

Walrasian Pricing in Multi-Unit Auctions*

Simina Brânzei¹, Aris Filos-Ratsikas², Peter Bro Miltersen³, and Yulong Zeng⁴

- 1 Hebrew University of Jerusalem, Jerusalem, Israel
simina.branzei@gmail.com
- 2 Oxford University, Oxford, United Kingdom
aris.filos-ratsikas@cs.ox.ac.uk
- 3 Aarhus University, Aarhus, Denmark
bromille@cs.au.dk
- 4 Tsinghua University, Beijing, China
cengyl13@mails.tsinghua.edu.cn

Abstract

Multi-unit auctions are a paradigmatic model, where a seller brings multiple units of a good, while several buyers bring monetary endowments. It is well known that Walrasian equilibria do not always exist in this model, however compelling relaxations such as Walrasian envy-free pricing do. In this paper we design an optimal envy-free mechanism for multi-unit auctions with budgets. When the market is even mildly competitive, the approximation ratios of this mechanism are small constants for both the revenue and welfare objectives, and in fact for welfare the approximation converges to 1 as the market becomes fully competitive. We also give an impossibility theorem, showing that truthfulness requires discarding resources, and in particular, is incompatible with (Pareto) efficiency.

1998 ACM Subject Classification I.2.11 [Distributed Artificial Intelligence] Multiagent Systems, J.4 [Social and Behavioral Sciences] Economics, F.2 [Analysis of Algorithms and Problem Complexity]

Keywords and phrases mechanism design, multi-unit auctions, Walrasian pricing, market share

Digital Object Identifier 10.4230/LIPIcs.MFCS.2017.80

1 Introduction

Auctions are procedures for allocating goods that have been studied in economics in the 20th century, and which are even more relevant now due to the emergence of online platforms. Major companies such as Google and Facebook make most of their revenue through auctions, while an increasing number of governments around the world use spectrum auctions to allocate licenses for electromagnetic spectrum to companies. These transactions involve hundreds or thousands of participants with complex preferences, reason for which auctions require more careful design and their study has resurfaced in the computer science literature.

* Simina Brânzei was supported by the ISF grant 1435/14 administered by the Israeli Academy of Sciences and Israel-USA Bi-national Science Foundation (BSF) grant 2014389 and the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 740282). Aris Filos-Ratsikas was supported by the ERC Advanced Grant 321171 (ALGAME). A part of this work was done when Simina Brânzei was visiting the Simons Institute for the Theory of Computing.

In this paper we study a paradigmatic model known as multi-unit auctions with budgets, in which a seller brings multiple units of a good (e.g. apples), while the buyers bring money and have interests in consuming the goods. Multi-unit auctions have been studied in a large body of literature due to the importance of the model, which already illustrates complex phenomena [16, 6, 18, 17, 19].

The main requirements from a good auction mechanism are usually computational efficiency, revenue maximization for the seller, and simplicity of use for the participants, the latter of which is captured through the notion of truthfulness. An important property that is often missing from auction design is fairness, and in fact for the purpose of maximizing revenue it is useful to impose higher payments to the buyers that are more interested in the goods. However, there are studies showing that customers are unhappy with such discriminatory prices (see, e.g., [1]), which has led to a body of literature focused on achieving fair pricing [25, 22, 14, 23, 38].

A remarkable solution concept that has been used for achieving fairness in auctions comes from free markets, which are economic systems where the prices and allocations are not designed by a central authority. Instead, the prices emerge through a process of adjusting demand and supply such that everyone faces the same prices and the buyers freely purchase the bundles they are most interested in. When the goods are divisible, an outcome where supply and demand are perfectly balanced – known as competitive (or Walrasian) equilibrium [39] – always exists under mild assumptions on the utilities and has the property that the participants face the same prices and can freely acquire their favorite bundle at those prices. The competitive equilibrium models outcomes of large economies, where the goods are divisible and the participants so small (infinitesimal) that they have no influence on the market beyond purchasing their most preferred bundle at the current prices. Unfortunately, when the goods are indivisible, the competitive equilibrium does not necessarily exist (except for small classes of valuations see, e.g., [30, 24]) and the induced mechanism – the Walrasian mechanism [3, 13] – is generally manipulable.

A solution for recovering the attractive properties of the Walrasian equilibrium in the multi-unit model is to relax the clearing requirement of the market equilibrium, by allowing the seller to not sell all of the units. This solution is known as (Walrasian) envy-free pricing [25], and it ensures that all the participants of the market face the same prices¹, and each one purchases their favorite bundle of goods. An envy-free pricing trivially exists by pricing the goods infinitely high, so the challenge is finding one with good guarantees, such as high revenue for the seller or high welfare for the participants.

We would like to obtain envy-free pricing mechanisms that work well with strategic participants, who may alter their inputs to the mechanism to get better outcomes. To this end, we design an optimal truthful and envy-free mechanism for multi-unit auctions with budgets, with high revenue and welfare in competitive environments. Our work can be viewed as part of a general research agenda of *simplicity in mechanism design* [27], which recently proposed item pricing [4, 23] as a way of designing simpler auctions while at the same time avoiding the ill effects of discriminatory pricing [22, 1]. Item pricing is used in practice all over the world to sell goods in supermarkets or online platforms such as Amazon, which provides a strong motivation for understanding it theoretically. Other recent notions of simplicity in mechanism design include the menu-size complexity [26], the competition complexity [20], and verifiability of mechanisms (e.g. that the participants can easily convince themselves that the mechanism has a property, such as being truthful [9, 31]).

¹ The term envy-free pricing has also been used when the pricing is per-bundle, not per-item. We adopt the original definition of [25] which applies to unit-pricing, due to its attractive fairness properties [22].

