sent to an open auction on an ad exchange. We consider the problem of designing preferred deals (inventory, price, quantity) in the presence of general convex constraints, including budget constraints, and propose an approximation algorithm to maximize the revenue obtained from the deals. We then evaluate our algorithm using auction data from a major advertising exchange and our empirical results show that the algorithm achieves around 95% of the optimal revenue.

1 Introduction

Display advertising is the practice of placing text, banner, or video ads on publisher websites for the purpose of raising brand awareness via ad views, known as *impressions*. The display advertising industry is a major source of revenue for online publishers and Internet companies. In 2017, the revenue of this market exceeded \$27.5 billion for the U.S. alone [PwC, 2018]. Traditionally, display advertising has been mainly sold via reservation contracts or real-time bidding. In a reservation contract, the publisher buys a fixed amount of impressions on a segment of a publisher’s inventory, over some amount of time (e.g., one month), and the publisher commits to sending the advertiser the impressions [Feige *et al.*, 2008]. In real-time bidding, advertisers bid through an ad exchange to appear on the publisher’s website as the impressions arrive. Recently, an intermediate form of contract known as a *preferred deal* has grown in popularity [Digiday, 2018]. Under a preferred deal, the advertiser and the publisher agree on a fixed impression price and quantity (as in a reservation), but the advertiser is now given priority access to choose among impressions in real time. Impressions that the advertiser declines to buy can be sent on to other preferred deals with lower priority, or to an ad exchange.

Mirrokní and Nazerzadeh [2017] recently initiated a formal study of preferred deals. They design an approximation algorithm, called the auction-adjusted greedy algorithm (*AAG*),

for the problem of designing preferred deals for advertisers based on their historical bids on an ad exchange, in order to optimize for revenue. The idea is that the advertisers’ bids on the exchange (normally in second-price auctions) provide information about their distributions of value for different impressions, and this information can be used to construct and optimize deals going forward.

In practice, advertisers typically have strict budget limits on the amount they are able to spend on various parts of their marketing campaigns, and possibly other constraints such as volume constraints on the number of impressions from different subcategories of inventory (e.g., to ensure a diversity of impressions). In this paper, we design an approximation algorithm for optimizing revenue through preferred deals in the presence of budget and general convex constraints. To do so, we cast the *AAG* algorithm into a more general algorithmic framework which takes in an oracle that allocates impressions to optimize *welfare*, and results in a constant-factor approximation algorithm for the problem of *optimal-revenue* deal design. By using a *constraint-aware social welfare* oracle, we obtain constant-factor approximations in settings where buyer valuations are independent or additively correlated, where the constraint-aware social welfare is the social welfare respecting all constraints. For arbitrary correlation structures, we prove a negative result showing that no constant-factor approximation is possible.

To efficiently implement the social welfare oracle in practice, we use a reduction by Cai *et al.* [2012] from *ex-post* to *interim* impression allocation rules, which leads to a large-scale convex program. To our knowledge, this is the first practical implementation of this reduction. We evaluate our algorithm using auction data from the Google Ad Exchange, and compare it against the original *AAG* algorithm as well as second-price auctions (with and without reserve prices) as benchmarks. We find that the original *AAG* algorithm performs well at high budget levels, while second-price auctions perform well at low budget levels, but only our new algorithm achieves high revenue across the board, typically reaching at least 95% of the optimum.

Related Work. We briefly discuss research closely related to ours in this section. For a review of the literature on display advertising, readers are encouraged to refer to Korrula *et al.* [2016] and Choi *et al.* [2017] for recent surveys. As mentioned, Mirrokni and Nazerzadeh [2017] first for-

mally framed the preferred deals problem and considered the setting without constraints. The main reason behind the high revenue potential of preferred deals is bundling. Ghosh *et al.* [2007] show that it is computationally hard to find the optimal bundling and propose approximation algorithms. Babaioff *et al.* [2014] demonstrate that the strategy of either selling each item separately or selling them all as a grand bundle can obtain a constant approximation of the optimal revenue for selling multiple items to a single buyer. Our results provide a constant-factor approximation for the multi-buyer case with general convex constraints under the assumption that the items are identically and independently drawn, which may be of independent interest.

Another related line of work is auction design with budget constraints, which was initiated by Che and Gale [1998] (see also [Benoît and Krishna, 2001; Che and Gale, 2000]). Pai and Vohra [2014] characterize revenue-maximizing auctions for a single unit. For multi-unit settings, Dobzinski *et al.* [2012] show that there is no Pareto optimal and incentive compatible mechanism if the budgets are private, and for public budgets, they give a unique mechanism that is Pareto optimal and incentive compatible, which was later extended to divisible units by Bhattacharya *et al.* [2010]. Our work is also related to the recent study of liquid welfare (see [Dobzinski and Paes Leme, 2014; Lu and Xiao, 2017; Azar *et al.*, 2017]).

2 Preliminaries

In the preferred deals problem there is a publisher with a sequence of impressions, denoted as \mathcal{I} , to sell to n buyers. Let $\mathbf{v}(k) = (v_1(k), v_2(k), \dots, v_n(k))$ be a valuation profile of the buyers for the k -th impression. In this paper, we consider an environment where buyer i has a public budget constraint B_i . As usual, we use $-i$ to denote buyers other than i . We assume the buyer is risk-neutral with quasi-linear utility: the utility from purchasing such an impression with valuation v at price p is $v - p$. However, if the buyer's total spend exceeds B_i , his utility becomes $-\infty$ (in other words, the budget is a hard constraint on spend).

