

Revenue Maximizing Envy-free Multi-unit Auctions with Budgets

Michal Feldman, The Hebrew University of Jerusalem and Harvard University

Amos Fiat, Tel Aviv University

Stefano Leonardi, Sapienza University of Rome

Piotr Sankowski, University of Warsaw and Sapienza University of Rome

We study envy-free (EF) mechanisms for multi-unit auctions with budgeted agents that approximately maximize revenue. In an EF auction, prices are set so that every bidder receives a bundle that maximizes her utility amongst all bundles; We show that the problem of revenue-maximizing EF auctions is NP-hard, even for the case of identical items and additive valuations (up to the budget). The main result of our paper is a novel EF auction that runs in polynomial time and provides a approximation of $1/2$ with respect to the revenue-maximizing EF auction. A slight variant of our mechanism will produce an allocation and pricing that is more restrictive (so called item pricing) and gives a $1/2$ approximation to the optimal revenue within this more restrictive class.

Categories and Subject Descriptors: J.4 [Social and behavioral sciences]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: mechanism design, multi-unit auction, revenue maximization, envy freeness

1. INTRODUCTION

This paper deals with multi-unit auctions with budgets. *I.e.*, selling multiple identical items to a variety of bidders, each with their own valuation per item and budget. This setting was studied by [Dobzinski et al. 2008] who gave a Pareto-optimal mechanism for this problem that is incentive compatible with respect to valuations. A natural example for this framework is online advertisement, where content pages (such as Yahoo or CNN) attempt to maximize the revenue collected from advertisers.

Price discrimination [Carlton and Perloff 2005] refers to any nonuniform pricing policy used by a firm with market power to maximize its profits. Price discrimination is profitable because consumers who value the good more are willing to pay more.

Price discrimination may have some drawbacks:

- (1) Customer resentment: In a 28 month study, covering 50,000 customers, [Anderson and Simester 2010] found that customers who felt cheated due to price discrimination “... react by making fewer subsequent purchases from the firm. The effect is largest among the firm’s most valuable customers: those whose prior purchases

This work was partially supported by the ERC StG project PAA1 no. 259515, by the Israel Science Foundation (grants number 1219/09 and 0603806441), by the Google Inter-university center for Electronic Markets and Auctions, and the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 274919.

Author’s addresses: M. Feldman, Harvard School of Engineering and Applied Sciences, Computer Science Department, 33 Oxford Street, Cambridge, MA 02138; A. Fiat, Tel Aviv University; S. Leonardi, Sapienza University of Rome; P. Sankowski, University of Warsaw and Sapienza University of Rome.

Permission to make digital or hard copies 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 credit is 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’12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

were most recent and at the highest prices.” Similar sentiment has been observed in queues, where people prefer longer waits and no queue jumping over shorter, but unfair, queues [Avi-Itzhak et al. 2007].

- (2) It may violate price discrimination laws: The Robinson-Patman Act of 1936 is a US federal law that prohibits certain forms of price discrimination by requiring that the seller offer the same prices to customers at a given level of trade. The European Community Competition Law also forbids some forms of price discrimination [Geradin and Petit 2006].

Unfortunately, the various auctions now in use, e.g., the VCG mechanism (not useful in the context of budgets), or the Google Auction of TV Ads [Nisan et al. 2009] (simultaneous ascending auction), all produce “unfair” prices. By unfair we mean that the same product gets sold at different prices. Moreover, if we insist on incentive compatible multi-unit auctions (with respect to valuation), the “unfair” [Dobzinski et al. 2008] auction is the only possibility.

Thus, we turn to multi-unit envy-free (EF) mechanisms that are not incentive compatible. Generally, envy-free pricing (with budgets) allows one to set arbitrary bundle prices. One can define the following envy-free notions, ordered from the most acceptable to the least acceptable in terms of “customer experience”:

- (1) Envy free item pricing: given valuations and budgets, the mechanism sets a price of p per item. An agent chooses whatever number of items she wants, at this price. *I.e.*, an agent with value v per item and budget b will purchase nothing if $v < p$, will purchase $\lfloor b/p \rfloor$ items if $v > p$, and will be indifferent to purchasing any $0 \leq i \leq \lfloor b/p \rfloor$ items if $v = p$.
Envy free item pricing must choose p so that there is no over demand. However, envy free prices need not (and indeed, in some cases, cannot) clear the market, so envy free pricing is not Walrasian pricing.
- (2) Envy free item pricing with *maximal quantity limits* — (h, p) pricing: like envy free item pricing, but no agent is allowed to purchase more than h items. *I.e.*, with valuation v per item and budget b an agent will choose to purchase no items if $v < p$ and will purchase $\min(h, \lfloor b/p \rfloor)$ items otherwise. If $v = p$ (the utility for the agent is zero), then the agent is indifferent between purchasing any $0 \leq i \leq \lfloor b/p \rfloor$ items. Here, (h, p) need to be such that the auctioneer will never run short of items.
- (3) Envy free item pricing with *minimum and maximal quantity limits* — (ℓ, h, p) pricing: The agent will purchase nothing if $v < p$ or if $b < \ell \cdot p$, and will purchase $\min(\lfloor b/p \rfloor, h)$ items otherwise. Agents with $v = p$ are indifferent between purchasing any $0 \leq i \leq \lfloor b/p \rfloor$ items. As before, (ℓ, h, p^*) are such that there is no over-demand.
- (4) Envy-free bundle pricing: On sale is a set of “shrinkwrapped” bundles of items, each with its own price tag. These bundle prices are envy-free in that every agent with budget b and per item value v can be assigned a “best possible” bundle B amongst bundles with $\text{price}(B) \leq b$ — maximizing $(\# \text{ items in } B) \cdot v - \text{price}(B)$ — without agents fighting over bundles. Agents whose best possible bundle leaves them with zero utility are indifferent between purchasing the bundle and not. Envy-free bundle pricing is more general than envy-free item pricing or its variants above. Throughout the paper, envy-free (EF) without additional quantifiers refers to envy-free bundle prices.
- (5) Arbitrary pricing: The pricing of bundles is arbitrary, different agents may get different prices for identical bundles. Agents can (and will) complain about unfairness, as loudly as possible.

Item pricing, (h, p) -pricing and (ℓ, h, p) -pricing are examples of *proportional bundle pricing*. I.e., where price per unit does not depend on the quantity purchased by agents. There are other more complicated examples of proportional pricing, e.g., when one can purchase only a prime number of items.

We remark that proportional pricing *does not* fall into any of the categories of price discrimination as defined in [Carlton and Perloff 2005] and as viewed under European Community law [Geradin and Petit 2006] (1st, 2nd and 3rd degrees of price discrimination).

Pricing schemes 1 – 4 all produce envy-free allocations, but may produce different revenue, and may not be equally acceptable to the consumer. Unfortunately, pricing schemes that offer more revenue are less acceptable: item pricing is a special case of (h, p) pricing, which is a special case of (ℓ, h, p) pricing, which is itself a special case of envy-free bundle pricing, which is (obviously) a special case of unfair pricing (arbitrary prices). Let $\text{Rev}(\alpha)$ denote the maximal revenue attainable by pricing schemes of type α , then we have:

$$\text{Rev}\left(\begin{array}{c} \text{EF} \\ \text{item} \\ \text{pricing} \end{array}\right) \leq \text{Rev}\left(\begin{array}{c} (h, p) \\ \text{pricing} \end{array}\right) \leq \text{Rev}\left(\begin{array}{c} (\ell, h, p) \\ \text{pricing} \end{array}\right) \leq \text{Rev}\left(\begin{array}{c} \text{EF} \\ \text{bundle} \\ \text{pricing} \end{array}\right) \leq \text{Rev}\left(\begin{array}{c} \text{“unfair”} \\ \text{pricing} \end{array}\right).$$

In Section 3, the above inequalities are shown to be strict. In particular, we show examples where:

- Item pricing obtains no more than $1/m$ of the revenue of the best (h, p) -pricing.
- (h, p) -pricing obtains no revenue, whereas (ℓ, h, p) -pricing gets revenue > 0 .
- Proportional pricing obtains no more than $1/2$ of the revenue of the best envy-free bundle pricing.

