Walrasian Pricing in Multi-unit Auctions

Citation for published version:
Brânzei, S, Filos-Ratsikas, A, Miltersen, PB & Zeng, Y 2023, 'Walrasian Pricing in Multi-unit Auctions', *Artificial Intelligence*, vol. 322, 103961, pp. 1-17. https://doi.org/10.1016/j.artint.2023.103961

Digital Object Identifier (DOI):
10.1016/j.artint.2023.103961

Published In:
Artificial Intelligence

Highlights

Walrasian Pricing in Multi-unit Auctions

Simina Brânzei, Aris Filos-Ratsikas, Peter Bro Miltersen, Yulong Zeng

- We consider multi-unit auctions with budgets and design a best possible envy-free and prior-free mechanism.
- The mechanism obtains revenue and welfare that are close to optimal, within small constant factors; for welfare, the quality of approximation converges to 1 as the market becomes fully competitive.

Walrasian Pricing in Multi-unit Auctions

Simina Brânzei[a], Aris Filos-Ratsikas[b], Peter Bro Miltersen[c], Yulong Zeng[d]

[a] *Department of Computer Science, Purdue University, USA*
[b] *School of Informatics, University of Edinburgh, UK*
[c] *Department of Computer Science, Aarhus University, Denmark*
[d] *ASResearch, China*

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.

We design a best possible envy-free and prior-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 is thus incompatible with (Pareto) efficiency.

Keywords: multi-unit auctions, mechanism design, revenue, social welfare, approximation

1. Introduction

In a multi-unit auction, a seller brings multiple units of a good (e.g. chairs) and several buyers with budgets are interested in acquiring 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 [1, 2, 3, 4, 5]. Central 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. Fairness is an important property often missing from auction design, and in fact for the purpose of maximizing revenue it is useful to impose higher payments to very interested buyers. However, there are studies showing that customers are unhappy with such discriminatory prices (see, e.g., [6]), which led to a body of literature focused on achieving fair pricing [7, 8, 9, 10, 11].

The competitive (aka market or Walrasian) [12] equilibrium models the allocation of resources in an exchange economy at the steady state, where supply equals demand. When the goods are divisible, the Walrasian equilibrium exists under mild assumptions on the utilities and is considered very fair, as the prices are the same for everyone and each agent can freely acquire their favorite bundle at those prices. Unfortunately, when the goods are indivisible, the competitive equilibrium does not necessarily exist (except for small classes of valuations see, e.g., [13, 14]) and the induced mechanism – the Walrasian mechanism [15, 16] – is generally manipulable.

A solution for recovering the attractive fairness properties of the Walrasian equilibrium in the multi-unit auction 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 [7], and it ensures that all the participants of the market face the same prices[1], 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 a best possible truthful and envy-free mechanism for multi-unit auctions with budgets, which achieves high revenue and welfare in competitive environments.

Our work can be viewed as part of a general research agenda of *simplicity in mechanism design* [17], which proposed item pricing [18, 10] as a way of designing simpler auctions while at the same time avoiding the ill effects of discriminatory pricing [8, 6]. Item pricing is used in practice all over the world to sell goods in supermarkets or online platforms, and thus there is a strong motivation to understand it theoretically.

[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 [7] which applies to unit-pricing, due to its attractive fairness properties [8].

1.1. Our Results

We briefly describe the setting and our main result next.

Multi-unit auction with budgets. Consider the setting where a seller tries to sell m identical units of a good to n buyers. Each buyer i has a value v_i per unit and a budget B_i. We refer to this auction as a *linear multi-unit auction*, because the value of an agent for x units of the good is $v_i \cdot x$.

Walrasian (envy-free) pricing. The seller sets a price p for each unit. The utility of buyer is quasilinear up to the budget constraint, i.e., the utility of buyer i for purchasing x units at price p is $u_i = v_i \cdot x - p \cdot x$, if $p \cdot x \leq B_i$, and $u_i = -\infty$ otherwise. Given price p per unit, each buyer purchases a bundle of goods that maximizes their utility among the bundles they can afford at that price. This method of pricing is known as Walrasian or *envy-free pricing.*

Dominant strategy truthfulness. The seller needs to elicit information from the buyers about their valuations in order to set the price. However, the buyers may misreport their valuations to obtain better allocations for less money, thus compromising the seller's revenue. Our goal is to design dominant strategy truthful mechanisms, where the buyers have no incentives to lie about their valuations. The truthful mechanisms are in the *prior-free* setting, i.e. do not require any prior distribution assumptions.

Market share. We evaluate the efficiency of mechanisms using the notion of *market share*, which captures the maximum buying power of any individual buyer in the market; see Section 3.2 for the formal definition.