Let $x_i : \mathcal{I} \rightarrow [0, 1]^n$ with $\sum_i x_i(k) \leq 1$ be an allocation rule such that $x_i(k)$ specifies the probability that buyer i gets the k -th impression. For convenience, let the norm of the allocation rule be $|x_i| = \sum_{k \in \mathcal{I}} x_i(k)$. Denote buyer i 's expected average valuation subject to x_i by $\mathbb{E}_{\mathcal{I}}[v_i | x_i] = \sum_k v_i(k) \cdot \frac{x_i(k)}{|x_i|}$. In addition to the budget constraints, we consider an environment in which each buyer i has a personalized constraint $\mathcal{P}_i \subseteq [0, 1]^{|\mathcal{I}|}$.

Definition 1 (Personalized Constraint). *An allocation rule x is feasible only if for each buyer i , $x_i \in \mathcal{P}_i$. We assume that \mathcal{P}_i is convex and $\vec{0} \in \mathcal{P}_i$ where $\vec{0}$ is a rule that does not allocate anything to a buyer, i.e., $\vec{0}(k) = 0$ for all $k \in \mathcal{I}$.*

Our personalized constraints are general enough to capture many practical requirements, such as overall volume constraints or weighted constraints on different subcategories of inventory (e.g., different user demographics). Let I_i be a finite set of *different* impressions from the perspective of buyer i . In other words, the set \mathcal{I} of all impressions can be parti-

tioned into $|I_i|$ subsets of impressions such that within each subset, the impressions are *identical* according to buyer i .

Definition 2 (Identical Impressions). *From the perspective of buyer i , impressions k and k' are identical if their values are the same, i.e., $v_i(k) = v_i(k')$, and they are anonymous in \mathcal{P}_i , i.e., for any $x_i \in \mathcal{P}_i$ and all $0 \leq \delta \leq x_i(k) + x_i(k')$, $x_i^\delta \in \mathcal{P}_i$ where $x_i^\delta(k) = \delta$, $x_i^\delta(k') = x_i(k) + x_i(k') - \delta$, and $x_i^\delta(j) = x_i(j)$ for all $j \notin \{k, k'\}$.*

Anonymity means that the two impressions can be treated as two copies of the same impression when checking whether the allocation rule satisfies the personalized constraint. We use $f_i : I_i \rightarrow [0, 1]$ to denote the (empirical) marginal distribution of buyer i 's impressions: for $\kappa_i \in I_i$, $f_i(\kappa_i) = \sum_k \mathbf{1}\{k = \kappa_i\} / |\mathcal{I}|$.

We consider three types of correlation: no correlation, additive correlation, general correlation. In an environment with no correlation, we assume buyer i 's distribution f_i is independent of other buyers' distributions f_{-i} . For additive correlation, we allow the valuations to be correlated by a common value. More precisely, $v_i(k) = \eta_0 + \nu_i(k)$ where ν_i is independent of $\nu_j(k)$ for any $j \neq i$ and $\eta_0 \geq 0$ is a random common component. An environment with general correlation means that there can be arbitrary correlation between valuations.

Definition 3 (Preferred Deal). *A preferred deal $\langle p, \mu, \mathcal{I} \rangle$ for a sequence of impressions \mathcal{I} is specified by a price per impression p and a minimum purchase requirement μ . By accepting the deal, the buyer obtains priority access to the impressions in \mathcal{I} and agrees to buy at least a μ fraction of the impressions, each at price p .*

By *priority access*, we mean that the buyer is given the opportunity to buy the impression before it is presented to any other buyer, in contrast to other selling mechanisms like auctions where buyers compete simultaneously. Given a deal $\langle p, \mu, \mathcal{I} \rangle$ for buyer i , since the price is fixed, the buyer will always cherry-pick the impressions (i.e., buy impressions with highest valuations) to maximize her utility, subject to her personalized constraint and the minimum purchase requirement.

Definition 4 (Cherry-picking). *Let CP_i be a cherry-picking function that returns the welfare-optimal allocation rule for buyer i , formally,*

$$\text{CP}_i(\mathcal{I}) = \arg \max_{x_i \in \mathcal{P}_i} |x_i| \cdot \mathbb{E}_{\mathcal{I}}[v_i | x_i].$$

Moreover, let $\text{CP}_i(\mathcal{I}, \mu)$ be a cherry-picking function satisfying the purchase requirement μ exactly, formally,

$$\text{CP}_i(\mathcal{I}, \mu) = \arg \max_{\substack{x_i \in \mathcal{P}_i \\ |x_i| = \mu |\mathcal{I}|}} \mathbb{E}_{\mathcal{I}}[v_i | x_i].$$

If $\mathcal{P}_i \cap \{x_i \mid |x_i| = \mu |\mathcal{I}|\} = \emptyset$, then $\text{CP}_i(\mathcal{I}, \mu) = \vec{0}$.