On the other hand, the main result of the paper is a construction of an (ℓ, h, p) -pricing scheme that gives a $1/2$ -approximation to any envy-free pricing scheme. Since an (ℓ, h, p) -pricing scheme is, in particular, a proportional scheme, the last point above implies that our main result is tight with respect to any proportional scheme. It is interesting that within the broad class of proportional pricing schemes, the simple form of (ℓ, h, p) -pricing is sufficient to establish a tight bound.

Whereas item pricing is the most commonly used pricing scheme in supermarkets, limits on the minimal and maximum quantities for sale are also frequent in brick and mortar shops, and thus it seems that they should be acceptable to customers in an online shop. Moreover, as the prices and limits magically ensure that there is no shortage of goods, customers should find little reason to complain.

One may intuitively understand the limit of no more than h items as a form of “wartime rationing”, there is not enough supply to meet all demand at the price chosen (creating envy), so the upper limit of h items reduces demand. The lower bound of ℓ can be interpreted as follows: there is too much demand from the very poor, and any envy free pricing will have too little revenue with them around. It is better (in terms of revenue) to remove the poor from the picture but setting a minimal quantity for sale — by definition, the poor cannot be envious (they cannot afford to purchase anything), so they can be ignored and then prices set to maximize revenue amongst the “rich”.

Envy-free allocations were defined by [Foley 1967] (see also [Varian 1974]). A key property of such allocations is that no one envies anyone else. Informally, no agent wants to switch his allocation with another agent.

Unfortunately, some confusion now exists with respect to the definition of an envy-free allocation. In recent literature concerned with revenue maximization of combinatorial auctions, envy-free pricing has been identified with item pricing ([Guruswami

et al. 2005] and much subsequent work). Identical items must be given a per item price, and the allocation to an agent must be a set of (identical or not) items that maximize the agent utility, given these item prices.

Envy-freeness can be viewed as a generalization of *Walrasian* equilibria. The differences are as follows: (i) Walrasian pricing is for individual items and bundle prices are implicit, see discussion above for envy-free variants (ii) envy-free allocations do not insist that the market clears, in the context of multi-unit auctions this means that there may be unsold items.

For multi-unit auctions without budget constraints (the standard quasi-linear model), it is trivial to give an envy-free allocation: give all m items to some agent with maximal valuation, (say v_i), set the price of these m items to be mv_i . It is easy to see that this is also the outcome that maximizes both social welfare and revenue.

The situation becomes significantly more difficult if agents have budget constraints.

Indeed, there are simple cases in which no envy-free auction can sell any item at all. For example, consider the case of $m + 1$ identical bidders, each of which has budget 1 and valuation 2. As all bidders are identical, and there are only m items for $m + 1$ bidders, there must be some bidder that gets nothing and this bidder will be envious of any other bidder that gets something. Ergo, the only envy-free allocation is to allocate nothing and extract no revenue.

1.1. Our Contributions

The envy-free revenue maximization problem seeks to find the [bundle priced] envy-free outcome that maximizes the seller's revenue.

We show the following:

- (1) We compute a price p so that selling at an [envy-free] item price of p gives revenue no less than $1/2$ the revenue of *any item pricing* (Section 4.1).
- (2) If the number of items is \geq the number of bidders, we compute h, p , so that the [envy-free] (h, p) -pricing scheme produces revenue within a factor of $1/2$ of *any envy-free bundle pricing* (Section 5).
- (3) If the number of items is less than the number of bidders, we compute ℓ, h, p , so that the [envy-free] (ℓ, h, p) -pricing scheme produces revenue within a factor of $1/2$ of *any envy-free bundle pricing* (Section 4).

We also show that the envy-free revenue maximization problem is NP-hard. This is in contrast to the same problem in a quasi-linear world, where finding the envy-free revenue-maximizing outcome is trivial. Surprisingly, even without requiring computational efficiency, this problem is far from being straightforward. We provide an algorithm that solves this problem in doubly exponential time; in particular, in $O(\text{poly}(n)2^{n+m})$ time (where the input size is $O(n + \log(m))$).

1.2. Related Work

Most of the recent work on envy-free allocations, and many of the hardness results, are in the context of item pricing [Guruswami et al. 2005; Demaine et al. 2006; Briest and Krysta 2006; Cheung and Swamy 2008; Balcan et al. 2008].

Determining the pricing for revenue optimal item priced envy-free allocations is hard for a wide variety of combinatorial auction settings, and there are various in-approximability results as well as a variety of approximation algorithms. Revenue optimal bundle priced envy-free allocations can be computed in polynomial time in some special cases (see [Fiat and Wingarten 2009]).

In many of the problems, one can distinguish between the limited supply [Guruswami et al. 2005; Cheung and Swamy 2008] setting (generally more difficult) and

the unlimited supply [Guruswami et al. 2005; Balcan et al. 2008; Demaine et al. 2006] (typically, easier). We consider the limited supply setting.

Other recent work on envy-free allocations show a strong connection between envy-freeness and incentive compatibility [Hartline and Yan 2011] in the context of auctions. Other papers that consider the combination of incentive compatibility and envy-freeness study *capacitated* agents (i.e., agents that cannot receive more than a limited number of objects) [Cohen et al. 2011].

[Cohen et al. 2010] study the problem of makespan minimization in job scheduling applications, and gave (an almost tight) logarithmic bounds (in the number of machines) on the approximation that can be achieved (even without considering computational issues).

Regarding multi-unit auctions with budgets, [Dobzinski et al. 2008] considered a setting identical to ours, with a focus on incentive compatibility. They showed that a variant of Ausubel’s clinching auction (in which the price gradually increases, and whenever the combined demand of the other bidders decreases strictly below available supply, then bidder i “clinches” the remaining quantity at the current price) is incentive compatible and Pareto-optimal, if the budgets are public knowledge. However, no incentive-compatible auction that always produces a Pareto-optimal allocation exists for the setting in which budgets and values are private information. These results were extended by [Fiat et al. 2011] to single-valued combinatorial auctions.

Most closely related to our work is the work by [Kempe et al. 2009] who defined the notion of *budget friendly* allocations. These are item prices, but with the following twist: agents can only be charged a payment that is strictly smaller than their budget; thus agent 1 cannot be considered to envy agent 2 if the price paid by agent 2 is equal to agent 1’s budget, irrespective of the valuation. For this notion, [Kempe et al. 2009] have shown how to compute maximum prices for fixed allocation.

2. PRELIMINARIES

A *Multi-unit auction with budgeted bidders* \mathcal{A} is formally depicted by the four-tuple

$$\mathcal{A} = \langle n, m, \mathbf{b}, \mathbf{v} \rangle,$$

where $n \in \mathbb{Z}_{>0}$ stands for the number of *bidders*, $m \in \mathbb{Z}_{>0}$ stands for the number of identical indivisible *goods*, $\mathbf{b} \in \mathbb{R}^n$ is a vector $\mathbf{b} = (b_1, \dots, b_n)$ of bidders’ budgets, and $\mathbf{v} \in \mathbb{R}^n$ is a vector $\mathbf{v} = (v_1, \dots, v_n)$ of bidders’ per-item values. Throughout the paper, we use the notation $[n]$ to denote the set $\{1, \dots, n\}$. We also denote the set of bidders by N .