1.1 Our Results

Our model is a multi-unit auction with budgets, in which a seller owns m identical units of an item. Each buyer i has a budget B_i and a value v_i per unit. The utilities of the buyers are quasi-linear up to the budget cap, while any allocation that exceeds that cap is unfeasible.

We deal with the problem of designing envy-free pricing schemes for the strongest concept of incentive compatibility, namely dominant strategy truthfulness. The truthful mechanisms are in the *prior-free* setting, i.e. they do not require any prior distribution assumptions. We evaluate the efficiency of mechanisms using the notion of *market share*, s^* , which captures the maximum buying power of any individual buyer in the market. A market share of at most 50% roughly means that no buyer can purchase more than half of the resources when competition is maximal, i.e. at the minimum envy-free price. Our main theorem can be summarized as follows.

► **Theorem 1** (Main Theorem (informal)). *For linear multi-unit auctions with known monetary endowments:*

- *There exists no (Walrasian) envy-free mechanism that is both truthful and non-wasteful.*
- *There exists a truthful (Walrasian) envy-free auction, which attains a fraction of at least $\max\left\{2, \frac{1}{1-s^*}\right\}$ of the optimal revenue and at least $1 - s^*$ of the optimal welfare on any market, where $0 < s^* < 1$ is the market share. This mechanism is optimal for both the revenue and welfare objectives when the market is even mildly competitive (i.e. with market share $s^* \leq 50\%$), and its approximation for welfare converges to 1 as the market becomes fully competitive.*

In the statement above, optimal means that there is no other truthful envy-free auction mechanism with a better approximation ratio. A mechanism is non-wasteful if it allocates as many units as possible at a given price. The impossibility theorem implies in particular that truthfulness is incompatible with Pareto efficiency. Our positive results are for *known* budgets, similarly to [16]. In the economics literature budgets are viewed as hard information (quantitative), as opposed to the valuations, which represent soft information and are more difficult to verify (see, e.g., [37]).

1.2 Related Work

The multi-unit setting has been studied in a large body of literature on auctions ([16, 6, 18, 17, 19]), where the focus has been on designing truthful auctions with good approximations to some desired objective, such as the social welfare or the revenue. Quite relevant to ours is the paper by [16], in which the authors study multi-unit auctions with budgets, however with no restriction to envy-free pricing or even item-pricing. They design a truthful auction (that uses discriminatory pricing) for *known budgets*, that achieves near-optimal revenue guarantees when the influence of each buyer in the auction is bounded, using a notion of buyer *dominance*, which is conceptually close to the market share notion that we employ. Their mechanism is based on the concept of clinching auctions [2].

Attempts at good prior-free truthful mechanisms for multi-unit auctions are seemingly impaired by their general impossibility result which states that truthfulness and efficiency are essentially incompatible when the budgets are *private*. Our general impossibility result is very similar in nature, but it is not implied by the results in [16] for the following two reasons: (a) our impossibility holds for *known* budgets and (b) our notion of efficiency is weaker, as it is naturally defined with respect to envy-free allocations only. This also means that our impossibility theorem is not implied by their uniqueness result, even for two buyers.

Multi-unit auctions with budgets have also been considered in [17] and [6], and without budgets ([19, 5, 18]); all of the aforementioned papers do not consider the envy-freeness constraint.

The effects of strategizing in markets have been studied extensively over the past few years ([7, 8, 12, 33, 34]). For more general envy-free auctions, besides the multi-unit case, there has been some work on truthful mechanisms in the literature of envy-free auctions ([25]) and ([28]) for *pair envy-freeness*, a different notion which dictates that no buyer would want to swap its allocation with that of any other buyer [32]. It is worth noticing that there is a body of literature that considers envy-free pricing as a purely optimization problem (with no regard to incentives) and provides approximation algorithms and hardness results for maximizing revenue and welfare in different auction settings [22, 15].

It is worth mentioning that the good approximations achieved by our truthful mechanism are a *prior-free setting* ([29]), i.e. we don't require any assumptions on prior distributions from which the input valuations are drawn. Good prior-free approximations are usually much harder to achieve and a large part of the literature is concerned with auctions under distributional assumptions, under the umbrella of *Bayesian mechanism design* ([10, 11, 29, 35]).

2 Preliminaries

In a linear multi-unit auction with budgets there is a set of buyers, denoted by $N = \{1, \dots, n\}$, and a single seller with m indivisible units of a good for sale. Each buyer i has a valuation $v_i > 0$ and a budget $B_i > 0$, both drawn from a discrete domain \mathbb{V} of rational numbers: $v_i, B_i \in \mathbb{V}$. The valuation v_i indicates the value of the buyer for one unit of the good.

An *allocation* is an assignment of units to the buyers denoted by a vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{Z}_+^n$, where x_i is the number of units received by buyer i . We are interested in feasible allocations, for which: $\sum_{i=1}^n x_i \leq m$.

The seller will set a price p per unit, such that the price of purchasing ℓ units is $p \cdot \ell$ for any buyer. The interests of the buyers at a given price are captured by the demand function.

► **Definition 2 (Demand).** The *demand* of buyer i at a price p is a set consisting of all the possible bundle sizes (number of units) that the buyer would like to purchase at this price:

$$D_i(p) = \begin{cases} \min\{\lfloor \frac{B_i}{p} \rfloor, m\}, & \text{if } p < v_i \\ 0, \dots, \min\{\lfloor \frac{B_i}{p} \rfloor, m\}, & \text{if } p = v_i \\ 0, & \text{otherwise.} \end{cases}$$

If a buyer is indifferent between buying and not buying at a price, then its demand is a set of all the possible bundles that it can afford, based on its budget constraint.