Let $\text{CPW}_i(\mathcal{I})$ be the welfare obtained by cherry-picking, i.e. $\text{CPW}_i(\mathcal{I}) = |\text{CP}_i(\mathcal{I})| \cdot \mathbb{E}_{\mathcal{I}}[v_i | \text{CP}_i(\mathcal{I})]$, and similarly, $\text{CPW}_i(\mathcal{I}, \mu) = \mu |\mathcal{I}| \cdot \mathbb{E}_{\mathcal{I}}[v_i | \text{CP}_i(\mathcal{I}, \mu)]$. Given a deal $\langle p, \mu, \mathcal{I} \rangle$, the buyer will choose $\mu^* |\mathcal{I}|$ impressions such that

$$\mu^* = \arg \max_{\mu \leq \mu' \leq B_i / (p |\mathcal{I}|)} \mu' \cdot (\mathbb{E}_{\mathcal{I}}[v_i | \text{CP}_i(\mathcal{I}, \mu')] - p)$$

and purchase the impressions according to $\text{CP}_i(\mathcal{I}, \mu^*)$.

The objective of the publisher is to design a series of preferred deals to maximize revenue. The publisher first selects a priority list $(\pi_1, \pi_2, \dots, \pi_n)$ of ordered buyer indices, and a pair (p_i, μ_i) of price and minimum purchase requirement for each buyer i . The publisher will offer deals to the buyers in sequence according to the priority list and a buyer will receive impressions that are left over from the buyers with higher priorities. More precisely, let $H_i(p, \mu, \mathcal{I})$ be the set of impressions purchased by buyer i when offered deal $\langle p, \mu, \mathcal{I} \rangle$. Then, the publisher will offer deals $(\langle p_{\pi_1}, \mu_{\pi_1}, \mathcal{I}_{\pi_1} \rangle, \langle p_{\pi_2}, \mu_{\pi_2}, \mathcal{I}_{\pi_2} \rangle, \dots, \langle p_{\pi_n}, \mu_{\pi_n}, \mathcal{I}_{\pi_n} \rangle)$ where $\mathcal{I}_{\pi_1} = \mathcal{I}$ and $\mathcal{I}_{\pi_{j+1}} = \mathcal{I}_{\pi_j} \setminus H_{\pi_j}(p_{\pi_j}, \mu_{\pi_j}, \mathcal{I}_{\pi_j})$.

For simplicity, we will omit the dependence on \mathcal{I} or \mathcal{I}_{π_j} when clear from context.

3 Auction-Adjusted Greedy Framework

We first introduce an algorithmic framework for computing preferred deals, which extends the algorithm proposed by Mirrokni and Nazerzadeh [2017] that works exclusively for an environment without any constraints. The framework reduces the problem of optimizing preferred deals to the problem of optimizing impression allocations.

Let S be the set of buyers and $\omega(\mathcal{I}, S)$ denote the maximum expected revenue that can be collected by offering a series of deals. Note that when $S = \{i\}$ is a singleton, the optimal approach is to offer a deal with $\mu_i = |\text{CP}_i|$ and $p_i = \min\{\mathbb{E}[v_i | \text{CP}_i], B_i / |\text{CP}_i|\}$.

Lemma 1. *For any μ such that $\text{CP}_i(\mu) \neq \vec{0}$, given a deal $\langle p, \mu \rangle$ with $p = \min\{\mathbb{E}[v_i | \text{CP}_i(\mu)], B_i / (\mu |\mathcal{I}|)\}$, purchasing impressions according to $\text{CP}_i(\mu)$ is buyer i 's best response.*

By Lemma 1, such a deal can extract the maximum social welfare subject to \mathcal{P}_i or buyer i 's budget. Therefore, $\omega(\mathcal{I}, \{i\}) = \min\{\text{CPW}_i, B_i\}$. For $|S| \geq 2$, we can define $\omega(\mathcal{I}, S)$ recursively as follows:

$$\omega(\mathcal{I}, S) = \max_{i \in S, \mu \in [0, 1]} \left\{ \min(\text{CPW}_i(\mu), B_i) + \omega(\mathcal{I} \setminus \text{CP}_i(\mu), S \setminus \{i\}) \right\}.$$

In this formula, we abuse notation and let $\text{CP}_i(\mu)$ be the set of impressions that buyer i would like to purchase, and therefore, $\mathcal{I} \setminus \text{CP}_i(\mu)$ is the set of remaining impressions after buyer i takes $\text{CP}_i(\mu)$. By Lemma 1, to extract revenue $\min(\text{CPW}_i(\mu), B_i)$ from buyer i , the publisher can offer a deal with $p_i = \min\{\mathbb{E}[v_i | \text{CP}_i(\mu)], B_i / (\mu |\mathcal{I}|)\}$ and $\mu_i = \mu$.

However, computing $\omega(\mathcal{I}, S)$ exactly takes exponential time. In particular, Mirrokni and Nazerzadeh [2017] show that it is NP-Hard to compute the optimal deals even without any constraints. To circumvent the hardness result, Mirrokni and Nazerzadeh [2017] propose an approximation algorithm, called auction-adjusted greedy algorithm (\mathcal{AAG}), to generate a series of deals when there are no constraints. We cast their algorithm in a more general framework that can handle general convex constraints by taking in an impression-allocation oracle. To define an impression-allocation oracle, we first define the notion of an impression-allocation rule.

Input: Impression \mathcal{I} , a set of buyers S ,
impression-allocation oracle \mathcal{O}

Output: Sequence of deals \mathcal{D} .

If $(|S| = 0)$ Return

Compute $\mathbf{a} = \mathcal{O}(S, \mathcal{I})$

Let $S_0 = \{j \in S | a_j = 0\}$ and $S_1 = S \setminus S_0$.

