Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Envy-Free Sponsored Search Auctions with Budgets

Bo Tang, Jinshan Zhang
Dept. of Computer Science, University of Liverpool
Liverpool, United Kingdom
{Bo.Tang, Jinshan.Zhang}@liverpool.ac.uk

Abstract

We study the problem of designing envy-free sponsored search auctions, where bidders are budget-constrained. Our primary goal is to design auctions that maximize social welfare and revenue — two classical objectives in auction theory. For this purpose, we characterize envy-freeness with budgets by proving several elementary properties including consistency, monotonicity and transitivity. Based on this characterization, we come up with an envy-free auction, that is both social-optimal and bidder-optimal for a wide class of bidder types. More generally, for all bidder types, we provide two polynomial time approximation schemes (PTASs) for maximizing social welfare or revenue, where the notion of envy-freeness has been relaxed slightly. Finally, in cases where randomization is allowed in designing auctions, we devise similar PTASs for social welfare or revenue maximization problems.

1 Introduction

Sponsored search advertising via auctions is one of the most popular ways of Internet monetization, which account for a major part of search engines' revenue. Consequently, design and analysis of such auctions has drawn a lot of attention in artificial intelligence and electronic commerce. For example, in the pioneering work by Varian [2007] and Edelman et al. [2007], the "generalized second price" (GSP) auctions, used by Google, have been modeled as position auctions having several desirable properties. Variants of this model have been extensively studied from both theoretical and practical points of view [Kuminov and Tennenholtz, 2009; Graepel *et al.*, 2010; He *et al.*, 2013].

In practice, budget constraints are given by bidders to specify their monetary affordability. The issues arising from budgets have received some attention in prior study on auction design [Ashlagi *et al.*, 2010; Dobzinski *et al.*, 2012]. It has been observed that, in various settings, the imposition of bidders' budgets changes the problems dramatically. To see this change in sponsored search, we point out that an advertiser's utility function, which is the difference between her valuation for the ad slot and the money she pays, is no longer a continuous function of her payment.

Fairness is one of the most important criteria in auction design. In economic theory, fairness can be explained as a free market without discriminatory pricing. More precisely, buyers should be allocated their most desired items under their budget constraints. Unfortunately, auctions currently in use may produce unfair outcomes for budget-constrained bidders (e.g., simultaneous ascending auctions [Nisan *et al.*, 2009]). It has been pointed out that the lack of fairness may lead to worse customer experience and fewer subsequent purchases from the firm [Anderson and Simester, 2010].

In this paper, we adopt a concept of fairness called envy-freeness in auctions. Envy-freeness is a classical criterion for analyzing mechanisms [Foley, 1967], which has been widely studied in artificial intelligence [Bouveret and Lang, 2005; Othman and Sandholm, 2010]. An outcome is envy-free if no buyer can improve her utility via exchanging her allocated items and payment with others. In sponsored search, envy-freeness is equivalent to the above explanation of fairness in a free market.

Our key objective is to design "optimal" envy-free auctions with respect to two classical objective functions in auction theory — social welfare and revenue. To this end, we characterize envy-freeness with budget-constrained bidders by proving several elementary properties including consistency, monotonicity and transitivity for all envy-free outcomes. Based on this characterization, we have designed an envy-free auction which is both social-optimal (maximizing social welfare) and bidder-optimal (maximizing every bidder's utility) for a wide class of bidder types. More generally, for all bidder types, we devise two polynomial-time approximation schemes (PTASs) to maximize social welfare or revenue where the notion of envy-freeness has been relaxed slightly. Furthermore, we present similar PTASs for designing optimal randomized auctions.

Our work is closely related to existing results on designing envy-free auctions for budget-constrained bidders. We review them briefly and identify their differences from our work.

Ashlagi et al. [2010] first investigated the effect of budget constraints in designing sponsored search auctions. They showed that a modification of the Generalized English Auction introduced by Edelman et al. [2007] is envy-free and Pareto optimal (a weaker condition than social welfare maximizing). However, the mechanism they developed cannot provide any guarantee for social welfare or revenue.

Aggarwal et al. [2009] and Dütting et al. [2011] developed more expressive auctions with discontinuous utilities including the case with budgets. They presented envy-free and bidder-optimal mechanisms for matching markets that is a more general setting than sponsored search. Unfortunately, no bidder-optimal auction always maximizes social welfare even in the context of sponsored search (see Example 3.7).

Feldman et al. [2012] first studied the revenue maximization problem in designing envy-free auctions for budget-constrained bidders. However they worked in a multi-unit setting with identical items and multi-demand buyers.

Devanur et al. [2013] analyzed clinching auctions introduced by Ausubel [2004] in the context of sponsored search. They showed that, when all bidders have a common budget, the clinching auction is envy-free and approximates optimal social welfare and revenue within a constant factor. Instead, our results aim at optimizing social welfare and revenue without the common budget assumption.

Revenue maximization in envy-free auctions has been extensively studied in a more general setting. Guruswami et al. [2005] showed that it is APX-hard to maximize the revenue among envy-free outcomes in matching markets. After that, both positive [Chen and Deng, 2010] and negative results [Briest, 2008] have been proved for this problem without budget constraints. Relaxed notions of envy-freeness have been also used in solving several problems (e.g., [Cohler *et al.*, 2011], [Brânzei and Miltersen, 2013]).

Much effort has been made to design incentive compatible auctions for bidders with budgets. Dobzinski et al. [2012] generalized clinching auctions to multi unit environment with budget constraints. This result was also extended to ordinal environments including sponsored search scenario in [Goel *et al.*, 2012]. For multiple keyword sponsored search auctions with budgets, Colini-Baldeshi et al. [2012] developed a randomized incentive compatible and Pareto optimal mechanism. Dobzinski and Paes Leme [2014] designed truthful auctions to approximate "liquid welfare" within a constant factor in multi-unit environments. But these mechanisms may produce unfair outcomes for budget-constrained bidders.