► **Definition 3 (Utility).** The *utility* of buyer i given a price p and an allocation \mathbf{x} is

$$u_i(p, x_i) = \begin{cases} v_i \cdot x_i - p \cdot x_i, & \text{if } p \cdot x_i \leq B_i \\ -\infty, & \text{otherwise} \end{cases}$$

(Walrasian) Envy-free Pricing. An allocation and price (\mathbf{x}, p) represent a (Walrasian) *envy-free pricing* if each buyer is allocated a number of units in its demand set at price p , i.e. $x_i \in D_i(p)$ for all $i \in N$. A price p is an *envy-free price* if there exists an allocation \mathbf{x} such that (\mathbf{x}, p) is an envy-free pricing.

While an envy-free pricing always exists (just set $p = \infty$), it is not always possible to sell *all* the units in an envy-free way. We illustrate this through an example.

► **Example 4** (Non-existence of envy-free clearing prices). Let $N = \{1, 2\}$, $m = 3$, valuations $v_1 = v_2 = 1.1$, and $B_1 = B_2 = 1$. At any price $p > 0.5$, no more than 2 units can be sold in total because of budget constraints. At $p \leq 0.5$, both buyers are interested and demand at least 2 units each, but there are only 3 units in total.

Objectives. We are interested in maximizing the *social welfare* and *revenue* objectives attained at envy-free pricing. The *social welfare* at an envy-free pricing (\mathbf{x}, p) is the total value of the buyers for the goods allocated, while the *revenue* is the total payment received by the seller, i.e. $\mathcal{SW}(\mathbf{x}, p) = \sum_{i=1}^n v_i \cdot x_i$ and $\mathcal{REV}(\mathbf{x}, p) = \sum_{i=1}^n x_i \cdot p$.

Mechanisms. The goal of the seller will be to obtain money in exchange for the goods, however, it can only do that if the buyers are interested in purchasing them. The problem of the seller will be to obtain accurate information about the preferences of the buyers that would allow optimizing the pricing. Since the inputs (valuations) of the buyers are private, we will aim to design auction mechanisms that incentivize the buyers to reveal their true preferences [36].

An auction *mechanism* is a function $M : \mathbb{V}^n \rightarrow \mathbb{O} \times \mathbb{Z}_+^n$ that maps the valuations reported by the buyers to a price $p \in \mathbb{O}$, where \mathbb{O} is the space from which the prices are drawn², and an allocation vector $\mathbf{x} \in \mathbb{Z}_+^n$.

► **Definition 5** (Truthful Mechanism). A mechanism M is *truthful* if it incentivizes the buyers to reveal their true inputs, i.e. $u_i(M(\mathbf{v})) \geq u_i(M(v'_i, v_{-i}))$, for all $i \in N$, any alternative report $v'_i \in \mathbb{V}$ of buyer i and any vector of reports v_{-i} of all the other buyers.

Requiring incentive compatibility from a mechanism can lead to worse revenue, so our goal will be to design mechanisms that achieve revenue close to that attained in the pure optimization problem (of finding a revenue optimal envy-free pricing without incentive constraints).

Types of Buyers. The next definitions will be used extensively in the paper. Buyer i is said to be *hungry* at price p if $v_i > p$ and *semi-hungry* if $v_i = p$. Given an allocation \mathbf{x} and a price p buyer i is *essentially hungry* if it is either semi-hungry with $x_i = \min\{\lfloor B_i/p \rfloor, m\}$ or hungry. In other words, a buyer is essentially hungry if its value per unit is at least as high as the price per unit and, moreover, the buyer receives the largest non-zero element in its demand set.

3 An optimal envy-free and truthful mechanism

In this section, we present our main contribution, an envy-free and truthful mechanism, which is optimal among all truthful mechanisms and achieves small constant approximations to the optimal welfare and revenue. The approximation guarantees are with respect to the *market-share* s^* , which intuitively captures the maximum purchasing power of any individual buyer in the auction. The formal definition is postponed to the corresponding subsection.

² In principle the spaces \mathbb{V} and \mathbb{O} can be the same but for the purpose of getting good revenue and welfare, it is useful to have the price to be drawn from a slightly larger domain; see Section 3.

► **Theorem 6.** *There exists a truthful (Walrasian) envy-free auction, which attains a fraction of at least*

- $\max\left\{2, \frac{1}{1-s^*}\right\}$ of the optimal revenue, and
- $1 - s^*$ of the optimal welfare

on any market. This mechanism is optimal for both the revenue and welfare objectives when the market is even mildly competitive (i.e. with market share $s^* \leq 50\%$), and its approximation for welfare converges to 1 as the auction becomes fully competitive.

Consider the following mechanism.

All-or-Nothing:

Given as input the valuations of the buyers, let p be the minimum envy-free price and \mathbf{x} the allocation obtained as follows:

- For every hungry buyer i , set x_i to its demand.
- For every buyer i with $v_i < p$, set $x_i = 0$.
- For every semi-hungry buyer i , set $x_i = \lfloor B_i/p \rfloor$ if possible, otherwise set $x_i = 0$ taking the semi-hungry buyers in lexicographic order.

In other words, the mechanism always outputs the minimum envy-free price but if there are semi-hungry buyers at that price, they get either all the units they can afford at this price or 0, even if there are still available units, after satisfying the demands of the hungry buyers.