$\mathcal{D}_0 \leftarrow \{(p_j = 0, \mu_j = 0)\}_{j \in S_0}$

For each buyer $j \in S_1$

$\mu_j = a_j$

$p_j = \min\{\mathbb{E}[v_j | \text{CP}_j(\mu_j)], B_j / (\mu_j |\mathcal{I}|\})\}$

Choose a buyer $i \in \arg \max_{j \in S_1} p_j$

$\mathcal{D} \leftarrow (p_i, \mu_i) + \mathcal{AAG}(\mathcal{I} \setminus \text{CP}_i(\mu_i), S_1 \setminus \{i\}) + \mathcal{D}_0$

Return \mathcal{D}

Algorithm 1: \mathcal{AAG} Framework

Definition 5 (Impression-allocation Rule). *An impression allocation rule is $(a_1, \dots, a_n) \in [0, 1]^n$ with $\sum a_i \leq 1$ such that a_i specifies the fraction of impressions allocated to i .*

Definition 6 (Impression-allocation Oracle). *Given a set of buyers S and a sequence of impressions \mathcal{I} , an impression-allocation oracle \mathcal{O} returns an impression-allocation rule.*

Given an impression-allocation oracle \mathcal{O} , the \mathcal{AAG} framework (Algorithm 1) takes a set of buyers S and a sequence of impressions \mathcal{I} as input and outputs a series of deals. Intuitively, in each recursion, the \mathcal{AAG} algorithm obtains the minimum purchase requirement μ_j from the impression-allocation oracle \mathcal{O} and computes the associated price respecting μ_j as well as the constraints, such that buyer j will buy exactly all the impressions from $\text{CP}_j(\mu_j)$ according to Lemma 1. Finally, the algorithm selects a deal with maximum price to offer.

3.1 Constraint-aware Welfare and Oracle

To apply the \mathcal{AAG} framework in an environment with constraints, the key is to choose an impression-allocation oracle that respects all these constraints.

In an environment without any constraints, the maximum social welfare is a natural upper bound on the revenue of any deals. However, in the presence of budget constraints and other constraints, the maximum social welfare could be much larger than the revenue of the optimal deal: consider a simple case where there is only one buyer and a single impression, and the buyer with budget 1 has valuation m on the impression. In this example, the maximum social welfare is m but the maximum revenue, bounded by the budget, is 1.

Therefore, instead of social welfare, we use the *constraint-aware social welfare* as the benchmark, which generalizes the concept of liquid welfare defined exclusively for budget constraints [Dobzinski and Paes Leme, 2014]. Intuitively, constraint-aware social welfare is social welfare subject to the budget constraints and the personalized constraints.

Definition 7 (Constraint-aware Social Welfare). *Given a set of buyers S and a set of impressions \mathcal{I} , the constraint-aware social welfare $\text{CAW}(S, \mathcal{I})$ is computed according to the following convex program:*

$$\begin{aligned} \max \quad & \sum_i \sum_k v_i(k) \cdot x_i(k) \\ \text{s.t.} \quad & x_i \in \mathcal{P}_i \quad \forall i \\ & \sum_k v_i(k) \cdot x_i(k) \leq B_i \quad \forall i \\ & \sum_i x_i(k) \leq 1, x_i(k) \geq 0 \quad \forall i, k \end{aligned} \quad (\text{expost-CAW})$$

In the program, the first set of constraints ensure that the allocation rule satisfies all personalized constraints, while the second set of constraints ensure that the welfare contribution from each buyer does not exceed her budget constraint. The remaining constraints ensure that the allocation rule is valid. Let $x^{\text{CAW}(S, \mathcal{I})}$ be the allocation rule that maximizes $\text{CAW}(S, \mathcal{I})$.

Definition 8 (Constraint-aware Impression-allocation Oracle \mathcal{CAO}). *Given a set of buyers S and a sequence of impressions \mathcal{I} , the constraint-aware impression-allocation oracle \mathcal{CAO} returns an impression-allocation rule $\mathbf{a} = (a_1, \dots, a_n)$, where $a_i = \sum_k x_i^{\text{CAW}(S, \mathcal{I})}(k)/|\mathcal{I}|$.*

4 Approximation Guarantee

In this section, we present the formal proof for the approximation guarantee by applying the \mathcal{AAG} framework with constraint-aware social welfare oracle for additive correlations. We also provide a lower bound for general correlations.

4.1 Additive Correlation

Recall that under the additive correlation model, values take the form $v_i = \eta_0 + \nu_i$ where the ν_i are independent and η_0 is a random common component. We first consider the special case of no correlation ($\eta_0 \equiv 0$) and show that the \mathcal{AAG} algorithm with \mathcal{CAO} gives a $1/2$ -approximation.