A mechanism returns an *outcome* for every instance \mathcal{A} . Specifically, an outcome is depicted by a tuple $\langle \mathbf{k}, \mathbf{p} \rangle$, where $\mathbf{k} = (k_1, \dots, k_n)$ is an *allocation* vector, specifying for every bidder $i \in [n]$ the number of items, k_i , allocated to bidder i , and $\mathbf{p} = (p_1, \dots, p_n)$ is a payment vector, specifying for every bidder $i \in [n]$ her payment.

Bidders are assumed to have additive valuations up to their budget. Thus, given a multi-unit auction \mathcal{A} , and an outcome $\langle \mathbf{k}, \mathbf{p} \rangle$, the utility of bidder i is given by

$$u_i(k_i, p_i) = \begin{cases} v_i \cdot k_i - p_i & \text{if } p_i \leq b_i \\ -\infty & \text{otherwise} \end{cases}$$

A notion that will be used extensively in the sequel is the notion of a bidder’s *demand*. Given a multi-unit auction \mathcal{A} , we define two closely related notions of demand, as follows.

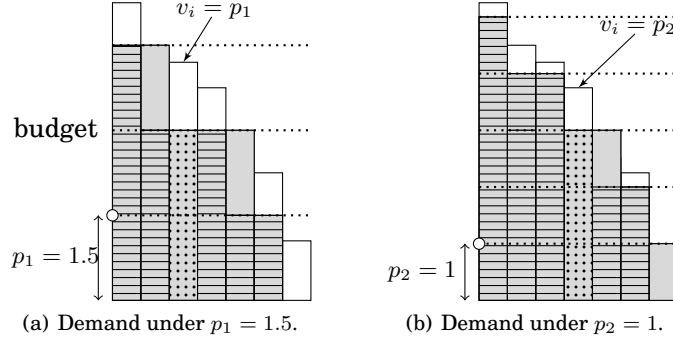


Fig. 1. An illustration of a multi-unit auction with 7 bidders. The x -axis represents the agents and the y -axis represents the agents' budget. The gray area is the (scaled by p) agent's demand ($D_i(p)$), the horizontal line pattern above represents the (scaled by p) agent's "demand-plus" ($D_i^+(p)$), and the dotted area is the (scaled by p) demand of value-limited agents (i.e., agents with $v_i = p$).

$$D_i[p, m, b] = \begin{cases} \min\{m, \lfloor b_i/p \rfloor\} & \text{if } p \leq v_i \\ 0 & \text{if } p > v_i \end{cases} \quad (1)$$

$$D_i^+[p, m, b] = \lim_{\epsilon \rightarrow 0^+} D_i[p + \epsilon, m, b]. \quad (2)$$

Essentially, $D_i[p, m, b]$ is equal to the number of items that bidder i wishes to purchase, and $D_i^+[p, m, b]$ is equal to the number of items bidder i would have wished to purchase if the price were increased by an infinitesimally small amount.

Given a price p , a bidder is said to be *value limited* (resp. *non-value limited*) if $v_i > p$ (resp. $v_i \leq p$).

Observe that $D_i^+[p, m, b] \leq D_i[p, m, b]$ for every bidder $i \in [n]$. For every non-value limited bidder $i \in [n]$ such that $b_i = cp$ for some $c \in \mathbb{Z}$, we have $D_i^+(p) = D_i(p) - 1$, whereas for non-value limited bidders such that $b_i \neq cp$ (for every $c \in \mathbb{Z}$), we have $D_i^+(p) = D_i(p)$.

The demands D_i and D_i^+ for value-limited and non value-limited agents, as a function of the price p , are illustrated in Figure 1.

The following definitions will be used in the analysis of the algorithms. Given a price p and an integer $k \in \mathbb{Z}$, let $A_k(p)$ be the set defined as $A_k(p) = \{i | b_i > kp \text{ and } v_i > p\}$, and let $B_k(p)$ be the set defined as $B_k(p) = \{i | b_i \geq kp \text{ and } v_i > p\}$. When clear in the context, we simply write A_k and B_k .

Given a multi-unit auction \mathcal{A} , an outcome is said to be *envy-free bundle pricing* (or simply EF) if for every pair of bidders $i, j \in [n]$, it holds that $u_i(k_i, p_i) \geq u_i(k_j, p_j)$; that is, no bidder wishes to switch her outcome with that of another.

An outcome is said to be *envy-free item pricing* if items are given item prices, and every bidder receives a bundle that maximizes her utility (given the prices). In the case of multi-unit auctions, this means that all items have the same per-item price p , and every bidder receives exactly $D_i(p)$ items.

The following observation can be easily verified.

PROPOSITION 2.1. *Given a price p , if all items are sold at a per-item price of p , then: (i) a bidder i who receives $D_i(p)$ items does not envy any other bidder, and (ii) a value-limited bidder does not envy any other bidder.*

Given an outcome (\mathbf{k}, \mathbf{p}) , the revenue of the auctioneer is given by the total payments of the bidders; i.e., $\sum_{i \in [n]} p_i$.

This raises the *revenue-maximization* problem: given a multi-unit auction \mathcal{A} , construct an EF outcome $\langle \mathbf{k}, \mathbf{p} \rangle$ that maximizes the auctioneer's revenue (among all EF outcomes).

A related optimization problem is the following: given a multi-unit auction \mathcal{A} and an allocation $\mathbf{k} = (k_1, \dots, k_n)$, construct a payment vector $\mathbf{p} = (p_1, \dots, p_n)$ that maximizes the auctioneer's revenue, under the constraint that the outcome $\langle \mathbf{k}, \mathbf{p} \rangle$ is EF. If no such payment vector exists, report infeasibility.

3. SEPARATION EXAMPLES

In this section we demonstrate the limitations and strengths of the different pricing schemes.

PROPOSITION 3.1. *For any $\epsilon > 0$, there exist bidders for which the best item pricing achieves no more than an ϵ fraction of the best (h, p) pricing.*

PROOF. Given ϵ , consider a multi-unit auction with m items and $n = m$ agents, such that $m > 4/\epsilon$. Let $m - 1$ agents have valuation 1.5 and budget 1, and let the last agent have valuation 2 and budget 4.

For this auction the optimal (h, p) -pricing is to set $h = 1$ and $p = 1$, extracting revenue m .

Any item pricing has to set the price per item to be > 1 — otherwise demand exceeds supply. However, this means that agents with budget equal to 1 cannot buy anything so the total revenue is ≤ 4 .

As $4/m < \epsilon$, we're done. \square

PROPOSITION 3.2. *For any $\epsilon > 0$, there exist bidders for which the best (ℓ, h, p) pricing achieves no more than an ϵ fraction of the best (ℓ, h, p) pricing.*

PROOF. Consider a multi-unit auction with four agents and two items, where three agents have value $v = 1$ and budget $b = 0.9$, and one agent has value $v = 0.9$ and budget $b = 1.8$. In this case the optimal (ℓ, h, p) -pricing is to set $\ell = 2$, $h = \infty$ and $p = 0.9$, extracting a revenue of 1.8. On the other hand, there exists no (h, p) -pricing that can extract any revenue, because for $p \leq 0.9$ the total demand is always higher than 2, whereas for $p > 0.9$, it is equal to zero. \square

Finally, we show that one cannot approximate a general envy-free pricing to within a factor better than $1/2$ by any proportional pricing. In this sense, our pricing schemes are the best possible.