When considering social welfare or revenue maximization, Incentive Compatibility and Envy-freeness are not compatible in several settings with budgets (see [Feldman *et al.*, 2012] and references therein). Example 3.7 shows that no truthful mechanism can always output social-optimal envy-free outcomes. Therefore, in this paper we forgo truthfulness and concentrate on envy-free outcomes.

2 Preliminaries

In sponsored search, m advertisement positions are allocated to n bidders. Let $N = \{1, \ldots, n\}$ and $M = \{1, \ldots, m\}$ denote the set of bidders and positions respectively. Each position $j \in M$ is characterized by a quality number q_j, representing the number of clicks it could provide. This model is able to describe the setting where the value of the click may depend on the position it appears. We assume positions are indexed in decreasing order of q_j, i.e. $q_1 \geq \ldots \geq q_m$. Each bidder $i \in N$ is associated with a pair (v_i, B_i), where v_i represents her monetary valuation per click and B_i is the maximum payment she could afford. Note that bidder i is only interested in getting a single position, which is also called *unit-demand*. So bidder i's valuation for position j is $v_i q_j$. We assume all v_i, B_i and q_j are integers. We also use $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$ to denote an instance of a sponsored search auction, where $\mathbf{v} = (v_1, \ldots, v_n)$, $\mathbf{B} = (B_1, \ldots, B_n)$ and $\mathbf{q} = (q_1, \ldots, q_m)$. An outcome of an auction can be represented by an allocation vector $\mathbf{x} = (x_1, \ldots, x_n)$ and a payment vector $\mathbf{p} = (p_1, \ldots, p_n)$, where x_i denotes the expected number of clicks allocated to bidder i and p_i is the corresponding payment. In deterministic auctions, an allocation vector $\mathbf{x}$ can be viewed as a matching between bidders and positions. That is, if bidder i is matched to position j then $x_i = q_j$. In randomized auctions, a bidder may be allocated a distribution over positions. More precisely, for each bidder i and position j, r_{ij} is the probability that i gets j such that $\sum_{i \in N} r_{ij} \leq 1$ and $\sum_{j \in M} r_{ij} \leq 1$. Thus, the allocation $x_i = \sum_{j \in M} r_{ij} q_j$ is the expected number of clicks that bidder i gets from the allocated distribution.

In this paper, we restrict our attention to outcomes $(\mathbf{x}, \mathbf{p})$ with the following properties: for all bidder $i \in N$, $p_i \leq B_i$ (*budget feasible*), $x_i v_i - p_i \geq 0$ (*individual rational*) and $p_i \geq 0$ (*no positive transfer*). So we omit the word "feasible" for brevity in the remainder of this paper.

To formalize fairness, we adopt the following notion of *envy-freeness*. A set of bidders $S \subseteq N$ is said to be *envy-free* in an outcome $(\mathbf{x}, \mathbf{p})$, if for any pair of bidders, $i, j \in S$, i does not envy j, i.e., $p_j \leq B_i$ implies $x_i v_i - p_i \geq x_j v_i - p_j$. An outcome is envy-free if all bidders are envy-free in it. Moreover, an auction is envy-free if its outcome for any instance is envy-free. In order to get approximation schemes, we also use a weaker notion of envy-freeness called *relaxed ϵ-envy-freeness*. An outcome $(\mathbf{x}, \mathbf{p})$ is said to be *relaxed ϵ-envy-free* if it is feasible and for any pair of bidders $i, j \in N$ such that $p_j \leq B_i - \epsilon$, $x_i v_i - p_i \geq x_j v_i - p_j - \epsilon$. We consider two classical objective functions — social welfare, and revenue. Given an outcome $(\mathbf{x}, \mathbf{p})$, the social welfare of the outcome is $\sum_{i \in N} v_i x_i$ and the revenue is $\sum_{i \in N} p_i$. We say that an envy-free auction A is social-optimal (or revenue-optimal) if for any instance, no envy-free auction can produce an outcome with higher social welfare (or revenue) than A. Note that our social-optimal envy-free auction may obtain lower social welfare than the optimal social welfare without envy-freeness.

3 Deterministic Auctions

In this section, we describe our approach to designing social-optimal or revenue-optimal deterministic auctions for budget-constrained bidders. Before presenting our auctions, we show several properties of envy-free outcomes.

Lemma 3.1 (Consistency). *For any envy-free outcome, the order of the allocation vector should be consistent with the order of the payment vector. That is, for any two bidders i and j, $x_i > x_j \Leftrightarrow p_i > p_j$ and $x_i = x_j \Leftrightarrow p_i = p_j$.*

Proof. We only show the proof for $x_i > x_j \Rightarrow p_i > p_j$; similar proofs can be derived for the other three cases. Assume the opposite that, there exists two bidders i and j such that

$x_i > x_j$ and $p_i < p_j$. It is easy to see j will envy i since $p_i < p_j \le B_j$ by budget feasibility. □

Intuitively, consistency means no bidder can get a better position than another without paying more. Furthermore, the bidders with the same budget can be characterized below.

Lemma 3.2 (Monotonicity). *Suppose bidders i and j have the same budget and i has higher value than j. Then i must get a better position and pay more than j in any envy-free outcome. That is, $v_i \ge v_j$ implies $x_i \ge x_j$ and $p_i \ge p_j$.*

Proof. Recall that the envy-free conditions for i and j are $v_i x_i - p_i \ge v_i x_j - p_j$ and $v_j x_j - p_j \ge v_j x_i - p_i$. By summing these two inequalities, we have $(v_i - v_j)(x_i - x_j) \ge 0$. Since $v_i \ge v_j$, we have $x_i \ge x_j$. Furthermore $p_i \ge p_j$ follows from the consistency (Lemma 3.1). □

Although the above lemma reveals the structures of envy-free outcomes for bidders with the same budget, we need the following transitivity of envy-freeness to characterize bidders with different budgets.