Theorem 1. *The expected revenue of the sequence of deals found by the \mathcal{AAG} algorithm with \mathcal{CAO} is at least $\frac{1}{2}\text{CAW}(S, \mathcal{I})$ when there is no correlation.*

Proof. Let $\text{SOL}(S, \mathcal{I})$ be the revenue generated by \mathcal{AAG} algorithm with \mathcal{CAO} with input S and \mathcal{I} . We prove by an induction on the size of S from $|S| = 1$ to $|S| = n$. When $S = \{i\}$, if $p_i = \mathbb{E}[v_i | \text{CP}_i(\mu_i)]$, then $\text{SOL}(S, \mathcal{I})$ is exactly $\text{CAW}(S, \mathcal{I})$; otherwise, $\text{SOL}(S, \mathcal{I})$ is B_i , but $\text{CAW}(S, \mathcal{I})$ must be equal to B_i in this case. Therefore, the base case is true. Assume the induction hypothesis is true for all $|S_1| \leq s$. Then for $|S_1| = s + 1$, for convenience, let

$$\alpha(S, \mathcal{I} | S', \mathcal{I}') = \sum_{k \in \mathcal{I}'} \sum_{i \in S'} v_i(k) \cdot x_i^{\text{CAW}(S, \mathcal{I})}(k)$$

Therefore, $\text{CAW}(S, \mathcal{I}) = \alpha(S, \mathcal{I} | S, \mathcal{I})$. Let i be the buyer that the algorithm offers a deal to with input (S, \mathcal{I}) . First:

$$\begin{aligned} \text{CAW}(S, \mathcal{I}) &= \alpha(S, \mathcal{I} | S, \mathcal{I}) \\ &= \alpha(S, \mathcal{I} | \{i\}, \mathcal{I}) + \alpha(S, \mathcal{I} | S \setminus \{i\}, \text{CP}_i(\mu_i)) \\ &\quad + \alpha(S, \mathcal{I} | S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i)). \end{aligned}$$

Notice that

$$\begin{aligned} \alpha(S, \mathcal{I} | S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i)) &\leq \text{CAW}(S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i)) \\ &\leq 2 \cdot \text{SOL}(S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i)) \end{aligned}$$

where the first inequality is due to the fact that $x_j^{\text{CAW}(S, \mathcal{I})}(k)$ constrained in $j \in S \setminus \{i\}$ and $k \in \mathcal{I} \setminus \text{CP}_i(\mu_i)$ is a valid solution for program (expost-CAW) with $S \setminus \{i\}$ and $\mathcal{I} \setminus \text{CP}_i(\mu_i)$ as input, and the second inequality is by the induction hypothesis.

Next, for any buyer j , notice that if $p_j = B_j / (\mu_j |\mathcal{I}|)$, then we have $p_j \cdot \mu_j \cdot |\mathcal{I}| = B_j \geq \alpha(S, \mathcal{I} | \{j\}, \mathcal{I})$. Otherwise, notice that $\mu_j = a_j = \sum_k x_i^{\text{CAW}(S, \mathcal{I})}(k) / |\mathcal{I}|$ and $p_j = \mathbb{E}_{\mathcal{I}}[v_j | \text{CP}_j(\mu_j)]$ is the average valuation of a_j fraction of impressions due to cherry-picking. Therefore, we have $\alpha(S, \mathcal{I} | \{j\}, \mathcal{I}) \leq p_j \cdot \mu_j \cdot |\mathcal{I}|$ for all j . As for $\alpha(S, \mathcal{I} | S \setminus \{i\}, \text{CP}_i(\mu_i))$, since there is no correlation,

$$\begin{aligned} \alpha(S, \mathcal{I} | S \setminus \{i\}, \text{CP}_i(\mu_i)) &= \mu_i \cdot \alpha(S, \mathcal{I} | S \setminus \{i\}, \mathcal{I}) \\ &\leq \mu_i \cdot \alpha(S \setminus \{i\}, \mathcal{I} | S \setminus \{i\}, \mathcal{I}) \\ &\leq \mu_i \cdot \alpha(S, \mathcal{I} | S, \mathcal{I}) \end{aligned}$$

where the equality uses the fact that there is no correlation, the first inequality is due to the fact that $x_j^{\text{CAW}(S, \mathcal{I})}$ constrained in $j \in S \setminus \{i\}$ is a valid solution for program (expost-CAW) with $S \setminus \{i\}$ and \mathcal{I} as input, and the last inequality is because of the fact that the welfare does not decrease by introducing an additional buyer since $\vec{0} \in \mathcal{P}_i$. In addition,

$$\begin{aligned} \alpha(S, \mathcal{I} | S, \mathcal{I}) &= \sum_j \alpha(S, \mathcal{I} | \{j\}, \mathcal{I}) \\ &\leq \sum_j p_j \cdot \mu_j \cdot |\mathcal{I}| \leq \max_j p_j \cdot |\mathcal{I}| \leq p_i \cdot |\mathcal{I}| \end{aligned}$$

where the second-to-last inequality is due to the fact that $\sum_j \mu_j = \sum_j a_j \leq 1$. Therefore, we have $\alpha(S, \mathcal{I} | S \setminus \{i\}, \text{CP}_i(\mu_i)) \leq p_i \cdot \mu_i \cdot |\mathcal{I}|$. To sum up, we have

$$\begin{aligned} \text{CAW}(S, \mathcal{I}) &= \alpha(S, \mathcal{I} | \{i\}, \mathcal{I}) + \alpha(S, \mathcal{I} | S \setminus \{i\}, \text{CP}_i(\mu_i)) \\ &\quad + \alpha(S, \mathcal{I} | S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i)) \\ &\leq 2(p_i \cdot \mu_i \cdot |\mathcal{I}| + \text{SOL}(S \setminus \{i\}, \mathcal{I} \setminus \text{CP}_i(\mu_i))) \\ &= 2\text{SOL}(S, \mathcal{I}) \quad \square \end{aligned}$$

Theorem 2. *The expected revenue of the sequence of deals found by the \mathcal{AAG} algorithm with \mathcal{CAO} is at least $\frac{1}{3}\text{CAW}(S, \mathcal{I})$ with additive correlation.*

4.2 General Correlation

However, when the correlation can be arbitrary, we show that even in an environment without any constraints, the performance of preferred deals can be arbitrarily bad.