PROPOSITION 3.3. *There exists no proportional pricing that approximates the revenue-maximizing envy-free bundle pricing to within a factor $> 1/2$.*

PROOF. Consider a multi-unit auction with m items and $n = m - 1$ agents, where valuations and budgets are as follows:

- Let A be a set of $m - 2$ agents with valuation and budget equal to 1, and
- Let the last agent, agent $m - 1$, have valuation and budget equal to m .

In this case the optimal envy-free allocation can extract revenue of $2m - 2$ by selling one item to each agent in A at a price of 1, and two items to agent $m - 1$ at a price of m .

On the other hand, any proportional pricing scheme can collect revenue of at most m : either by **(a)** selling one item at a price of 1 to each of the agents in A , and selling agent i 2 items at a (bundle) price of 2, or by **(b)** selling 2 items to agent i at a total price of m , and selling nothing to the agents in A .

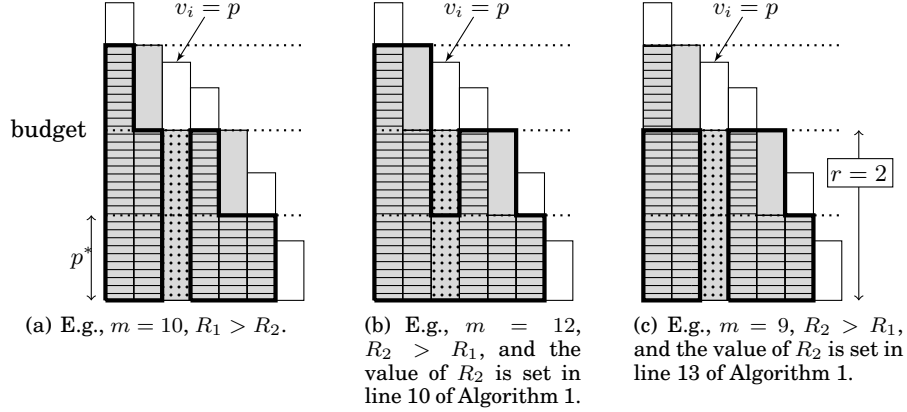


Fig. 2. Various possible outcomes of Algorithm 1, for a multi-unit auction with 7 agents. The envy-free allocation is represented by the thick border.

Hence, there exists no proportional price auction approximating the best envy free revenue to within a factor $\geq (\frac{1}{2} + \frac{1}{m})$. \square

4. $\frac{1}{2}$ -APPROXIMATE ALGORITHM FOR THE CASE OF MANY ITEMS ($m \geq n$)

In this section, we construct a $\frac{1}{2}$ -approximate algorithm for the problem of revenue-maximizing envy-free multi-unit auction, for the case where $m \geq n$. Algorithm 1 computes values h and p , the envy-free multi unit auction will be to use (h, p) -pricing: the per item price is p , but bidders may not purchase more than h items.

At the heart of this algorithm lies the price p^* , defined as

$$p^* = \min \left\{ p \geq 0 \mid \sum_{i=1}^n D_i^+(p) \leq m \right\}.$$

That is, p^* is the minimal price p such that the total $D_i^+(p)$ does not exceed supply. It follows that $\sum_{i \in [n]} D_i(p^*) > m$ and $\sum_{i \in [n]} D_i^+(p^*) \leq m$.

The algorithm computes two values, R_1 and R_2 , and chooses the larger of the two, obtaining revenue $\max(R_1, R_2)$. The pricing and allocation are done as follows:

— R_1 (set in line 6) is the revenue that is obtained by selling $D_i(p^* + \epsilon^*)$ items to bidder i , at a price of $p^* + \epsilon^*$ per item.

Note that this is equivalent to setting a $(h = \infty, p = p^* + \epsilon^*)$ pricing scheme. See Figure 2(a) for an illustration.

— R_2 can be one of two values, depending on the if statement in line 9 of the algorithm:

— if $\sum_{j \geq 1} t_j \leq m$, then R_2 is set in line 10, and the revenue R_2 is obtained by the allocation that sells all items at a price of p^* per item, as follows: non-value limited bidders get sold their complete demand ($D_i(p^*)$), and any left-over item is sold at a price of p^* to (arbitrary) value-limited bidders.

Note that this is a $(h = \infty, p = p^*)$ pricing scheme. See Figure 2(b) for an illustration.

— if $\sum_{j \geq 1} t_j > m$, R_2 is set in line 13, and the revenue R_2 is the result of the allocation that sells $\min(D_i(p^*), r)$ items to non-value limited bidders at a price of p^* per item, where r is set in line 12.

Note here that this is equivalent to using a $(h = r, p = p^*)$ pricing scheme. See Figure 2(c) for an illustration.

Algorithm 1 Multi-Unit Auction with Budgets ($m \geq n$)

1: **procedure** MULTI-UNIT AUCTION WITH BUDGETS(v, b, m)
 ▷ Input: $\mathcal{A} = \langle n, m, \mathbf{b}, \mathbf{v} \rangle$. Output: h, p for (h, p) pricing scheme.

2: Let

$$p^* = \min \left\{ p \geq 0 \mid \sum_{i=1}^n D_i^+(p) \leq m \right\}.$$

3: Set $R_1 = 0, R_2 = 0$.

4: **if** $\sum_{i=1}^n D_i^+(p^*) > 0$ **then** ▷ See Figure 2(a).

5: Let

$$\epsilon^* = \max \left\{ \epsilon > 0 \mid \sum_{i=1}^n D_i^+(p^*) = \sum_{i=1}^n D_i(p^* + \epsilon) \right\}.$$

6: Set

$$R_1 = (p^* + \epsilon^*) \cdot \sum_{i=1}^n D_i^+(p^*).$$

 ▷ Revenue R_1 is obtained setting an item price of $p^* + \epsilon^*$ ▷ — equivalently by a $(h = \infty, p = p^* + \epsilon^*)$ pricing scheme.

7: **end if**

8: Let

$$t_j = |\{i \mid D_i(p^*) \geq j, v_i \neq p^*\}|.$$

9: **if** $\sum_{j \geq 1} t_j \leq m$ **then**
 ▷ See Figure2(b), In total we will sell exactly m items using a $(h = \infty, p = p^*)$ pricing scheme.

10: Set

$$R_2 = m \cdot p^*.$$

11: **else**

12: Let $r = \max\{\ell \mid t_\ell > 0, \sum_{j=1}^\ell t_j \leq m\}$.

 ▷ this is an $(h = r, p = p^*)$ pricing scheme. See Figure2(c)

13: Set

$$R_2 = p^* \cdot \sum_{j=1}^r t_j.$$

14: **end if**

15: **end procedure**

The main result of this section is stated in both Theorem 4.1 that establishes the envy-freeness of the allocation, and in Theorem 4.2 that gives the approximation ratio.

THEOREM 4.1. *The allocation computed by Algorithm 1 is envy-free with (h, p^*) pricing.*

PROOF. Let us first consider the case when $R_1 \geq R_2$. The revenue R_1 (set in line 6) is obtained by selling $D_i(p^* + \epsilon^*)$ items to bidder i , at a fixed price of $p^* + \epsilon^*$ per item, where the demand of all bidders is saturated. Based on Observation 2.1, the outcome is envy-free without quantity limits.

If $R_2 > R_1$, then the allocation depends on the if statement in line 9 of Algorithm 1.

If the condition in line 9 is met, then the revenue R_2 is obtained by selling all items at a price of p^* . Non-value limited bidders (i.e., $v_i > p^*$) get their complete demand ($D_i(p^*)$), and any left-over item is sold at a price of p^* to (arbitrary) value-limited bidders. Based on Observation 2.1, the outcome is envy-free and does not require quantity limits.