► **Lemma 7.** *The minimum envy-free price does not exist when the price domain is \mathbb{R} .*

Proof. If the price can be any real number, consider an auction with $n = 2$ buyers, $m = 2$ units, valuations $v_1 = v_2 = 3$ and budgets $B_1 = B_2 = 2$. At any price $p \leq 1$, there is overdemand since each buyer is hungry and demands at least 2 units, while there are only 2 units in total. At any price $p \in (1, 2]$, each buyer demands at most one unit due to budget constraints, and so all the prices in the range $(1, 2]$ are envy-free. This is an open set, and so there is no minimum envy-free price. Note however, that by making the output domain discrete, e.g. with 0.1 increments starting from zero, then the minimum envy-free price output is 1.01. At this price each buyer purchases 1 unit. ◀

Given the example above, we will consider the discrete domain \mathbb{V} as an infinite grid with entries of the form $k \cdot \epsilon$, for $k \in \mathbb{N}$ and some sufficiently small³ ϵ . For the output of the mechanism, we will assume a slightly finer grid, e.g. with entries $k \cdot \delta = k(\epsilon/2)$, for $k \in \mathbb{N}$. The minimum envy-free price can be found in time which is polynomial in the input and $\log(1/\epsilon)$, using binary search⁴ and the mechanism is optimal with respect to discrete domain that we operate on. Operating on a grid is actually without loss of generality in terms of the objectives; even if we compare to the optimal on the continuous domain, if our discretization is fine enough, we don't lose any revenue or welfare. This is established by the following theorem; the proof is omitted due to lack of space (see full version).

³ For most of our results, any discrete domain is sufficient for the results to hold; for some results we will need to a number of grid points that polynomial in the size of the input grid.

⁴ In the full version, we describe a faster procedure that finds the minimum envy-free without requiring to do binary search over the grid.

► **Theorem 8.** *When the valuation and budget of each buyer are drawn from a discrete grid with entries $k \cdot \epsilon$, and the price is drawn from a finer grid with entries $k \cdot \epsilon/2$, for $k \in \mathbb{N}$, then the welfare and revenue loss of the ALL-OR-NOTHING mechanism due to the discretization of the output domain is zero. The mechanism always runs in time polynomial in the input and $\log(1/\epsilon)$.*

Truthfulness of the All-or-Nothing Mechanism

The following theorem establishes the truthfulness of ALL-OR-NOTHING.

► **Theorem 9.** *The ALL-OR-NOTHING mechanism is truthful.*

Proof. First, we will prove the following statement. If p is any envy-free price and p' is an envy-free price such that $p \leq p'$ then the utility of any essentially hungry buyer i at price p is at least as large as its utility at price p' . The case when $p' = p$ is trivial, since the price (and the allocation) do not change. Consider the case when $p < p'$. Since p is an envy-free price, buyer i receives the maximum number of items in its demand. For a higher price p' , its demand will be at most as large as its demand at price p and hence its utility at p' will be at most as large as its utility at p .

Assume now for contradiction that Mechanism ALL-OR-NOTHING is not truthful and let i be a deviating buyer who benefits by misreporting its valuation v_i as v'_i at some valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, for which the minimum envy-free price is p . Let p' be the new minimum envy free price and let \mathbf{x} and \mathbf{x}' be the corresponding allocations at p and p' respectively, according to ALL-OR-NOTHING. Let $\mathbf{v}' = (v'_i, v_{-i})$ be the valuation profile after the deviation.

We start by arguing that the deviating buyer i is essentially hungry. First, assume for contradiction that i is neither hungry nor semi-hungry, which means that $v_i < p$. Clearly, if $p' \geq p$, then buyer i does not receive any units at p' and there is no incentive for manipulation; thus we must have that $p' < p$. This implies that every buyer j such that $x_j > 0$ at price p is hungry at price p' and hence $x'_j \geq x_j$. Since the demand of all players does not decrease at p' , this implies that p' is also an envy-free price on instance \mathbf{v} , contradicting minimality of p .

Next, assume that buyer i is semi-hungry but not essentially hungry, which means that $v_i = p$ and $x_i = 0$, by the allocation of the mechanism. Again, in order for the buyer to benefit, it has to hold that $p' < p$ and $x'_i > 0$ which implies that $x'_i = \lfloor B_i/p' \rfloor$, i.e. buyer i receives the largest element in its demand set at price p' . But then, since $p' < p$ and p' is an envy-free price, buyer i could receive $\lfloor B_i/p \rfloor$ units at price p without violating the envy-freeness of p , in contradiction with each buyer i being essentially hungry at p .

From the previous two paragraphs, the deviating buyer must be essentially hungry. This means that $x_i > 0$ and $v_i \geq p$. By the discussion in the first paragraph of the proof, we have $p' < p$. Since $x_i > 0$, the buyer does not benefit from reporting v'_i such that $v'_i < p'$. Thus it suffices to consider the case when $v'_i \geq p'$. We have two subcases:

- $v'_i > p$: Buyer i is essentially hungry at price p according to v_i and hungry at price p' according to v'_i . The reports of the other buyers are fixed and B_i is known; similarly to above, price p' is an envy-free price on instance \mathbf{v} , contradicting the minimality of p .
- $v'_i = p'$: Intuitively, an essentially hungry buyer at price p is misreporting its valuation as being lower trying to achieve an envy-free price p' equal to the reported valuation. Since $v'_i = p'$, Mechanism ALL-OR-NOTHING gives the buyer either as many units as it can afford at this price or zero units. In the first case, since p' is envy-free and B_i is known,

buyer i at price p' receives the largest element in its demand set and since the valuations of all other buyers are fixed, p' is also an envy-free price on input \mathbf{v} , contradicting the minimality of p . In the second case, the buyer does not receive any units and hence it does not benefit from misreporting.