Lemma 3.3 (Transitivity). *Suppose that for three bidders i, j and k, $B_i = B_j$ and $v_i \ge v_j$. If $x_k < x_j$ or ($x_k = x_j$ and $v_k \le v_j$), then $\{i, j\}$ and $\{j, k\}$ are envy-free implies $\{i, k\}$ is envy-free.*

Proof. We first prove that i does not envy k. By consistency and monotonicity, we have $x_i \ge x_j \ge x_k$ and $p_i \ge p_j \ge p_k$. Due to the assumption that $\{j, k\}$ is envy-free, we have $v_j(x_j - x_k) \ge p_j - p_k$ by rearranging the terms. Since $v_i \ge v_j$ and $x_j \ge x_k$, we get $v_i(x_j - x_k) \ge v_j(x_j - x_k) \ge p_j - p_k$. So the utility of bidder i is $v_i x_i - p_i \ge v_i x_j - p_j \ge v_i x_k - p_k$. Therefore, i does not envy k. It remains to show k does not envy i. We consider two cases. (a) $p_j > B_k$, then we have $p_i \ge p_j > B_k$. So the payment of bidder i exceeds bidder k's budget and the envy-freeness follows. (b) $p_j \le B_k$: We first show $v_j \ge v_k$ in this case. By rearranging the terms in the envy-free conditions for j and k, we get $v_k(x_j - x_k) \le p_j - p_k \le v_j(x_j - x_k)$. If $x_j > x_k$, we have $v_j \ge v_k$ by the above inequality. Otherwise $v_j \ge v_k$ follows from the assumption that $x_j = x_k$ and $v_j \ge v_k$. Then we have $v_k(x_j - x_i) \ge v_j(x_j - x_i) \ge p_j - p_i$ since $v_k \le v_j$, $x_j \le x_i$ and $v_j x_j - p_j \ge v_j x_i - p_i$. By using the envy-freeness between j and k and rearranging the terms in the above inequalities, we have $v_k x_k - p_k \ge v_k x_j - p_j \ge v_k x_i - p_i$. This shows that k does not envy i. □

The next lemma shows that, in envy-free outcomes, we can restrict our attention to integer payments.

Lemma 3.4. *Suppose $(\mathbf{x}, \mathbf{p})$ is an envy-free outcome for budget constrained bidders. Then there exists a non-negative integer price vector $\mathbf{p}'$ such that $(\mathbf{x}, \mathbf{p}')$ is also an envy-free outcome and $p'_i \ge p_i$ for all bidder i.*

Proof. Recall that we assume all v_i, B_i and q_j are integers. Given any outcome $(\mathbf{x}, \mathbf{p})$, let $p'_i = \lceil p_i \rceil$, for any $i \in N$. It is easy to check that $(\mathbf{x}, \mathbf{p}')$ still satisfies all constraints and must be an envy-free outcome. □

3.1 Proportional Bidder Types

Here we use the properties of envy-free outcomes to design a social-optimal envy-free auction for proportional bidders. The bidders are called proportional if for any bidders $i, j \in N$, $B_i > B_j$ implies $v_i \ge v_j$. Thus, we can order the bidders such that $i < j$ implies $B_i \ge B_j$ and $v_i \ge v_j$. For convenience, we define $q_j = 0$ for all $j > m$ and $B_{n+1} = -1$.

Lemma 3.5. *In any envy-free outcome with proportional bidders, bidder i's minimum payment for getting position j is*

$$p_{ij}^{\min} \triangleq \min_{k>i}\{B_k + 1 + \sum\nolimits_{\ell=i}^{k-2} v_{\ell+1}(q_{j+\ell-i} - q_{j+\ell-i+1})\}$$

Proof. Given any envy free outcome $(\mathbf{x}, \mathbf{p})$, we have for any bidder i, $p_i \ge \min\{B_{i+1} + 1, p_{i+1} + v_{i+1}(x_i - x_{i+1})\}$ since $v_{i+1}x_{i+1} - p_{i+1} \ge v_{i+1}x_i - p_i$ if $p_i \le B_{i+1}$. The lemma follows by using the above rule from bidder n to bidder 1 inductively. For the base case, we have $p_n \ge 0$. For the inductive step, we have

$$\begin{aligned}
p_i &\ge \min\{B_{i+1} + 1, v_{i+1}(x_i - x_{i+1}) + p_{i+1}\} \\
&\ge \min\{B_{i+1} + 1, v_{i+1}(x_i - x_{i+1}) \\
&\quad + \min_{k>i+1}\{B_k + 1 + \sum\nolimits_{\ell=i+1}^{k-2} v_{l+1}(x_\ell - x_{\ell+1})\} \\
&= \min_{k>i}\{B_k + 1 + \sum\nolimits_{\ell=i}^{k-2} v_{\ell+1}(x_\ell - x_{\ell+1})\} \\
&\ge \min_{k>i}\{B_k + 1 + \sum\nolimits_{\ell=i}^{k-2} v_{\ell+1}(q_{j+\ell-i} - q_{j+\ell-i+1})\}
\end{aligned}$$

The equality comes from $\sum_{\ell=i}^{k-2} v_{\ell+1}(x_\ell - x_{\ell+1}) = 0$ when $k = i + 1$. The last inequality is from setting $x_\ell = q_{j+\ell-i}$ for all $\ell \ge j$ that minimizes $\sum_{\ell=i}^{k-2} v_{\ell+1}(x_\ell - x_{\ell+1})$. □

Our social-optimal auction for proportional bidders can be described as Auction 1. Besides social optimality, we are able to show that Auction 1 is also *bidder-optimal*. That is for any instance $\mathcal{I}$, no bidder can improve her utility $v_i x_i - p_i$ in any envy-free outcome.