If the condition in line 9 is not met (i.e., it holds that $\sum_{j \geq 1} t_j > m$), then R_2 is set in line 13 of Algorithm 1. The revenue R_2 is obtained by selling $\min(D_i(p^*), r)$ items to non-value limited bidders at a price of p^* per item, where r is set in line 12 — equivalently, we sell using a $(h = r, p^*)$ pricing scheme. In this case, the only way an agent i can get less items than another agent j is if agent i cannot afford the price that agent j is charged, therefore there is no envy. \square

We now turn to establish the desired approximation ratio,

THEOREM 4.2. *The outcome computed by Algorithm 1 is a $\frac{1}{2}$ -approximation to the optimal envy-free revenue.*

PROOF.

Let OPT be an optimal envy-free outcome. Recall that $A_k = \{i | b_i > kp^* \text{ and } v_i > p^*\}$, and $B_k = \{i | b_i \geq kp^* \text{ and } v_i > p^*\}$. We distinguish between the following two cases:

case (a): OPT allocates items only to bidders such that $b_i > p^*$ and $v_i > p^*$. Recall that this set is exactly the set A_1 . First, by the definition of R_1 we get:

$$R_1 > p^* \cdot \sum_{i \in [n]} D_i^+(p^*). \quad (3)$$

Now, observe that for every agent $i \in A_1$ we have

$$D_i^+(p^*) \geq 1. \quad (4)$$

We get:

$$\begin{aligned} OPT &\leq \sum_{i \in A_1} b_i \stackrel{[\text{by case (a)}]}{\leq} \sum_{i \in A_1} p^* \cdot (D_i^+(p^*) + 1) \\ &\stackrel{[\text{by Equation 4}]}{\leq} \sum_{i \in A_1} p^* \cdot 2D_i^+(p^*) \stackrel{[\text{by Equation 3}]}{<} 2R_1 \leq 2ALG, \end{aligned}$$

as required.

case (b): The optimal envy-free auction allocates at least one item to a bidder i such that $i \notin A_1$; i.e., $\min\{b_i, v_i\} \leq p^*$. Recall that B_k is defined as $B_k = \{i | b_i \geq kp^* \text{ and } v_i > p^*\}$, and that $t_k = |B_k|$.

In order to obtain a bound on R_2 we first claim that

$$R_2 \geq (m - t_2) \cdot p^*. \quad (5)$$

To prove this, consider two cases.

If $\sum_{j \geq 1} t_j > m$, then from the definition of r it follows that $\sum_{j \leq r} t_j + t_{r+1} > m$. We get $\sum_{j \leq r} t_j > m - t_{r+1} \geq m - t_2$, where the last inequality follows since for every $j \geq 2$, it holds that $t_j \leq t_2$. It follows that $R_2 \geq (m - t_2) \cdot p^*$, as required.

If $\sum_{j \geq 1} t_j \leq m$, all m items are sold at a price of p^* , thus we get $R_2 = m \cdot p^* \geq (m - t_2) \cdot p^*$, as desired.

We next consider the structure of the optimal allocation in case (b); i.e., where at least one item is sold to a bidder such that $\min\{b_i, v_i\} \leq p^*$. The price of this item cannot exceed p^* . Hence, any other bidder that is charged more than p^* must receive at least two items (otherwise, he is envious).

We consider that the maximum revenue that OPT can collect from agents in the different sets is as follows:

- For every $i \in A_1/B_1$, OPT can collect from i at most p^* by selling to i a single item.
- For every $i \in B_1/B_2$, OPT can collect from i at most $2p^*$ by selling to i two items, or at most p^* by selling to i a single item.
- For every $i \in B_2$, OPT can collect from i at least $2p^*$ by selling to i two items, but no more than $p^*(D_i^+(p^*) + 1)$.

It is easy to verify that the optimal allocation is never worse off by selling two items to every agent in B_2 first, then selling two items to every agent in B_1/B_2 . Hence, an upper bound on OPT is the sum of the revenue obtained by selling two items to each agent in B_2 (amounting to $\sum_{i \in B_2} p^* \cdot (D_i^+(p^*) + 1)$) and the revenue obtained by selling the remaining items, at a price of p^* per item (amounting to $(m - 2t_2) \cdot p^*$). It follows that

$$\begin{aligned} OPT &\leq \sum_{i \in B_2} p^* \cdot (D_i^+(p^*) + 1) + (m - 2t_2) \cdot p^* = \sum_{i \in B_2} p^* \cdot D_i^+(p^*) + t_2 p^* + (m - 2t_2) \cdot p^* \\ &= \sum_{i \in B_2} p^* \cdot D_i^+(p^*) + (m - t_2) \cdot p^* < R_1 + R_2, \end{aligned}$$

where the last inequality follows from Equations (3) and (5).

Since our algorithm chooses the maximum of R_1 and R_2 , it follows that $ALG = \max(R_1, R_2) \geq \frac{R_1 + R_2}{2} > OPT/2$, and the $\frac{1}{2}$ -approximation follows. \square

4.1. Approximating Optimal Item-Pricing

The algorithm above produces an allocation that achieves revenue at least 1/2 of any bundle-priced envy-free allocation. As expected, the allocation produced is not an item-priced allocation. The problem is the scenario where the high budget agents will not get their full demand.

This is not a problem for (bundle-priced) envy-freeness, there is no one to envy, but it is a problem for item-priced envy-free allocations, in which agents must get an allocation that maximizes their utility.

A slight modification of this algorithm will give us an item-priced envy-free allocation that approximates the best revenue amongst such allocation to within a factor of 1/2, irrespective of n vs. m .

In the context of item-priced allocations, it is obvious that no such allocation can set a price less than p^* (since the demand is too high). The only scenario in which the mechanism above does not produce an item-priced envy-free allocation is when $\sum_j t_j > m$. The modification to the algorithm would be that in this case, we will sell at an item price of $p^* + \epsilon$, in which case $\sum D_i^+(p^*)$ items will be sold. Note that we have enough items to do so.

This will no longer be a good approximation to the optimal (bundle-priced) envy-free allocation, but it will be a 1/2-approximation to the optimal item-priced envy-free allocation.

The reason is that no item-priced allocation can sell at a price of less than or equal to p^* (due to too high demand). Hence, any item priced allocation must give up on all

the agents for which $D_i^+(p^*) = 0$. When we compare our revenue to the budget of those agents for which $D_i^+(p^*) > 0$, we get at least $1/2$ of their budget.

5. $\frac{1}{2}$ -APPROXIMATE ALGORITHM FOR THE CASE OF MANY AGENTS ($n > m$)

In the case when $m \geq n$ we observed that $t_1 \leq m$ and so the value of R_2 was non-zero. In the general case, however, the value of R_2 in Algorithm 1 might be zero, and this might be a problem for the approximation ratio. This happens exactly for the example given in Proposition 3.2. One can easily verify that in this case, Algorithm 1 collects no revenue at all (in particular, $p^* = 0.9, t_1 = 3$, and consequently $R_1 = R_2 = 0$), whereas the optimal EF outcome earns revenue 1.8 by assigning two items to the unique agent at a total price of 1.8. In order to overcome this problem, a more sophisticated algorithm is in place.

Algorithm Multi-Unit-EF: At the heart of the algorithm is Algorithm 2, which is a generalization of Algorithm 1, where items are sold but in bundles of size at least k . In particular, our algorithm runs Algorithm 2 m times, for $k = 1, \dots, m$, and returns the maximum revenue achieved in any iteration.