Thus there are no improving deviations, which concludes the proof of the theorem. ◀

Performance of the All-or-Nothing Mechanism

Next, we show that the mechanism has a good performance for both objectives. We measure the performance of a truthful mechanism by the standard notion of approximation ratio, i.e.

$$\text{ratio}(M) = \sup_{\mathbf{v} \in \mathbb{R}^n} \frac{\max_{\mathbf{x}, p} \mathcal{OBJ}(\mathbf{v})}{\mathcal{OBJ}(M(\mathbf{v}))},$$

where $\mathcal{OBJ} \in \{SW, \mathcal{REV}\}$ is either the social welfare or the revenue objective. Obviously, a mechanism that outputs a pair that maximizes the objectives has approximation ratio 1. The goal is to construct truthful mechanisms with approximation ratio as close to 1 as possible.

We remark here that for the approximation ratios, we only need to consider valuation profiles that are not “trivial”, i.e. input profiles for which at any envy-free price, no hungry or semi-hungry buyers can afford a single unit and hence the envy-free price can be anything; on trivial profiles, both the optimal price and allocation and the price and allocation output by Mechanism ALL-OR-NOTHING obtain zero social welfare or zero revenue.

Market Share. A well-known notion for measuring the competitiveness of a market is the *market share*, understood as the percentage of the market accounted for by a specific entity (see, e.g., [21], Chapter 2).

In our model, the maximum purchasing power (i.e. number of units) of any buyer in the auction occurs at the minimum envy-free price, p_{min} . By the definition of the demand, there are many ways of allocating the semi-hungry buyers, so when measuring the purchasing power of an individual buyer we consider the maximum number of units that buyer can receive, taken over the set of all feasible maximal allocations at p_{min} . Let this set be \mathcal{X} . Then the market share of buyer i can be defined as:

$$s_i = \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{x_i}{\sum_{k=1}^n x_k} \right).$$

Then, the *market share* is defined as $s^* = \max_{i=1}^n s_i$. Roughly speaking, a market share $s^* \leq 1/2$ means that a buyer can never purchase more than half of the resources.

► **Theorem 10.** *The ALL-OR-NOTHING mechanism approximates the optimal revenue within a factor of 2 whenever the market share, s^* , is at most 50%.*

Proof. Let OPT be the optimal revenue, attained at some price p^* and allocation \mathbf{x} , and $\mathcal{REV}(AON)$ the revenue attained by the ALL-OR-NOTHING mechanism. By definition, mechanism ALL-OR-NOTHING outputs the minimum envy-free price p_{min} , together with an allocation \mathbf{z} . For ease of exposition, let $\alpha_i = B_i/p_{min}$ and $\alpha_i^* = B_i/p^*$, $\forall i \in N$. There are two cases, depending on whether the optimal envy-free price, p^* , is equal to the minimum envy-free price, p_{min} :

Case 1: $p^* > p_{min}$. Denote by L the set of buyers with valuations at least p^* that can afford at least one unit at the optimal price. Note that the set of buyers that get allocated

at p_{min} is a superset of L . Moreover, the optimal revenue is bounded by the revenue attained at the (possibly infeasible) allocation where all the buyers in L get the maximum number of units in their demand. These observations give the next inequalities:

$$\mathcal{REV}(AON) \geq \sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_{min} \quad \text{and} \quad OPT \leq \sum_{i \in L} \lfloor \alpha_i^* \rfloor \cdot p^*.$$

Then the revenue is bounded by:

$$\begin{aligned} \frac{\mathcal{REV}(AON)}{OPT} &\geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_{min}}{\sum_{i \in L} \lfloor \alpha_i^* \rfloor \cdot p^*} \geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_{min}}{\sum_{i \in L} \alpha_i^* \cdot p^*} = \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor \cdot p_{min}}{\sum_{i \in L} B_i} \\ &= \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor}{\sum_{i \in L} \alpha_i} \geq \frac{\sum_{i \in L} \lfloor \alpha_i \rfloor}{\sum_{i \in L} 2 \lfloor \alpha_i \rfloor} = \frac{1}{2}, \end{aligned}$$

where we used that the auction is non-trivial, i.e. for any buyer $i \in L$, $\lfloor \alpha_i \rfloor \geq 1$, and so $\alpha_i \leq \lfloor \alpha_i \rfloor + 1 \leq 2 \lfloor \alpha_i \rfloor$.

Case 2: $p^* = p_{min}$. The hungry buyers at p_{min} , as well as the buyers with valuations below p_{min} , receive identical allocations under ALL-OR-NOTHING and the optimal allocation, \mathbf{x} . However there are multiple ways of assigning the semi-hungry buyers to achieve an optimal allocation. Recall that \mathbf{z} is the allocation made by ALL-OR-NOTHING. Without loss of generality, we can assume that \mathbf{x} is an optimal allocation with the property that \mathbf{x} is a superset of \mathbf{z} and the following condition holds:

- the number of buyers not allocated under \mathbf{z} , but that are allocated under \mathbf{x} , is minimized.