Theorem 3. *There exists an instance in which preferred deals can generate revenue at most $O(1/n)$ of the maximum social welfare even without any constraints.*

Note that it is trivial to achieve a $1/n$ -approximation by offering the buyer $i^* = \arg \max_i \mathbb{E}[v_i \vec{1}]$ a deal with price $\mathbb{E}[v_i \vec{1}]$ and minimum purchase requirement 1, where $\vec{1}$ always allocates the impressions: $\vec{1}(k) = 1$ for all k .

Proof of Theorem 3. We create an example as follows. Let z to be some positive real number. For each $0 \leq \ell < n$, we create $n \cdot z^{1-\frac{1}{2^\ell}}$ impressions and denote by Q_ℓ the set of these impressions. Among Q_ℓ , for each $1 \leq i \leq n$, there are $z^{1-\frac{1}{2^\ell}}$ impressions such that buyer i has valuation $z \cdot z^{\frac{1}{2^\ell}}$ while

the other buyers have valuation $z^{\frac{1}{2^{\ell+1}}} + 1$. The maximum social welfare is $\mathbb{E}[\max_i v_i] = \sum_{\ell=0}^{n-1} n \cdot z^{1-\frac{1}{2^\ell}} \cdot z^{\frac{1}{2^\ell}} = n^2 \cdot z$.

However, notice that for any buyer, they will cherry-pick the impression in the order of Q_0, Q_1, \dots, Q_{n-1} . Given any deal, assume the first impression buyer i picks (the impression with highest v_i) is left_i and the last impression (the impression with lowest v_i) is right_i . Let $A_\ell = |\{\text{left}_i \mid \text{left}_i \in Q_\ell\}| + |\{\text{right}_i \mid \text{right}_i \in Q_\ell\}|$ be the number of first or last impressions in Q_ℓ . Note that if $A_\ell = 0$, all impressions in Q_ℓ is chosen by only one buyer, which contributes revenue at most $z^{1-\frac{1}{2^\ell}} \cdot z^{\frac{1}{2^\ell}} + (n-1) \cdot z^{1-\frac{1}{2^\ell}} \cdot (z^{\frac{1}{2^{\ell+1}}} + 1) \leq z + 2(n-1) \cdot z^{1-\frac{1}{2^{\ell+1}}}$. If $A_\ell > 0$, then in the best scenario, these A_ℓ buyers pick all $z^{1-\frac{1}{2^\ell}}$ impressions that they have valuation $z^{\frac{1}{2^\ell}}$ from Q_ℓ . Therefore, it can contribute at most $A_\ell \cdot z^{1-\frac{1}{2^\ell}} \cdot z^{\frac{1}{2^\ell}} + (n-A_\ell) \cdot z^{1-\frac{1}{2^\ell}} \cdot (z^{\frac{1}{2^{\ell+1}}} + 1) \leq A_\ell \cdot z + 2(n-1) \cdot z^{1-\frac{1}{2^{\ell+1}}}$. Note that $\sum_\ell A_\ell = 2n$. Thus, as z goes to ∞ , the ratio between the revenue generated by preferred deals and the maximum social welfare is $\lim_{z \rightarrow \infty} \frac{2n \cdot z + (n-1) \cdot \sum_{\ell=0}^{n-1} z^{1-\frac{1}{2^{\ell+1}}}}{n^2 \cdot z} = \frac{3}{n}$. \square

4.3 Efficient Implementation of the Oracle

Although \mathcal{AAG} algorithm with \mathcal{CAO} provides good approximation, the computational complexity of \mathcal{CAO} depends on the number of impressions, which could be huge in practice. To overcome this hurdle, we describe an efficient algorithm to implement \mathcal{CAO} . The complexity of our algorithm is polynomial in $\sum_i |I_i|$, the number of different impressions.

In fact, writing down the *ex-post* allocation rule $x_i : \mathcal{I} \rightarrow [0, 1]$ takes $O(|\mathcal{I}|)$ time, which is already unacceptable. To alleviate this problem, we consider an *interim* allocation rule $\chi_i : I_i \rightarrow [0, 1]$ for $1 \leq i \leq n$ such that $\chi_i(\kappa_i)$ represents the probability that buyer i receives impression $\kappa_i \in I_i$. Given an *ex-post* allocation rule x , we can compute the interim allocation rule χ by $\chi_i(\kappa_i) \cdot f_i(\kappa_i) \cdot |\mathcal{I}| = \sum_{k: \kappa_i = \kappa_i} x_i(k)$. We can now use the following program to compute the interim allocation rule that maximizes the constraint-aware social welfare:

$$\begin{aligned} \max \quad & \sum_i \sum_{\kappa_i \in I_i} v_i(\kappa_i) \cdot f_i(\kappa_i) \cdot \chi_i(\kappa_i) \\ \text{s.t.} \quad & \chi_i \in \mathcal{P}_i \quad \forall i \\ & \sum_{\kappa_i \in I_i} v_i(\kappa_i) \cdot f_i(\kappa_i) \cdot \chi_i(\kappa_i) \cdot |\mathcal{I}| \leq B_i \quad \forall i \\ & \chi \text{ is a feasible interim allocation rule} \end{aligned}$$

where χ is feasible if and only if there exists a feasible *ex-post* allocation rule x that can induce χ , where an *ex-post* allocation rule x is feasible if $\sum_i x_i(k) \leq 1$ for all $k \in \mathcal{I}$ and $x_i(k) \geq 0$ for all i and $k \in \mathcal{I}$. The oracle then returns $a_i = \sum_{\kappa_i \in I_i} \chi_i(\kappa_i) \cdot f_i(\kappa_i)$. We can now apply the techniques developed by Cai *et al.* [2012] to obtain an efficient separation oracle to verify whether χ is feasible. In theory, a program with an efficient separation oracle can be solved efficiently by the ellipsoid method, while in practice, constraint generation can perform well, as evidenced by our implementation in the next section.