In a given iteration k , the values of R_1^k and R_2^k are obtained from the following outcomes:

- The revenue R_1^k (set in line 6) is obtained by selling $D_i^+(p^*)$ items to bidders i for which $D_i^+(p^*) \geq k$, at a price of $p_1^* + \epsilon^*$ per item. This is an $(\ell = k, h = \infty, p = p^* + \epsilon^*)$ pricing scheme.
- R_2^k can be one of two values, depending on the if statement in line 9 of the algorithm:
 - if $k \cdot t_k + \sum_{j \geq k+1} t_j \leq m$, then R_2^k is set in line 14, where the revenue is obtained by the allocation that sells all items at a price of p^* per item, as follows: non-value limited bidders, with demand $D_i(p^*) \geq k$, get sold their complete demand ($D_i(p^*)$), and the left-over items are sold to value-limited bidders, who can afford k or more items, up to their budget, for a per-item price of p^* . One can easily verify that, by the definition of p^* , at the end of this process there are at most $k - 1$ unallocated items. This is equivalent to an $(\ell = k, h = \infty, p = p^*)$ pricing scheme.
 - if $\sum_{j \geq 1} t_j > m$, then R_2^k is set in line 17, where the revenue is the result of the allocation that sells $\min(D_i(p^*), r)$ items to non-value limited bidders with $D_i(p^*) \geq k$ at a price of p^* per item, where r is set in line 16. This is equivalent to a $(\ell = k, h = r, p = p^*)$ pricing scheme.

One may notice that Algorithm 1 is essentially a special case of Algorithm 2 with $k = 1$.

We now prove the following.

THEOREM 5.1. *The allocation computed by Algorithm Multi-Unit-EF is envy-free with (ℓ, h, p^*) pricing.*

PROOF. In order to prove that the algorithm is EF, we show that the outcome of the procedure is EF for every k . For simplicity of exposition, we omit the superscript k and write simply R_1 and R_2 . Let us first consider the case where $R_1 \geq R_2$. In this case, every agent with $D_i^+(p^*) \geq k$ gets her complete demand at a price of $p^* + \epsilon^*$ per item. On the one hand, no agent envies any other agent who gets less items, since all items are priced the same. On the other hand, no agent envies any other agent who gets more items, because this exceeds her budget affordable by the agent.

Assume now that $R_2 > R_1$. Then, the allocation we choose depends on the if statement in line 9 of Algorithm 2.

Algorithm 2 Multi-Unit Auction with Budgets (General case)

1: **procedure** MULTI-UNIT AUCTION WITH BUDGETS (GENERAL CASE)(v, b, m, k)2: **Let**

$$p^* = \min \left\{ p \geq 0 \left| \sum_{i \text{ s. t. } D_i^+(p) \geq k} D_i^+(p) \leq m \right. \right\}.$$

3: **Set** $R_1 = 0, R_2 = 0$.4: **if** $\sum_{i \text{ s. t. } D_i^+(p^*) \geq k} D_i^+(p^*) > 0$ **then**5: **Let**

$$\epsilon^* = \max \left\{ \epsilon > 0 \left| \sum_{i \text{ s. t. } D_i^+(p^*) \geq k} D_i^+(p^*) = \sum_{i \text{ s. t. } D_i^+(p^*) \geq k} D_i(p^* + \epsilon) \right. \right\}.$$

6: **Set**

$$R_1^k = \sum_{i \text{ s. t. } D_i^+(p^*) \geq k} (p^* + \epsilon^*) \cdot D_i^+(p^*).$$

7: **end if**8: **Let**

$$t_j = |\{i \mid D_i(p^*) \geq j, v_i \neq p^*\}|.$$

9: **if** $k \cdot t_k + \sum_{j \geq k+1} t_j \leq m$ **then**10: **Let** $e = k \cdot t_k + \sum_{j \geq k+1} t_j$ 11: **Define** i_q to be the q 'th index such that $D_{i_q}(p^*) \geq k, v_{i_q} = p^*$.12: **Let** q_{\max} be the number of indices for which this holds.13: **Let** q^* be the minimal index in $0, \dots, q_{\max}$ for which

$$\sum_{q=1, \dots, q^*} D_{i_q}(p^*) > m - e - k.$$

14: **Set**

$$R_2^k = p^* \cdot \min \left(m, e + \sum_{q=1, \dots, q^*} D_{i_q}(p^*) \right).$$

15: **else**16: **Let** $r = \max\{\ell \mid t_\ell > 0, k \cdot t_k + \sum_{j=k+1}^\ell t_j \leq m\}$.17: **Set**

$$R_2^k = p^* \cdot \left(k \cdot t_k + \sum_{j=k+1}^r t_j \right).$$

18: **end if**19: **end procedure**

If the condition in line 9 is true, then the revenue R_2 is obtained by selling all items at a price of p^* per item. Every non-value limited bidder i such that $D_i(p^*) \geq k$ receives her complete demand, $D_i(p^*)$; in this way, $e = k \cdot t_k + \sum_{j \geq k+1} t_j$ items are sold. Next, we sell the remaining $m - e$ items to value-limited agents with demand $D_i(p^*) \geq k$ one by one. However, we need to sell in bundles of at least k items, so when the number of

items left is less than k we cannot assign these items to the next value-limited agents. In this case at most $k - 1$ items are left unsold.

It follows from Observation 2.1 and from the indifference of the value limited agents that no one is envious. In addition, since every agent that receives any quantity > 0 gets at least k items, agents i such that $D_i(p^*) < k$ cannot envy anybody.

If the condition in line 9 does not hold; i.e., $\sum_{j \geq 1} t_j > m$, then R_2 is set in line 13 of Algorithm 1. The revenue R_2 is obtained by selling $\min(D_i(p^*), r)$ items to non-value limited bidders with $D_i(p^*) \geq k$ at a price of p^* per item, where r is set in line 16. In this case, the only way agent i can get less items than agent j is if agent i cannot afford the price that agent j is charged, therefore there is no envy.

Observe that this allocation is equivalent to a $(\ell = k, h = r, p = p^*)$ pricing. \square

We next establish the desired approximation ratio.

THEOREM 5.2. *The outcome computed by Algorithm Multi-Unit-EF achieves a $\frac{1}{2}$ -approximation to the optimal envy-free revenue.*

PROOF. Let OPT be an optimal envy-free outcome, and let k^* be the minimum number of items sold to any bidder in OPT. We claim that the outcome computed in phase k^* of our algorithm (i.e., for $k = k^*$) extracts at least $1/2$ of the revenue collected by OPT. Once again, for simplicity of presentation, we omit the superscript k^* and write R_1 and R_2 instead of $R_1^{k^*}$ and $R_2^{k^*}$, respectively. We distinguish between two cases.

case (a): OPT allocates items only to bidders such that $b_i > k^*p^*$ and $v_i > p^*$; i.e., to agents in A_{k^*} . In this case $OPT \leq \sum_{i \in A_{k^*}} b_i$, and by the definition of R_1 we get:

$$R_1 > p^* \cdot \sum_{i=1}^n D_i^+(p^*). \quad (6)$$

Observe that for every agent i in A_{k^*} , we have $D_i^+(p^*) \geq 1$, so we get:

$$\begin{aligned} OPT &\leq \sum_{i \in A_{k^*}} b_i \stackrel{\text{[by def. of } D_i^+]}{\leq} \sum_{i \in A_{k^*}} p^* (D_i^+(p^*) + 1) \\ &\stackrel{\text{[by } D_i^+(p^*) \geq 1]}{\leq} \sum_{i \in A_{k^*}} 2p^* D_i^+(p^*) \stackrel{\text{[by (6)]}}{\leq} 2R_1 \leq 2ALG. \end{aligned}$$

The $1/2$ -approximation follows.

case (b): OPT allocates items to a bidder i such that $b_i \leq k^*p^*$ or $v_i \leq p^*$. Note that, based on the definition of k^* , this means that bidder i receives at least k^* items in OPT.