We argue that \mathbf{x} allocates at most one buyer more compared to \mathbf{z} . Assume by contradiction that there are at least two semi-hungry buyers i and j , such that $0 < x_i < \lfloor \alpha_i \rfloor$ and $0 < x_j < \lfloor \alpha_j \rfloor$. Then we can progressively take units from buyer j and transfer them to buyer i , until either buyer i receives $x'_i = \lfloor \alpha_i \rfloor$, or buyer j receives $x'_j = 0$. Hence we can assume that the set of semi-hungry buyers that receive non-zero, non-maximal allocations in the optimal solution \mathbf{x} is either empty or a singleton. If the set is empty, then ALL-OR-NOTHING is optimal. Otherwise, let the singleton be ℓ ; denote by \tilde{x}_ℓ the maximum number of units that ℓ can receive in any envy-free allocation at p_{min} . Since the number of units allocated by any maximal envy-free allocation at p_{min} is equal to $\sum_{i=1}^n x_i$, but $x_\ell \leq \tilde{x}_\ell$, we get:

$$\frac{x_\ell}{\sum_{i=1}^n x_i} \leq \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = s_i^*.$$

Thus

$$\begin{aligned} \frac{\mathcal{REV}(AON)}{OPT} &= \frac{OPT - x_\ell \cdot p_{min}}{OPT} \geq \frac{OPT - \tilde{x}_\ell \cdot p_{min}}{OPT} = 1 - \frac{\tilde{x}_\ell \cdot p_{min}}{\sum_{i=1}^n x_i \cdot p_{min}} \\ &= 1 - \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = 1 - s_i^* \geq 1 - s^* \end{aligned}$$

Combining the two cases, the bound follows. This completes the proof. \blacktriangleleft

► **Corollary 11.** *The performance of the ALL-OR-NOTHING mechanism is $\max\{2, 1/(1 - s^*)\}$ on any market (i.e. with market share $0 < s^* < 1$).*

Proof. From the proof of Theorem 10, since the arguments of Case 1 do not use the market share s^* , it follows that the ratio of ALL-OR-NOTHING for the revenue objective can alternatively be stated as $\max\{2, 1/(1 - s^*)\}$ and therefore it degrades gracefully with the increase in the market share. \blacktriangleleft

The next theorem establishes that the approximation ratio for welfare is also constant.

► **Theorem 12.** *The approximation ratio of Mechanism ALL-OR-NOTHING with respect to the social welfare is at most $1/(1-s^*)$, where the market share $s^* \in (0, 1)$. The approximation ratio goes to 1 as the market becomes fully competitive.*

Proof. For social welfare we have, similarly to Theorem 10, that

$$\begin{aligned} \frac{SW(\text{AON})}{OPT} &= \frac{OPT - x_\ell \cdot v_\ell}{OPT} \geq \frac{OPT - \tilde{x}_\ell \cdot v_\ell}{OPT} = 1 - \frac{\tilde{x}_\ell \cdot v_\ell}{\sum_{i=1}^n x_i \cdot v_i} \geq 1 - \frac{\tilde{x}_\ell \cdot v_\ell}{\sum_{i=1}^n x_i \cdot v_\ell} \\ &= 1 - \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = 1 - s_i^* \geq 1 - s^*, \end{aligned}$$

where OPT is now the optimal welfare, \mathbf{x} the corresponding allocation at OPT , and we used the fact that $v_\ell \leq v_i$ for all $i \in L$. ◀

Finally, ALL-OR-NOTHING is optimal among all truthful mechanisms for both objectives whenever the market share s^* is at most $1/2$.

► **Theorem 13.** *Let M be any truthful mechanism that always outputs an envy-free pricing scheme. Then the approximation ratio of M for the revenue and the welfare objective is at least $2 - \frac{4}{m+2}$.*

Proof. Consider an auction with equal budgets, B , and valuation profile \mathbf{v} . Assume that buyer 1 has the highest valuation, v_1 , buyer 2 the second highest valuation v_2 , with the property that $v_1 > v_2 + \epsilon$, where ϵ is set later. Let $v_i < v_2$ for all buyers $i = 3, 4, \dots, n$. Set B such that $\lfloor \frac{B}{v_2} \rfloor = \frac{m}{2} + 1$ and ϵ such that $\lfloor \frac{B}{v_2 + \epsilon} \rfloor = \frac{m}{2}$. Informally, the buyers can afford $\frac{m}{2} + 1$ units at prices v_2 and $v_2 + \epsilon$. Note that on this profile, Mechanism ALL-OR-NOTHING outputs price v_2 and allocates $\frac{m}{2} + 1$ units to buyer 1. For a concrete example of such an auction, take $m = 12$, $v_1 = 1.12$, $v_2 = 1.11$ (i.e. $\epsilon = 0.01$) and $B = 8$ (the example can be extended to any number of units with appropriate scaling of the parameters).

Let M be any truthful mechanism, p_M its price on this instance, and p^* the optimal price (with respect to the objective in question). The high level idea of the proof, for both objectives, is the following. We start from the profile \mathbf{v} above, where $p_{min} = v_2$ is the minimum envy-free price, and argue that if $p^* \neq v_2$, then the bound follows. Otherwise, $p^* = v_2$, case in which we construct a series of profiles $\mathbf{v}, \mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \dots, \mathbf{v}^{(k)}$ that only differ from the previous profile in the sequence by the reported valuation $v_2^{(j)}$ of buyer 2. We argue that in each such profile, either the mechanism allocates units to buyer 1 only, case in which the bound is immediate, or buyer 2 is semi-hungry. In the latter case, truthfulness and the constraints on the number of units will imply that any truthful mechanism must allocate to buyer 2 zero items, yielding again the required bound.

First, consider the social welfare objective. Observe that for the optimal price p^* on profile \mathbf{v} , it holds that $p^* = v_2$. We have a few subcases:

Case 1: $p_M < v_2$. Then M is not an envy-free mechanism, since in this case there would be over-demand for units.

Case 2: $p_M > v_2$: Then M allocates units only to buyer 1, achieving a social welfare of at most $(\frac{m}{2} + 1)v_2$. The maximum social welfare is $m \cdot v_2$, so the approximation ratio of M is at least $\frac{m}{(\frac{m}{2} + 1)} = 2 - \frac{4}{m+2}$.