5 Empirical Evaluation

We evaluate the performance of our framework by focusing on an environment with budget constraints only, and therefore, our constraint-aware social welfare is known as *liquid*

welfare in the literature [Dobzinski and Paes Leme, 2014]. We use data collected from the Google Ad Exchange (AdX) over a period of one day in summer 2018. The data consists of bids from real-time second price auctions for specific ad spaces online, known as inventory units. Because second-price auctions are truthful, we can construe the bids as the advertisers' values for each impression and use the data for the purpose of preferred deal optimization. An advertiser's bid for an inventory unit can vary throughout the day due to variations in context (e.g., the user viewing the ad).

5.1 Data Set and Experiment Setup

We run our experiment on 5 high-volume inventory units for the day in question. Recall that the complexity of our algorithm mainly depends on the number of buyer-bid pairs. Therefore, we discretize the bids to cents and only consider the top 50 most frequent buyer-bid pairs. In particular, for each inventory unit, we first compute the frequencies for distinct buyer-bid pairs. Then, we sort the pairs in decreasing order of frequency and retain the top 50 pairs. Finally, in our experiment, we take the first 100K auctions in which at least two of the top 50 buyer-bid pairs appear in the auction to form our instances. These are the impressions to be allocated via preferred deals. Our algorithm can scale well beyond 100K impressions and we choose 100K because it represents a sizable instance where it is still possible to exactly compute the maximum liquid welfare using LP (expost-CAW), so that we are able to report exact approximation numbers.

For each inventory unit, we conduct experiments parametrized by a budget ratio $r \in [0.1, 1.5]$. This parameter controls the ratio between the expected total budgets and the social welfare, which is the summation of highest bids over all auctions. For a fixed setting of r , we repeat the experiment 50 times and generate the budget as follows: for each run, (1) compute the contribution s_i to the social welfare for each buyer i , by summing buyer i 's bids over all auctions that it wins; (2) set buyer i 's budget to a value uniformly drawn from $[0, 2 \cdot s_i \cdot r]$, so that the mean of the generated budget is $s_i \cdot r$, proportional to the buyer's social welfare contribution. By our construction, as r increases, the budget constraint becomes looser and one should expect the liquid welfare benchmark to increase as r increases. Finally, for each budget ratio r and each method, we report the ratio between the average of the revenue (welfare) for the 50 runs and the social welfare.

We compare the auction-adjusted greedy algorithm with budget (AAGBudget) against three other methods. The first is the basic second-price auction. In a second-price auction with budget constraints, the buyer's submitted bid is adjusted by taking the minimum between his bid and his remaining budget. We call a second-price auction *naive* if it is without reserve (Naive SPA). In a second-price auction with *optimal* personalized reserves (Optimal SPA), we optimize reserve prices by finding the revenue-maximizing reserve for each buyer. We also consider the original auction-adjusted greedy algorithm (AAGNoBudget) proposed in [Mirrokni and Nazarzadeh, 2017], which ignores the budget.

Inventory ID	1	2	3	4	5
Liquid Welfare	72.1%	74.8%	74.1%	78.6%	83.7%
AAGBudget	71.8%	70.9%	73.8%	75.9%	78.8%
AAGNoBudget	44.4%	41.2%	40.2%	45.2%	54.5%
Naive SPA	5.5%, 99.7%	34.8%, 92.4%	22.2%, 93.8%	34.3%, 95.3%	33.8%, 94.8%
Optimal SPA	61.2%, 66.2%	44.9%, 86.7%	66.5%, 71.6%	50.9%, 87.1%	62.1%, 82.1%

Table 1: Revenue and welfare (respectively) as the fraction of the social welfare for $r = 1$. For the Liquid Welfare, AAGBudget, and AAGNoBudget, revenue is equal to welfare.

Revenue against social welfare



Welfare against social welfare




Figure 1: Revenue and Welfare Plot: the x -axis is for the budget ratio r and the y -axis is the ratio between the revenue (welfare) of different methods and the social welfare

5.2 Empirical Results

We implemented our algorithm and experiments in Python 2.7, using the Glop linear programming solver¹ to compute the liquid welfare within our algorithm. Each run of the experiment takes roughly 30 seconds on a single CPU. Figure 1 demonstrates the revenue and welfare trends for a single inventory unit as the budget ratio r is varied (trends were very similar across all inventory units), and Table 1 summarizes the results for $r = 1$. From Figure 1 we see that our algorithm’s (AAGBudget) revenue performance closely tracks the optimal revenue achievable as captured by the liquid welfare for all values of r . The original AAG algorithm that does not take budget into account (AAGNoBudget) performs poorly especially when the budgets are small, and improves as the budgets increase. This is because, when the budget is small, the AAGNoBudget algorithm might offer a deal that violates buyer’s budget. Given such a deal, the buyer may reject the deal, resulting in 0 revenue. This is why AAGBudget shows 0 revenue up to $r = 0.5$. Recall that for AAGBudget and AAGNoBudget, revenue is equivalent to welfare.