In what follows, we make a few observations regarding the structure of the optimal allocation. Since k^* items are allocated to an agent i such that $b_i \leq k^*p^*$ or $v_i \leq p^*$, the price of this set of items cannot exceed k^*p^* . Hence, in order for OPT to collect more revenue than k^*p^* from any other bidder $j \neq i$, agent j must be allocated at least $k^* + 1$ items (otherwise, agent j will envy agent i). The maximal revenue OPT can collect from an agent i , depending on its type, is as follows:

- if $i \in N/B_{k^*}$, then OPT can get at most k^*p^* from i by selling her k^* items.
- if $i \in B_{k^*}/B_{k^*+1}$, then OPT can get at most $(k^* + 1)p^*$ from i by selling her $k^* + 1$ items, or at most k^*p^* by selling her k^* items.

— if $i \in B_{k^*+1}$, then OPT can get from i no more than $p^*(D_i^+(p^*) + 1)$.

Based on the last observation, OPT cannot be worse off by selling $k^* + 1$ items to every agent in B_{k^*+1} before selling $k^* + 1$ items to any agent in B_{k^*}/B_{k^*+1} . This is because any agent that gets $k^* + 1$ items should be charged the same. However, one can get at least as much from agents in B_{k^*+1} as from agents in B_{k^*} , because $p^*(D_i^+(p^*) + 1) \geq (k^* + 1)p^*$. Hence, we can assume without loss of generality that OPT sells $k^* + 1$ items to each agent in B_{k^*+1} , and sells the remaining items to other agents, collecting at most p^* per item. Consequently, we get:

$$OPT \leq \sum_{i \in B_{k^*+1}} p^* \cdot (D_i^+(p^*) + 1) + (m - (k^* + 1)t_{k^*+1}) \cdot p^*. \quad (7)$$

We distinguish between two cases (depending on the if statement in line 9). Again, we write R_1 (resp., R_2) instead of $R_1^{k^*}$ (resp., $R_2^{k^*}$) for simplicity.

case (b1): It holds that $k \cdot t_k + \sum_{j \geq k+1} t_j \leq m$. In this case it is easy to verify the following two lower bounds on R_2 :

$$R_2 \geq p^* \cdot \sum_{i \in A_{k^*}} D_i(p^*); \quad R_2 \geq (m + 1 - k^*) \cdot p^*.$$

Again, we distinguish between two cases.

If $t_{k^*+1} \geq 1$, then we get:

$$\begin{aligned} OPT &\leq \sum_{i \in B_{k^*+1}} p^* \cdot D_i^+(p^*) + p^* t_{k^*+1} + (m - (k^* + 1)t_{k^*+1}) \cdot p^* \\ &= \sum_{i \in B_{k^*+1}} p^* \cdot D_i^+(p^*) + (m - k^* t_{k^*+1}) \cdot p^* \leq 2R_2, \end{aligned}$$

where the first inequality follows by substituting $\sum_{i \in B_{k^*+1}} 1 = t_{k^*+1}$ into (7), and the last inequality follows by the two lower bounds on R_2 developed above and substituting $t_{k^*+1} \geq 1$.

Now assume that $t_{k^*+1} = 0$. If $m < 2k^*$, then our algorithm collects at least $k^* p^*$ from one agent, whereas OPT can collect at most $(k^* + 1)p^*$ from one agent.

On the other hand, when $m \geq 2k^*$, OPT sells items to at most $\lfloor \frac{m}{k^*} \rfloor$ agents, extracting at most $(k^* + 1)p^*$ from each one of them; we get:

$$\begin{aligned} OPT &\leq \left\lfloor \frac{m}{k^*} \right\rfloor (k^* + 1) \cdot p^* \leq \left\lfloor \frac{m(k^* + 1)}{k^*} \right\rfloor \cdot p^* = \left\lfloor m + \frac{1}{k^*} \right\rfloor \cdot p^* \\ &\leq (m + 1) \cdot p^* \stackrel{[\text{by } m \geq 2k^*]}{\leq} (2m - 2k^* + 1) \cdot p^* < 2R_2, \end{aligned}$$

where the last inequality follow from the second lower bound on R_2 . The 1/2-approximation follows.

case (b2): It holds that $k \cdot t_k + \sum_{j \geq k+1} t_j > m$. Recall that we are in the case where OPT sells k^* items to an agent i such that $b_i \leq k^* p^*$ or $v_i \leq p^*$. Observe that when agent i gets k^* items, all agents in B_{k^*} have to get k^* items as well; otherwise, some agent in B_{k^*} would envy i . Hence, there are more items than k^* multiplied by the number of agents in B_{k^*} ; i.e., $m \geq k^* \cdot t_{k^*}$. This implies that the value of r computed in line 16 is at least 1. On the other hand, for $j \geq k^* + 1$, we know that $t_j \leq t_{k^*+1}$. It follows that

$$R_2 \geq (m - t_{k^*+1}) \cdot p^*. \quad (8)$$

For integral $D_i^+(p^*)$ we have:

$$D_i^+(p^*) + 1 \leq k^* \cdot \left(\left\lfloor \frac{D_i^+(p^*)}{k^*} \right\rfloor + 1 \right).$$

Substituting the last inequality into (7) we obtain:

$$\begin{aligned} OPT &\leq \sum_{i \in B_{k^*+1}} p^* \cdot k^* \cdot \left(\left\lfloor \frac{D_i^+(p)}{k^*} \right\rfloor + 1 \right) + (m - (k^* + 1)t_{k^*+1}) \cdot p^* \\ &\stackrel{[\text{by } \sum_{i \in B_{k^*+1}} 1 = t_{k^*+1}]}{\leq} \sum_{i \in B_{k^*+1}} p^* \cdot k^* \cdot \left\lfloor \frac{D_i^+(p^*)}{k^*} \right\rfloor + t_{k^*+1} k^* p^* + (m - (k^* + 1)t_{k^*+1}) \cdot p^* \\ &= \sum_{i \in B_{k^*+1}} p^* \cdot k^* \cdot \left\lfloor \frac{D_i^+(p^*)}{k^*} \right\rfloor + (m - t_{k^*+1}) \cdot p^* < R_1 + R_2, \end{aligned}$$

where the last inequality follows from (6) and (8). The 1/2-approximation follows. \square

6. REVENUE MAXIMIZING ENVY-FREE MULTI-UNIT ALLOCATIONS

In this section we consider the problem of computing revenue-maximizing envy-free pricing, given an allocation of items to agents. We show that this can be done in polynomial time.

THEOREM 6.1. *For every multi-unit auction, given an allocation, the problem of finding the revenue-maximizing envy-free payments can be solved in $O(n^2)$ time.*

PROOF. This result follows from a result due to [Kempe et al. 2009], as follows. The paper by Kempe et al. considers the case of a unit-demand auction over heterogeneous items (i.e., where agents have different valuations for different items). They also consider the “budget friendly” scenario where agent 1 cannot envy agent 2 unless the budget for agent 1 is strictly larger than the payment made by agent 2. Fortunately, it is easy to adapt their algorithm to the envy-free setting, where one does not insist upon “strictly larger”, as above, but only upon “greater than or equal to”.

Once the allocation is fixed, we can restate the allocation as being an allocation of heterogeneous items, two allocations of a different number of items are viewed as distinct item types.