Case 3: $p_M = v_2$: Let x_2 be the number of units allocated to buyer 2 at price v_2 ; note that since buyer 2 is semi-hungry at v_2 , any number of units up to $\frac{m}{2} - 1$ is a valid allocation. If $x_2 = 0$, then M allocates units only to buyer 1 at price v_2 and for the same reason as in Case 2, the ratio is greater than or equal to $2 - \frac{4}{m+2}$; so we can assume $x_2 \geq 1$.

Next, consider valuation profile $\mathbf{v}^{(1)}$ where for each buyer $i \neq 2$, we have $v_i^{(1)} = v_i$, while for buyer 2, $v_2 < v_2^{(1)} < v_2 + \epsilon$. By definition of B , the minimum envy-free price on $\mathbf{v}^{(1)}$ is $v_2^{(1)}$. Let $p_M^{(1)}$ be the price output by M on valuation profile $\mathbf{v}^{(1)}$ and take a few subcases:

- (a) $p_M^{(1)} > v_2^{(1)}$: Then using the same argument as in Case 2, the approximation is at least $2 - \frac{4}{m+2}$.
- (b) $p_M^{(1)} < v_2^{(1)}$: This cannot happen because by definition of the budgets, $v_2^{(1)}$ is the minimum envy-free price.
- (c) $p_M^{(1)} = v_2^{(1)}$: Let $x_2^{(1)}$ be the number of units allocated to buyer 2 at profile $\mathbf{v}^{(1)}$; we claim that $x_2^{(1)} \geq 2$. Otherwise, if $x_2^{(1)} \leq 1$, then on profile $\mathbf{v}^{(1)}$ buyer 2 would have an incentive to report v_2 , which would move the price to v_2 , giving buyer 2 at least as many units (at a lower price), contradicting truthfulness.

Consider now a valuation profile $\mathbf{v}^{(2)}$, where for each buyer $i \neq 2$, it holds that $v_i^{(2)} = v_i^{(1)} = v_i$ and for buyer 2 it holds that $v_2^{(1)} < v_2^{(2)} < v_2 + \epsilon$. For the same reasons as in Cases a-c, the behavior of M must be such that:

- the price output on input $\mathbf{v}^{(2)}$ is $v_2^{(2)}$ (otherwise M only allocates to buyer 1, and the bound is immediate), and
- the number of units $x_2^{(2)}$ allocated to buyer 2 is at least 3 (otherwise truthfulness would be violated).

By iterating through all the profiles in the sequence constructed in this manner, we arrive at a valuation profile $\mathbf{v}^{(k)}$ (similarly constructed), where the price is $v_2^{(k)}$ and buyer 2 receives at least $m/2$ units. However, buyer 1 is still hungry at price $v_2^{(k)}$ and should receive at least $\frac{m}{2} + 1$ units, which violates the unit supply constraint. This implies that in the first profile, \mathbf{v} , M must allocate 0 units to buyer 2 (by setting the price to v_2 or to something higher where buyer 2 does not want any units). This implies that the approximation ratio is at least $2 - \frac{4}{m+2}$.

For the revenue objective, the argument is exactly the same, but we need to establish that at any profile \mathbf{v} or $\mathbf{v}^{(i)}$, $i = 1, \dots, k$ that we construct, the optimal envy-free price is equal to the second highest reported valuation, i.e. v_2 or $v_2^{(i)}$, $i = 1, \dots, k$ respectively. To do that, choose v_1 such that $v_1 = v_2 + \delta$, where $\delta > \epsilon$, but small enough such that $\lfloor \frac{B}{v_2 + \delta} \rfloor = \lfloor \frac{B}{v_2} \rfloor$, i.e. any hungry buyer at price $v_2 + \delta$ buys the same number of units as it would buy at price v_2 . Furthermore, ϵ and δ can be chosen small enough such that $(\frac{m}{2} + 1)(v_2 + \delta) < m \cdot v_2$, i.e. the revenue obtained by selling $\frac{m}{2} + 1$ units to buyer 1 at price $v_2 + \delta$ is smaller than the revenue obtained by selling $\frac{m}{2} + 1$ units to buyer 1 and $\frac{m}{2} - \epsilon$ units to buyer 2 at price v_2 . This establishes the optimal envy-free price is the same as before, for every profile in the sequence and all arguments go through.

Given that we are working over a discrete domain, for the proof to go through, it suffices to assume that there are m points of the domain between v_1 and v_2 , which is easily the case if the domain is not too sparse. Specifically, for the concrete example presented at the first paragraph of the proof, assuming that the domain contains all the decimal floating point numbers with up to two decimal places suffices. ◀

4 Impossibility Results

In this section, we state our impossibility results, which imply that truthfulness can only be guaranteed when there is some kind of wastefulness; a similar observation was made in [6] for a different setting.

► **Theorem 14.** *There is no Pareto efficient, truthful mechanism that always outputs an envy-free pricing, even when the budgets are known.*

The proof of the theorem is left for the full version. The next theorem provides a stronger impossibility result. First, we provide the necessary definitions. A buyer i on profile input v is called *irrelevant* if at the minimum envy-free price p on v , the buyer can not buy even a single unit. A mechanism is called *in-range* if it always outputs an envy-free price in the interval $[0, v_j]$ where v_j is the highest valuation among all buyers that are not irrelevant. Finally, a mechanism is *non-wasteful* if at a given price p , the mechanism allocates as many items as possible to the buyers. Note that Pareto efficiency implies in-range and non-wastefulness, but not the other way around. In a sense, while Pareto efficiency also determines the price chosen by the mechanism, non-wastefulness only concerns the allocation given a price, whereas in-range only restricts prices to a “reasonable” interval.