Theorem 3.6. *For instances with proportional bidders, Auction 1 is envy-free, bidder-optimal and social-optimal.*

Proof. First, we show the outcome of Auction 1 is envy-free. For any two bidders i and i' such that $i < i'$, we use j and j' to denote the position they get in the outcome of Auction 1. It is not hard to see that $j < j'$ and $j - i \le j' - i'$ from the process of the auction. So it suffices to prove (a) $v_i q_j - p_i \ge v_i q_{j'} - p_{i'}$

Auction 1: Auction for Proportional Bidders

Input: An instance $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$
Output: An outcome $(\mathbf{x}, \mathbf{p})$
Initialize $\mathbf{x} = \mathbf{0}$, $\mathbf{p} = \mathbf{0}$ and $i = 1$;
for *Position* $j = 1$ **to** m **do**
  if $p_{ij}^{\min} \le B_i$ **then**
    Set $x_i = q_j$ and $p_i = p_{ij}^{\min}$;
    Consider next bidder by increasing i by one;

and (b) $v_{i'}q_{j'} - p_{i'} \geq v_{i'}q_j - p_i$ if $p_i \leq B_{i'}$. For the inequality (a), let $k > i'$ be the bidder such that $p_{i'} = p_{ij}^{\min}(k)$. So $v_i q_j - p_i$ is at least

$$
\begin{aligned}
& v_i q_j - \left(B_k + 1 + \sum_{\ell=i}^{k-2} v_{\ell+1}(q_{j+\ell-i} - q_{j+\ell-i+1})\right) \\
=& \sum_{\ell=i}^{k-2} (v_\ell - v_{\ell+1}) q_{j+\ell-i} + v_{k-1} q_{j+k-i-1} - B_k - 1 \\
\geq & \sum_{\ell=i}^{k-2} (v_\ell - v_{\ell+1}) q_{j'+\ell-i'} + v_{k-1} q_{j'+k-i'-1} - B_k - 1 \\
\geq & (v_i - v_{i'}) q_{j'} + \sum_{\ell=i'}^{k-2} (v_\ell - v_{\ell+1}) q_{j'+\ell-i'} \\
& + v_{k-1} q_{j'+k-i'-1} - B_k - 1 \\
=& (v_i - v_{i'}) q_{j'} + v_{i'} q_{j'} - p_{i'} = v_i q_{j'} - p_{i'}
\end{aligned}
$$

For the inequality (b), since $p_i \leq B_{i'}$, $i' - i = j' - j$. This is because for any bidder ℓ between i and i', we have $p_\ell \leq p_i \leq B_{i'} \leq B_\ell$ when Auction 1 tries to allocate position $j + \ell - i$ to bidder ℓ. Then by a simple calculation, we can show that p_i equals to $p_{i'} + \sum_{k=i}^{i'-1} v_{k+1}(q_{k+j-i} - q_{k+j-i+1}) \geq p_{i'} + v_{i'}(q_j - q_{j'})$. Envy-freeness follows by rearranging the terms.

Next we show Auction 1 is social-optimal. Suppose there exists an envy-free outcome $(\mathbf{x}', \mathbf{p}')$ such that some bidder can get a better position than $(\mathbf{x}, \mathbf{p})$ outputted by Auction 1. Let i be the first bidder (with the smallest index) who gets a better position. Let j be the position that i gets from $(\mathbf{x}', \mathbf{p}')$. By the process of Auction 1, for all position j' such that $q_{j'} > q_j$, either j' is allocated to a bidder preceding i or $p_{ij'}^{\min} > B_i$. Thus, there must exist a bidder i' who gets a position j' such that $q_{j'} < q_j$ in $\mathbf{x}'$ and $i' < i$, otherwise i cannot get a better position. So $v_{i'} \geq v_i$ and $B_{i'} \geq B_i$ by the order of bidders. Since $(\mathbf{x}', \mathbf{p}')$ is envy-free, we have $v_{i'} x'_{i'} - p'_{i'} \geq v_{i'} x'_i - p'_i$ and $v_i x'_i - p'_i \geq v_i x'_{i'} - p'_{i'}$ since $j' > j$ implies $p'_i \geq p'_{i'}$ by Lemma 3.1 (Consistency). By summing the above two inequalities, we have $(v_{i'} - v_i)(x'_{i'} - x'_i) \geq 0$. Since $x'_{i'} = q_{j'} < q_j = x'_i$, we must have $v_{i'} \leq v_i$. Combine this with $v_{i'} \geq v_i$, we have $v_i = v_{i'}$. Thus we can modify $(\mathbf{x}', \mathbf{p}')$ by swapping the allocation and payment of bidder i and i' without changing the social welfare. By repeating this modification, we get Auction 1 is social-optimal.

Finally, we show Auction 1 is bidder-optimal. Assume to the contrary that there exists another envy-free outcome where some bidder can improve her utility. Let bidder i be the first bidder who gets higher utility from $(\mathbf{x}', \mathbf{p}')$ than that from $(\mathbf{x}, \mathbf{p})$, the outcome of Auction 1. Let j and j' be the positions that i gets from $(\mathbf{x}, \mathbf{p})$ and $(\mathbf{x}', \mathbf{p}')$ respectively. Case 1: $j' \geq j$. By Lemma 3.5, $p'_i \geq p_{ij'}^{\min}$. So it suffices to show that $v_i q_j - p_i \geq v_i q_{j'} - p_{ij'}^{\min}$. For any $k > i$, bidder i's utility $v_i q_j - p_i$ is at least

$$
\begin{aligned}
& v_i q_j - \sum_{l=i}^{k-2} v_{l+1}(q_{j+l-i} - q_{j+l-i+1}) - B_k - 1 \\
=& \sum_{l=i}^{k-2} (v_l - v_{l+1}) q_{j+l-i} + v_{k-1} q_{j+k-i-1} - B_k - 1 \\
\geq & \sum_{l=i}^{k-2} (v_l - v_{l+1}) q_{j'+l-i} + v_{k-1} q_{j'+k-i-1} - B_k - 1 \\
=& v_i q_{j'} - \sum_{l=i}^{k-2} v_{l+1}(q_{j'+l-i} - q_{j'+l-i+1}) - B_k - 1
\end{aligned}
$$