The Bellman-Ford like (but *poly time*) algorithm of [Kempe et al. 2009] allows us to compute a vector $p \in \mathbb{R}^n$, of maximal payments, in $O(n^2)$ time. I.e., p_i (the payment by agent i) is an upper bound on the payment by agent i in *any* envy-free allocation. Such prices clearly maximize revenue. \square

With Theorem 6.1 at hand, one can compute the revenue-maximizing outcome (i.e., both allocation and payment), by trying out all possible allocations, as summarized in the following theorem (whose proof is deferred to the full version of the paper):

THEOREM 6.2. *Given a multi-unit auction, the problem of finding the revenue-maximizing envy-free outcome can be solved in doubly exponential¹ $O(\text{poly}(n)2^{n+m})$ time.*

¹Since the input size is $O(n + \log(m))$, 2^m is doubly exponential in the input.

Note that the optimal pricing computed by the above procedure need not be proportional pricing.

7. NP-HARDNESS

In this section we establish the NP-hardness of the revenue-maximizing envy-free auction problem. Due to lack of space, we only outline the proof; the full proof is deferred to the full version.

THEOREM 7.1. *The problem of finding the revenue-maximizing envy-free auction in multi-unit auctions is NP-hard.*

The idea of the proof is to reduce the k -subset sum problem to the revenue maximizing envy-free allocation for multi-unit auction. In the k -subset sum problem we are given a multiset $X = \{x_1, \dots, x_l\}$ of positive integers and an integer s , and we are asked to find a k -element subset of X that sums up to s . We say that the k -subset sum problem is *uniform* if $\frac{s}{k+1} < x_i < \frac{s}{k-1}$, and $x_i \geq 5$, for $i = 1, \dots, l$. By scaling, one can show that the uniform k -subset sum problem is NP-hard as well.

Given a uniform instance of k -subset sum, we build an instance of the auction in the following way. For $i = 1, \dots, l$, let us define $s_0 = 0$ and $s_i = \sum_{j=1}^i x_j$. The instance has $n = (8k + 1)k$ agents. Let ϵ be some arbitrary, but very small, real number. For $i = 1, \dots, k$, we construct the following agent types:

- type A.** one agent A_i with budget $b_i^A = s_i$ and valuation $v_i^A = 1 + \epsilon$,
- type B.** $4k$ agents B_i with budget $b_i^B = (s_i - 1)(1 + \epsilon)$ and valuation $v_i^B = \infty$,
- type C.** $4k$ agents C_i with budget $b_i^C = s_i(1 + \epsilon)$ and valuation $v_i^C = 1 + \epsilon$.

Finally, we set the number of items to $m = (8k + 1) \left(\sum_{i=1}^k s_i \right) - 8ks_k + k + s$. One can show that there exists an envy-free allocation that extracts a revenue of at least m if and only if the uniform k -subset sum problem has a solution. One can essentially show that due to envy-freeness the players of type A can be assigned either $s_i - x_i$ or s_i items, whereas the other agents need to get in total $(8k + 1) \left(\sum_{i=1}^k s_i \right) - 8ks_k + k$ items. Hence, one needs to choose type A agents in such a way that the corresponding x_i 's sum up to s .

8. CONCLUSIONS AND OPEN PROBLEMS

In this paper we have studied the problem of envy-free pricing in multi-unit auctions. We have defined an hierarchy of envy-free pricing schemes, and studied the revenue that can be extracted by the different schemes. Our main result is an (ℓ, h, p) -pricing scheme that always extracts at least $1/2$ of the revenue that can be extracted by any general (bundle-pricing) envy-free scheme. We also show that this is tight with respect to proportional pricing schemes. Our analysis and results suggest some interesting directions for future research:

- Does the revenue-maximization envy-free pricing problem admit a PTAS?
- Does the revenue-maximization problem remain hard also with respect to proportional pricing schemes?
- Can we achieve similar approximation results for more complicated combinatorial auction structures?
- What other forms of interesting envy-free pricing schemes should be considered?

REFERENCES

- ANDERSON, E. T. AND SIMESTER, D. I. 2010. Price stickiness and customer antagonism. *The Quarterly Journal of Economics* 125, 2, 729–765.
- AVI-ITZHAK, B., LEVY, H., AND RAZ, D. 2007. A resource allocation queueing fairness measure: properties and bounds. *Queueing Syst. Theory Appl.* 56, 65–71.
- BALCAN, M.-F., BLUM, A., AND MANSOUR, Y. 2008. Item pricing for revenue maximization. *EC*.
- BRIEST, P. AND KRZYSTA, P. 2006. Single-minded unlimited supply pricing on sparse instances. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. SODA '06. ACM, New York, NY, USA, 1093–1102.
- CARLTON, D. W. AND PERLOFF, J. M. 2005. *Modern Industrial Organization*. Prentice Hall.
- CHEUNG, M. AND SWAMY, C. 2008. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. *FOCS*, 35–44.
- COHEN, E., FELDMAN, M., FIAT, A., KAPLAN, H., AND OLONETSKY, S. 2010. Envy-free makespan approximation: extended abstract. In *ACM Conference on Electronic Commerce*, D. C. Parkes, C. Dellarocas, and M. Tennenholtz, Eds. ACM, 159–166.
- COHEN, E., FELDMAN, M., FIAT, A., KAPLAN, H., AND OLONETSKY, S. 2011. Truth, envy, and truthful market clearing bundle pricing. In *WINE*, N. Chen, E. Elkind, and E. Koutsoupias, Eds. Lecture Notes in Computer Science Series, vol. 7090. Springer, 97–108.
- DEMAINE, E. D., HAJIAGHAYI, M. T., FEIGE, U., AND SALAVATIPOUR, M. R. 2006. Combination can be hard: approximability of the unique coverage problem. In *SODA*. ACM Press, 162–171.
- DOBZINSKI, S., LAVI, R., AND NISAN, N. 2008. Multi-unit auctions with budget limits. In *FOCS*. IEEE Computer Society, 260–269.
- FIAT, A., LEONARDI, S., SAIA, J., AND SANKOWSKI, P. 2011. Single valued combinatorial auctions with budgets. See Shoham et al. [2011], 223–232.
- FIAT, A. AND WINGARTEN, A. 2009. Envy, multi envy, and revenue maximization. See Leonardi [2009], 498–504.
- FOLEY, D. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7, 45–98.
- GERADIN, D. AND PETIT, N. 2006. Price discrimination under ec competition law: Another antitrust doctrine in search of limiting principles? *Journal of Competition Law and Economics* 2, 3, 479–531.
- GURUSWAMI, V., HARTLINE, J. D., KARLIN, A. R., KEMPE, D., KENYON, C., AND MCSHERRY, F. 2005. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. SODA '05. 1164–1173.
- HARTLINE, J. AND YAN, Q. 2011. Envy, truth, and profit. See Shoham et al. [2011], 243–252.
- KEMPE, D., MU'ALEM, A., AND SALEK, M. 2009. Envy-free allocations for budgeted bidders. See Leonardi [2009], 537–544.
- LEONARDI, S., Ed. 2009. *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*. Lecture Notes in Computer Science Series, vol. 5929. Springer.
- NISAN, N., BAYER, J., CHANDRA, D., FRANJI, T., GARDNER, R., MATIAS, Y., RHODES, N., SELTZER, M., TOM, D., VARIAN, H. R., AND ZIGMOND, D. 2009. Google's auction for tv ads. In *ICALP (2)*, S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. E. Nikolettseas, and W. Thomas, Eds. Lecture Notes in Computer Science Series, vol. 5556. Springer, 309–327.
- SHOHAM, Y., CHEN, Y., AND ROUGHGARDEN, T., Eds. 2011. *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*. ACM.
- VARIAN, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 63–91.