► **Theorem 15.** *There is no in-range, non-wasteful and truthful mechanism that always outputs an envy-free pricing scheme, even when the budgets are known.*

We leave the proof for the full version. To prove the impossibility, we first obtain a necessary condition; any mechanism in this class must essentially output the minimum envy-free price (or the next highest price on the output grid). Then we can use this result to construct an example where the mechanism must leave some items unallocated in order to satisfy truthfulness.

5 Discussion

Our results show that it is possible to achieve good approximate truthful mechanisms, under reasonable assumptions on the competitiveness of the auctions which retain some of the attractive properties of the Walrasian equilibrium solutions. The same agenda could be applied to more general auctions, beyond the case of linear valuations or even beyond multi-unit auctions. It would be interesting to obtain a complete characterization of truthfulness in the case of private or known budgets; for the case of private budgets, we can show that a class of order statistic mechanisms are truthful, but the welfare or revenue guarantees for this case may be poor. Finally, in the full version, we present an interesting special case, that of *monotone auctions*, in which Mechanism ALL-OR-NOTHING is optimal among all truthful mechanisms for both objectives, regardless of the market share.

Acknowledgements. We would like to thank the MFCS reviewers for useful feedback.

References

- 1 E. Anderson and D. Simester. Price stickiness and customer antagonism. *Available at SSRN 1273647*, 2008.
- 2 L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *The American Economic Review*, 94(5):1452–1475, 2004.
- 3 M. Babaioff, B. Lucier, N. Nisan, and R. Paes Leme. On the efficiency of the walrasian mechanism. In *ACM EC*, pages 783–800. ACM, 2014.
- 4 M. F. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *ACM EC*, pages 50–59. ACM, 2008.
- 5 Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *TARK*, pages 72–87, 2003.

- 6 C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *ACM EC*, pages 44–51, 2005.
- 7 A. Borodin, O. Lev, and T. Strangway. Budgetary effects on pricing equilibrium in online markets. In *AAMAS*, 2016.
- 8 S. Brânzei, Y. Chen, X. Deng, A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. The fisher market game: Equilibrium and welfare. In *AAAI*, pages 587–593, 2014.
- 9 Simina Brânzei and Ariel D. Procaccia. Verifiably truthful mechanisms. In *ITCS*, pages 297–306, 2015.
- 10 Y. Cai, C. Daskalakis, and M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *FOCS*, pages 130–139, 2012.
- 11 Y. Cai, C. Daskalakis, and M. Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *SODA*, pages 578–595, 2013.
- 12 N. Chen, X. Deng, and J. Zhang. How profitable are strategic behaviors in a market? In *ESA*, pages 106–118. Springer, 2011.
- 13 M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *FOCS*, pages 35–44, 2008.
- 14 V. Cohen-Addad, A. Eden, M. Feldman, and A. Fiat. The invisible hand of dynamic market pricing. In *ACM EC*, pages 383–400, 2016.
- 15 R. Colini-Baldeschi, S. Leonardi, P. Sankowski, and Q. Zhang. Revenue maximizing envy-free fixed-price auctions with budgets. In *WINE*, pages 233–246. Springer, 2014.
- 16 S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *GEB*, 74(2):486–503, 2012.
- 17 S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *ICALP*, pages 392–404. Springer, 2014.
- 18 S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *ACM EC*, pages 346–351, 2007.
- 19 S. Dobzinski and N. Nisan. Multi-unit auctions: beyond roberts. *JET*, 156:14–44, 2015.
- 20 A. Eden, M. Feldman, O. Friedler, I. Talgam-Cohen, and M. Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In *ACM EC*, 2017.
- 21 P. Farris, N. Bendle, P. Pfeifer, and D. Reibstein. *Marketing metrics: The definitive guide to measuring marketing performance*. Pearson Education, 2010.
- 22 M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *ACM EC*, pages 532–549, 2012.
- 23 M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *SODA*, pages 123–135, 2015.
- 24 F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.
- 25 V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *SODA*, pages 1164–1173, 2005.
- 26 Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *ACM EC*, pages 565–566, 2013.
- 27 J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *ACM EC*, pages 225–234. ACM, 2009.
- 28 J. Hartline and Q. Yan. Envy, truth, and profit. In *ACM EC*, pages 243–252, 2011.
- 29 J. D. Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.
- 30 A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.
- 31 Shengwu Li. Obviously strategy-proof mechanisms, 2015. Working paper.

- 32 E. Markakis and O. Telelis. Envy-free revenue approximation for asymmetric buyers with budgets. In *SAGT*, pages 247–259, 2016.
- 33 R. Mehta and M. Sohoni. Exchange markets: Strategy meets supply-awareness. In *WINE*, pages 361–362. Springer, 2013.
- 34 R. Mehta, N. Thain, L. Végh, and A. Vetta. To save or not to save: The fisher game. In *WINE*, pages 294–307. Springer, 2014.
- 35 R. B. Myerson. Optimal auction design. *MOR*, 6(1):58–73, 1981.
- 36 N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge Univ. Press, (editors) 2007.
- 37 M. Petersen. Information: Hard and soft, 7 2004.
- 38 Shreyas Sekar. Posted pricing sans discrimination. In *IJCAI*, to appear, 2017.
- 39 L. Walras. Elements d’économie politique pure, ou theorie de la richesse sociale (in french), 1874. English translation: Elements of pure economics; or, the theory of social wealth. American Economic Association and the Royal Economic Society, 1954.