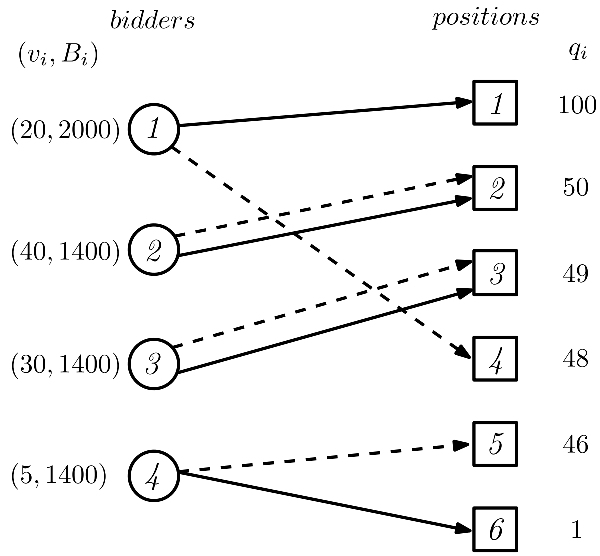

Figure 1: An illustration of Example 3.7. The left nodes represent bidders with value and budget pairs. The right squares are positions with qualities. The dashed arrows are the bidder-optimal allocations while the normal arrows are the social-optimal allocations.

By taking the minimum over all $k > i$, we prove $v_i q_j - p_i \geq v_i q_{j'} - p_{ij'}^{\min}$. Case 2: $j' < j$. Let bidder i' be the bidder who gets j' in $(\mathbf{x}, \mathbf{p})$. By the arguments in the previous paragraph, we have $i' < i$ and $v_i = v_{i'}$. Let u_i and u'_i denote the utility of bidder i in the outcome $(\mathbf{x}, \mathbf{p})$ and $(\mathbf{x}', \mathbf{p}')$. Since i is the first bidder who improves her utility, we have $u'_{i'} \leq u_{i'} = u_i < u'_i$ which contradicts the envy-freeness of $(\mathbf{x}', \mathbf{p}')$. Therefore Auction 1 is also bidder-optimal. ☐

Finally, we present an example to show if the bidders are not proportional, no bidder-optimal auction always maximizes social welfare.

Example 3.7. Consider four bidders with $v_1 = 20$, $v_2 = 40$, $v_3 = 30$, $v_4 = 5$, $B_1 = 2000$ and $B_2 = B_3 = B_4 = 1400$. There are also six positions with qualities $q_1 = 100$, $q_2 = 50$, $q_3 = 49$, $q_4 = 48$, $q_5 = 46$ and $q_6 = 1$. For convenience, we use B that equals 1400 to denote the budget of bidder 2. In order to compute the maximum social welfare, we consider two cases: **Case 1:** $p_1 \leq B$. By using Auction 1, we are able to show that the optimal allocation is $x_1 = q_4$, $x_2 = q_2$, $x_3 = q_3$ and $x_4 = q_5$, which gives the social welfare 4660. **Case 2:** $p_1 > B$. Suppose the allocation is $\mathbf{x}$, we can show that there exists a price vector $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is an envy-free outcome if and only if (i) $v_3(x_2 - x_3) + v_4(x_3 - x_4) \leq B$;(ii) $v_1(x_1 - x_2) + v_2(x_2 - x_3) + v_3(x_3 - x_4) + v_4 x_4 > B$; (iii) $x_1 \geq x_2 \geq x_3 \geq x_4$. By using these conditions, we can solve the social welfare maximization problem by simple calculations. As a result, the optimal social welfare is 5475 with the allocation $x_1 = q_1$, $x_2 = q_2$, $x_3 = q_3$ and $x_4 = q_6$.

This example shows that there does not exist a *bidder-optimal and social welfare maximizing auction.* To see this, note that in Case 2 where social welfare is maximized, bidder 1 can get utility at most $20 \cdot 100 - 1410 = 590$ (otherwise bidder 2 will envy bidder 1). However, in Case 1, we can compute an envy-free price vector for the bidders $\mathbf{p} = (p_1, p_2, p_3, p_4) = (10, 60, 30, 0)$ such that $(\mathbf{x}, \mathbf{p})$ is an envy-free outcome, where $(x_1, x_2, x_3, x_3) = (q_4, q_2, q_3, q_5)$. In this outcome, bidder 1 obtains utility $48 \cdot 20 - 10 = 950 > 590$. Hence, the social welfare maximizing outcome in this instance is not bidder-optimal. In fact, one can show that $(\mathbf{x}, \mathbf{p})$ is a bidder-optimal outcome for this instance. Similarly, it is easy to see bidder 1 can manipulate her budget or valuation to increase her own utility. That implies in this setting no truthful mechanism can always output social-optimal and envy-free outcomes.

3.2 All Bidder Types

In this section, we provide a polynomial-time approximation scheme (PTAS) for computing social-optimal or revenue-optimal deterministic envy-free auctions where the notion of envy-freeness has been relaxed slightly. More precisely, for any fixed parameter ϵ, our auction is a relaxed ϵB-envy-free and social-optimal auction that can be computed in time polynomial in n and m. The general idea of our PTAS can be described in four steps:

1. Partition the bidders into groups such that bidders with similar budget are in the same group.
2. Use Lemma 3.2 (Monotonicity) to characterize the structure of the bidders in the same group.
3. Use Lemma 3.3 (Transitivity) to establish the relations between different groups.
4. Compute the approximately-optimal envy-free auction via solving a dynamic programming.

For convenience, we say $\{i, k\}$ is ϵ-envy-free when $v_i x_i - p_i \geq v_i x_k - p_k - \epsilon$ if $p_k \leq B_i$ and $v_k x_k - p_k \geq v_k x_i - p_i - \epsilon$ if $p_i \leq B_k$ and the outcome is feasible for bidder i and bidder k. Given an instance $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$ and a parameter $\epsilon \in (0, 1)$, let $W = \{b \,|\, b = B_i$ for some bidder i or b is a multiple of $\epsilon B/n$ at most $B\}$. Clearly, the size of W is at most $n/\epsilon + n$.