An opposite revenue trend holds for the second-price auction benchmarks: their revenue performance decreases as r increases. This is because when the budget is small, the second price auction may be able to exhaust all the budgets within 100K auctions and therefore approach the liquid welfare optimum. For second-price auctions, especially the one with optimal reserves, there is a trade-off between revenue and efficiency. To understand why welfare can be higher than liquid welfare for the second-price auctions, recall that liquid

welfare only provides an upper bound on the revenue but not the welfare. If auction prices are consistently low but values are high, it is possible to achieve a high total welfare beyond the available budget, while respecting budget constraints.

Table 1 provides results, indexed against the optimal social welfare (unconstrained by budget) to interpret both the revenue and social welfare levels together. When comparing revenue directly against liquid welfare (the revenue optimum), AAGBudget’s performance ranges from 94–99% of the optimum across inventory units, whereas AAGNoBudget ranges from 54–65% and Optimal SPA from 60–90%. Our algorithm’s revenue performance is consistently close to the liquid welfare benchmark across all values of r and outperforms all other algorithms, which demonstrates that our algorithm has a stable and much better performance in practice than the theoretical $\frac{1}{2}$ -approximation guarantee.

6 Conclusions

We studied the problem of preferred deals with general constraints and provided an approximation algorithm to maximize the revenue obtained from the deals. We also validated our algorithm with experiments on real data, which showed that its empirical performance far surpasses the theoretical guarantees. One interesting future direction is to improve the approximation ratio or prove a matching lower bound for preferred deals, even in an environment without any constraints. Another interesting question is to understand how to combine and optimize the various ways used to sell display ad impressions together, including reservation contracts, preferred deals, and open auction.

¹See <https://developers.google.com/optimization/lp/glop>.

References

- [Azar *et al.*, 2017] Yossi Azar, Michal Feldman, Nick Gravin, and Alan Roytman. Liquid price of anarchy. In *International Symposium on Algorithmic Game Theory*, pages 3–15. Springer, 2017.
- [Babaioff *et al.*, 2014] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *55th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE, 2014.
- [Benoît and Krishna, 2001] Jean-Pierre Benoît and Vijay Krishna. Multiple-object auctions with budget constrained bidders. *The Review of Economic Studies*, 68(1):155–179, 2001.
- [Bhattacharya *et al.*, 2010] Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. Incentive compatible budget elicitation in multi-unit auctions. In *Proceedings of the 21st annual Symposium on Discrete Algorithms*, pages 554–572. SIAM, 2010.
- [Cai *et al.*, 2012] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th annual ACM Symposium on Theory of Computing*, pages 459–478. ACM, 2012.
- [Che and Gale, 1998] Yeon-Koo Che and Ian Gale. Standard auctions with financially constrained bidders. *The Review of Economic Studies*, 65(1):1–21, 1998.
- [Che and Gale, 2000] Yeon-Koo Che and Ian Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92(2):198–233, 2000.
- [Choi *et al.*, 2017] Hana Choi, Carl F. Mela, Santiago Balseiro, and Adam Leary. Online display advertising markets: A literature review and future directions. *Columbia Business School Research Paper*, (18-1), 2017.
- [Digiday, 2018] Digiday. GDPR is giving new life to programmatic guaranteed deals. <https://digiday.com/media/gdpr-giving-new-life-programmatic-guaranteed-deals/>, 2018.
- [Dobzinski and Paes Leme, 2014] Shahar Dobzinski and Renato Paes Leme. Efficiency guarantees in auctions with budgets. In *International Colloquium on Automata, Languages, and Programming*, pages 392–404. Springer, 2014.
- [Dobzinski *et al.*, 2012] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.
- [Feige *et al.*, 2008] Uriel Feige, Nicole Immorlica, Vahab Mirrokni, and Hamid Nazerzadeh. A combinatorial allocation mechanism with penalties for banner advertising. In *Proceedings of the 17th International Conference on World Wide Web*, pages 169–178. ACM, 2008.
- [Ghosh *et al.*, 2007] Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. Computing optimal bundles for sponsored search. In *International Workshop on Web and Internet Economics*, pages 576–583. Springer, 2007.
- [Korula *et al.*, 2016] Nitish Korula, Vahab Mirrokni, and Hamid Nazerzadeh. Optimizing display advertising markets: Challenges and directions. *IEEE Internet Computing*, 20(1):28–35, 2016.
- [Lu and Xiao, 2017] Pinyan Lu and Tao Xiao. Liquid welfare maximization in auctions with multiple items. In *International Symposium on Algorithmic Game Theory*, pages 41–52. Springer, 2017.
- [Mirrokni and Nazerzadeh, 2017] Vahab Mirrokni and Hamid Nazerzadeh. Deals or no deals: Contract design for online advertising. In *Proceedings of the 26th International Conference on World Wide Web*, pages 7–14, 2017.
- [Pai and Vohra, 2014] Malleesh Pai and Rakesh Vohra. Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425, 2014.
- [PwC, 2018] PwC. IAB internet advertising revenue report. https://www.iab.com/wp-content/uploads/2018/05/IAB-2017-Full-Year-Internet-Advertising-Revenue-Report.REV_.pdf, 2018.