We describe the auction in detail as follows. First of all, we round the bidders' budgets down to a closest multiple of ϵB, that is $B_i - \epsilon B \leq B'_i \leq B_i$. Then we solve the instance $\mathcal{I}' = (\mathbf{v}, \mathbf{B}', \mathbf{q})$ instead of $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$. For $\mathcal{I}'$, we partition the bidders into groups such that the bidders in the same group have the same budget. Clearly, the number of groups is at most $1/\epsilon$. For simplicity of presentation, we only present the auction as Auction 2 in which the bidders can be divided into two groups. It is straightforward to generalize the auction to many groups. After that, we sort bidders in each group with decreasing order of their valuations. We use G_1 and G_2 to denote bidders in these two groups and n_1 and n_2 be the size of G_1 and G_2 respectively. Let $R[i_1, i_2; j_1, j_2; w_1, w_2]$ be the optimal social welfare we can get from the first i_ℓ bidders in each group ℓ when the i_ℓth bidder in group ℓ gets position j_ℓ by paying w_ℓ for all $\ell = 1, 2$. For all $i_\ell = 0, \ldots, n_\ell$, $j_\ell = 0, \ldots, m$ and $w_\ell \in [0, B]$ that is a multiple of $\epsilon B/n$, we compute $R[i_1, i_2; j_1, j_2; w_1, w_2]$ in the following dynamic programming. Let bidder a_ℓ be the i_ℓth bidder in G_ℓ for $\ell = 1, 2$. So we have $x_{a_\ell} = q_{j_\ell}$ and

---

Auction 2: PTAS for Social-optimal Auctions

Input: An instance $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$ and a value $\epsilon \in (0, 1)$
Output: An outcome $(\mathbf{x}, \mathbf{p})$
$B \leftarrow \max_{i \in N} B_i$;
$W \leftarrow \{B_i\}_{i=1}^{n} \cup$ {multiples of $\epsilon B/n$ at most B};
Round B_i down to a multiple of ϵB for all $i \in N$;
Partition N into groups such that bidders in the same group have the same rounded budget;
/* *We only present the auction for two groups.* */
Sort bidders in each group with decreasing order of v_i;
$n_\ell \leftarrow$ the number of bidders in group ℓ;
Initialize $R[\mathbf{0}] = 0$;
for $i_1 = 0$ **to** n_1 *and* $i_2 = 0$ **to** n_2 *and* $i_1 i_2 \neq 0$ **do**
  for $j_1 = 0$ **to** m *and* $j_2 = 0$ **to** m **do**
    for $w_1 \in W$ *and* $w_2 \in W$ *in increasing order* **do**
      $a_\ell \leftarrow$ the i_ℓth bidder in group ℓ with $\ell = 1, 2$;
      $x_{a_\ell} \leftarrow q_{j_1}$ and $p_{a_\ell} \leftarrow w_\ell$;
      if (a_1, a_2) *are* $\epsilon B/n$*-envy-free and* $j_1 \neq j_2$ **then**
        Compute $R[i_1, i_2, j_1, j_2, w_1, w_2]$ by using $(\star)$;
      else $R[i_1, i_2, j_1, j_2, w_1, w_2] = -\infty$;;
SW $\leftarrow \max_{j_1, j_2, w_1, w_2}\{R[n_1, n_2, j_1, j_2, w_1, w_2]\}$;
Construct $(\mathbf{x}, \mathbf{p})$ by tracking back the computation of R;

$p_{a_\ell} = w_\ell$ for all $\ell = 1, 2$. If (a_1, a_2) are not $\frac{\epsilon B}{n}$-envy-free or $j_1 = j_2$, $R[i_1, i_2; j_1, j_2; w_1, w_2] = -\infty$.
Computation $(\star)$: If (a_1, a_2) are $\frac{\epsilon B}{n}$-envy-free and $j_1 \neq j_2$, consider two cases.
Case 1: $x_{a_1} < x_{a_2}$ or ($x_{a_1} = x_{a_2}$ and $v_{a_1} < v_{a_2}$). We set $R[i_1, i_2; j_1, j_2; w_1, w_2] = \max_{j < j_1, w \in W}\{R[i_1 - 1, i_2; j, j_2; w, w_2] + q_{j_1} v_{a_1} | (b, a_1)$ are $\frac{\epsilon B}{n}$-envy-free.$\}$ where b is the $(i_1 - 1)$th bidder in G_1 with $x_b = q_j$ and $p_b = w$.
Case 2: $x_{a_1} > x_{a_2}$ or ($x_{a_1} = x_{a_2}$ and $v_{a_1} \geq v_{a_2}$) We set $R[i_1, i_2; j_1, j_2; w_1, w_2] = \max_{j < j_2, w \in W}\{R[i_1, i_2 - 1; j_1, j; w_1, w] + q_{j_2} v_{a_2} | (b, a_2)$ are $\frac{\epsilon B}{n}$-envy-free.$\}$ where b is the $(i_2 - 1)$th bidder in G_2 with $x_b = q_j$ and $p_b = w$.
At last, the optimal social welfare is

$$\max_{j_1, j_2, w_1, w_2}\{R[n_1, n_2, j_1, j_2, w_1, w_2]\}.$$

Moreover the outcome $(\mathbf{x}, \mathbf{p})$ can be derived by tracking back the computation of R. It is not hard to see the same dynamic programming can be used to compute optimal revenue by changing $q_{j_1} v_{a_1}$ to w_1. So the outcome $(\mathbf{x}, \mathbf{p})$ is the output of our auction for the instance $\mathcal{I}$. Note that Auction 2 can be generalized when the number of groups is more than 2. So when we mention Auction 2 in next theorems, we refer the general auction for any number of groups. We summarize the result in the following theorem.

Theorem 3.8. *Given a parameter $\epsilon \in (0, 1)$, Auction 2 is a relaxed ϵB-envy-free and social-optimal auction. Moreover, the running time of the auction is $O((n^2 m/\epsilon)^{1/\epsilon})$.*

Proof. We first show that the outcome $(\mathbf{x}, \mathbf{p})$ outputted by Auction 2 is relaxed ϵB-envy-free. Since $B'_i \geq B_i - \epsilon B$, it suffices to show for any bidder i, k, $v_i x_i - p_i \geq v_i x_k - p_k - \epsilon B$ if $p_k \leq B'_i$ by the definition of relaxed ϵB-envy-freeness. It follows from a slight modification of Lemma 3.3 (Transitivity) for ϵ-envy-free transitive: given the same assumption, if i and j are ϵ_1-envy-free and j and k are ϵ_2-envy-free, we also have j and k are $(\epsilon_1 + \epsilon_2)$-envy-free. This is because the bidders in each group are considered in decreasing order of v_i and we check the $\epsilon B/n$-envy-freeness of the bidders that are considered consecutively in Auction 2.

Then we show Auction 2 is social-optimal. More specifically, no envy-free outcome can raise higher social welfare than the outcome of Auction 2. Given an envy-free outcome $(\mathbf{x}, \mathbf{p})$, we round p_i up to p'_i that is a value in W and round B_i down to B'_i that is a multiple of ϵB for all bidder $i \in N$. It is not difficult to see that any pair of bidders are $\frac{\epsilon B}{n}$-envy-free and $p'_i \leq B_i$. Hence the social welfare of $(\mathbf{x}, \mathbf{p}')$ is at most the social welfare of the outcomes outputted by Auction 2 by the optimality of the dynamic programming. So Auction 2 is social-optimal.

Finally, we estimate the running time of Auction 2. The total number of states in the dynamic programming is $(nm(n/\epsilon + n))^{1/\epsilon}$ and for each state we need $O(nm/\epsilon)$ to compute the entry. Therefore the overall running time is $(n^2 m(1 + 1/\epsilon))^{1+1/\epsilon} \approx (n^2 m/\epsilon)^{1/\epsilon}$ when ϵ is small. □

Corollary 3.9. *Given a parameter $\epsilon \in (0, 1)$, Auction 2 can be modified to a relaxed ϵB-envy-free and revenue-optimal auction with running time $O((n^2 m/\epsilon)^{1/\epsilon})$.*

4 Randomized Auctions

In this section, we study envy-free randomized auctions and present a PTAS for computing optimal envy-free outcomes. Unlike the deterministic case, the allocation of a randomized mechanism can be any real numbers. We also need that the allocation vector $\mathbf{x}$ can be implemented by a distribution of deterministic allocation rules. Formally, this condition is called the *majorization condition*, studied in [Feldman *et al.*, 2008; Goel *et al.*, 2012]. That is, a randomized allocation can be implemented iff the allocation vector is weakly majorized by the quality vector, i.e. $\sum_{i=1}^{k} x_i \leq \sum_{i=1}^{k} q_i$, for all $k = 1, \ldots, \min\{n, m\}$ where $\{x_1, \ldots, x_n\}$ are the bidders' allocations and sorted in decreasing order. Note that the allocations satisfying *majorization condition* form a convex set. However the set of all envy-free outcomes is not convex.

Example 4.1. Consider two bidders with $(v_1, B_1) = (1, 100)$ and $(v_2, B_2) = (2, 80)$ and one position with $q_1 = 100$. Consider the following two outcomes $x_1 = x_2 = 50, p_1 = p_2 = 50$ and $x'_1 = 100, x'_2 = 0, p'_1 = 100, p'_2 = 0$. It is easy to check that both outcomes are envy-free. However there is no pricing vector $\mathbf{p}$ such that $(\frac{\mathbf{x}+\mathbf{x}'}{2}, \mathbf{p})$ is an envy-free outcome.

Although the set of feasible outcomes is not convex, we are able to show that the set of all envy-free outcomes can be represented by a union of $n^{1/\epsilon}$ polytopes. Given an instance $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$ and a parameter $\epsilon \in (0, 1)$, let W be the set of all multiples of ϵB that is at most B. We use $d = 1/\epsilon$ to denote $|W|$ and $\{w_1, \ldots, w_d\}$ to denote the set W where the elements in W are ordered increasingly. For convenience, we add $w_0 = -1$ into W. Then we partition the set of bidders into d groups $\{G_1, \ldots, G_d\}$ such that for all bidder i in G_ℓ, $p_i \in (w_{\ell-1}, w_\ell]$. We use $\ell(i)$ to denote index of the group that i belongs to. We observe that the proofs of Lemma 3.1 and Lemma 3.2 do not use any fact about deterministic auctions. So the lemmas also hold for randomized auctions. By monotonicity, the number of all possible partitions for envy-free outcomes is at most $\binom{n}{d} \leq n^d$.

Given a partition $G = \{G_1, \ldots, G_d\}$, we are able to sort the bidders such that $i \leq j$ implies $p_i \geq p_j$. By consistency, we also have $p_i \geq p_j$ implies $x_i \geq x_j$. That is the majorization condition can be represented by n linear inequalities given the partition G. Moreover, the relaxed envy-freeness can be expressed as $v_i x_i - p_i \geq v_i x_j - p_j$ if $w_{\ell(j)} \leq B_i$ for all $i, j \in N$ since $p_j \leq B_i - \epsilon B$ implies $w_{\ell(j)} \leq B_i$. Note that the above conditions are also linear constraints given the partition G. So the set of all feasible envy-free outcomes can be represented by a polytope for a given partition. Let $f(x, p)$ be an objective function of $(\mathbf{x}, \mathbf{p})$, which is either the social welfare $\sum_i v_i x_i$ or the revenue $\sum_i p_i$. Given the partition G, we can compute the optimal objective via solving the following linear programming in polynomial time.

$$\begin{array}{lll} \max & f(x,p) & \\ \text{s.t.} & v_i x_i - p_i \geq v_i x_j - p_j & \forall i, j \text{ s.t. } w_{\ell(j)} \leq B_i \\ & \sum_{i=1}^{k} x_i \leq \sum_{i=1}^{k} q_j & \forall k \leq \min\{n, m\} \\ & p_i \in (w_{\ell-1}, w_\ell] & \forall \ell \leq d \text{ and } i \in G_\ell \\ & x_i \geq 0, p_i \in [0, B_i] & \forall i \in N \end{array}$$

Theorem 4.2. *Given an instance $\mathcal{I} = (\mathbf{v}, \mathbf{B}, \mathbf{q})$ and a parameter $\epsilon \in (0, 1)$, the randomized auction obtained via solving the above linear programming is a relaxed ϵB-envy-free and social-optimal (or revenue-optimal) auction running in time polynomial in $m, 1/\epsilon$ and $n^{1/\epsilon}$.*

Proof. By the optimality of linear programming, it suffices to show two claims. First, the solution computed by the linear program is relaxed ϵB-envy-free. Second, any envy-free outcome is a feasible outcome in the linear program. The first claim follows from the definition of relaxed ϵB-envy-freeness. For the second one, note that an envy-free outcome $(\mathbf{x}, \mathbf{p})$ is also relaxed envy-free. So the claim follows by dividing the bidders into groups according to $\mathbf{p}$. □

5 Conclusion

In this paper, we study the social welfare and revenue maximization problems in envy-free sponsored search auctions, for budget-constrained bidders. Our main contribution is a general approach to devise PTASs for maximizing social welfare and revenue in both deterministic and randomized settings. In addition, we provide an alternative understanding of the effects of budgets on envy-free auction design. There are several promising directions for future work. One of them is to generalize our auctions to more general settings, for instance matching markets and combinatorial auctions. Another promising direction is to consider fairness issues in some other settings such as TV advertising or labor markets.

Acknowledgment

We thank Paul W. Goldberg and the anonymous IJCAI reviewers for their useful comments. Jinshan Zhang is supported by EPSRC project EP/K01000X/1.

References

[Aggarwal *et al.*, 2009] G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. In *18th International World Wide Web Conference*, pages 241–241, April 2009.

[Anderson and Simester, 2010] E. T. Anderson and D. I. Simester. Price stickiness and customer antagonism. *Quarterly Journal of Economics*, 125(2), 2010.

[Ashlagi *et al.*, 2010] I. Ashlagi, M. Braverman, A. Hassidim, R. Lavi, and M. Tennenholtz. Position auctions with budgets: Existence and uniqueness. *The BE Journal of Theoretical Economics*, 10(1), 2010.

[Ausubel, 2004] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, pages 1452–1475, 2004.

[Bouveret and Lang, 2005] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods. In *Proceedings of the 19th international joint conference on Artificial intelligence*, pages 935–940. Morgan Kaufmann Publishers Inc., 2005.

[Brânzei and Miltersen, 2013] S. Brânzei and P. B. Miltersen. Equilibrium analysis in cake cutting. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 327–334. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[Briest, 2008] P. Briest. Uniform budgets and the envy-free pricing problem. In *Automata, Languages and Programming*, pages 808–819. Springer, 2008.

[Chen and Deng, 2010] N. Chen and X. Deng. Envy-free pricing in multi-item markets. In *Automata, Languages and Programming*, volume 6199, pages 418–429. Springer, 2010.

[Cohler *et al.*, 2011] Y. J. Cohler, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Optimal envy-free cake cutting. In *AAAI*, 2011.

[Colini-Baldeschi *et al.*, 2012] R. Colini-Baldeschi, M. Henzinger, S. Leonardi, and M. Starnberger. On multiple keyword sponsored search auctions with budgets. In *Automata, Languages, and Programming*, pages 1–12. Springer, 2012.

[Devanur *et al.*, 2013] N. R. Devanur, B. Q. Ha, and J. D. Hartline. Prior-free auctions for budgeted agents. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 287–304. ACM, 2013.

[Dobzinski and Paes Leme, 2014] S. Dobzinski and R. Paes Leme. Efficiency guarantees in auctions with budgets. In *Automata, Languages, and Programming*, pages 392–404. Springer, 2014.

[Dobzinski *et al.*, 2012] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[Dütting *et al.*, 2011] P. Dütting, M. Henzinger, and I. Weber. An expressive mechanism for auctions on the web. In *Proceedings of the 20th international conference on World wide web*, pages 127–136. ACM, 2011.

[Edelman *et al.*, 2007] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American economic review*, 97(1):242–259, 2007.

[Feldman *et al.*, 2008] J. Feldman, S. Muthukrishnan, E. Nikolova, and M. Pál. A truthful mechanism for offline ad slot scheduling. In *Algorithmic Game Theory*, pages 182–193. Springer, 2008.

[Feldman *et al.*, 2012] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 532–549. ACM, 2012.

[Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:4598, 1967.

[Goel *et al.*, 2012] G. Goel, V. Mirrokni, and R. Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 107–122. ACM, 2012.

[Graepel *et al.*, 2010] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 13–20, 2010.

[Guruswami *et al.*, 2005] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics, 2005.

[He *et al.*, 2013] D. He, W. Chen, L. Wang, and T.-Y. Liu. A game-theoretic machine learning approach for revenue maximization in sponsored search. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence(IJCAI)*, pages 206–212. AAAI Press, 2013.

[Kuminov and Tennenholtz, 2009] D. Kuminov and M. Tennenholtz. User modeling in position auctions: Reconsidering the gsp and vcg mechanisms. In *Autonomous Agents and Multiagent Systems*, AAMAS '09, pages 273–280, Richland, SC, 2009.

[Nisan *et al.*, 2009] N. Nisan, J. Bayer, D. Chandra, T. Franji, R. Gardner, Y. Matias, N. Rhodes, M. Seltzer, D. Tom, H. Varian, and D. Zigmond. Googles auction for tv ads. In *Automata, Languages and Programming*, volume 5556, pages 309–327. Springer, 2009.

[Othman and Sandholm, 2010] A. Othman and T. Sandholm. Envy quotes and the iterated core-selecting combinatorial auction. In *AAAI*, 2010.

[Varian, 2007] H. R. Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.