Main Theorem (informal) *For linear multi-unit auctions with known budgets:*

- *There exists no (Walrasian) envy-free mechanism that is both truthful and non-wasteful.*
- *There exists a truthful (Walrasian) envy-free mechanism, which attains a fraction of at least* $\min\left\{\frac{1}{2}, 1 - s^*\right\}$ *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* best possible *for both the revenue and welfare objectives when the market is mildly competitive (i.e. with market share* $s^* \leq 1/2$*), and its approximation for welfare converges to* 1 *as the market becomes fully competitive.*

In the statement above, *best possible* means that there is no other truthful envy-free auction mechanism with a better approximation ratio. Our benchmarks are with respect to the optimal outcomes (for revenue and welfare) that can be obtained at an envy-free price. A mechanism is non-wasteful if it allocates as many units to the buyers as possible at a given price.

The impossibility theorem implies in particular that truthfulness is incompatible with Pareto efficiency, as wasteful mechanisms cannot be Pareto efficient. Our positive results are for *known* budgets, similarly to [1]. 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., [19]).

Finally, we also analyze the performance of our mechanism on monotone auctions, where the approximation ratio is further improved.

1.2. Related Work

The multi-unit setting has been studied in a large body of literature on auctions [1, 2, 3, 4, 5], where the focus has been on designing truthful auctions with good approximations to some desired objective, such as the social welfare or the revenue.

Very relevant to ours is the work of Dobzinski, Lavi, and Nisan [1], in which the authors study multi-unit auctions with budgets, however with no restriction to envy-free pricing or even item-pricing. The work in [1] designs a truthful auction mechanism (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 from Ausubel [20].

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 is not implied by the results in [1] 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 [2, 4, 21], and without budgets [22, 3, 5]; all of the aforementioned papers do not consider the envy-freeness constraint.

Follow-up work to [1] by Bhattacharya, Conitzer, Munagala, and Xia [21] shows that in the presence of one infinitely indivisible good, it is possible to employ the adaptive clinching auction to achieve truthfulness in the case of private budget as well, when over-reporting the budgets is not possible.

Another relevant work is that of Abrams [23], which studies linear multi-unit auctions with private budgets, with the objective of approximating the maximum revenue without any distributional assumptions. The main result in [23] is a randomized mechanism that uses two different prices and achieves a constant approximation for that objective, under a notion of market dominance quite similar to ours. That notion, coined *bidder dominance* (different from the buyer dominance of [1] mentioned above) measures the fraction of the total revenue that a single buyer can be responsible for, as opposed to the fraction of units that a buyer can aquire, which is what the market share notion that we employ captures. In fact, the bidder dominance notion of [23] coincides with the *budget share* notion used in [24], where its relation to the market share is also briefly discussed. Crucially, the results of [23] do not employ item-pricing and do not have the envy-freeness constraint, both of which are central in our investigations. Additionally, our setting only concerns deterministic mechanisms. For those reasons, our results and those of [23] and the other aforementioned works are incomparable.

The effects of strategizing in markets and market-based auction mechanisms have been studied extensively over the past years [25, 26, 27, 28, 29]. 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 [7, 30] for *pair envy-freeness*, a different notion which dictates that no buyer would want to swap its allocation with that of any other buyer [31]. 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 [8, 32]. Incentives in auctions with budgets have also been considered [33, 34, 35] from the perspective of the Price of Anarchy [36] for the objective of the liquid welfare [4], a notion of the social welfare that incorporates the budgets into its definition, coined the "liquid price of Anarchy" in [33]. Again, these works study more general auction formats that multi-unit auctions, and without the envy-freeness constraint.

As mentioned above, the good approximations achieved by our truthful mechanism are in the *prior-free setting* [37], i.e. we do not require any assumptions on prior distributions from which the input valuations are drawn.

Good prior-free approximations are usually difficult to achieve and a large part of the literature is concerned with auctions under distributional assumptions, under the umbrella of *Bayesian mechanism design* [38, 39, 40, 41, 42, 37, 43].

Following the conference version of our paper [44], there have been follow-up works in multi-unit markets and envy-free pricing; we mention the most closely related ones. Flammini, Mauro, and Tonelli [45] consider a framework for capturing the setting of fair discriminatory pricing in multi-unit markets where the agents are related via an underlying graph and each agent is only aware of the prices of the neighboring agents. In the extreme case where the graph is complete, each agent must pay the same price per unit, while in the case where the graph has no edges, each agent can be charged a different price per unit. Flammini, Mauro, Tonelli, and Vinci [46] considered the envy-free pricing via an underlying graph and obtained bounds on the revenue for topologies inspired by social networks, such as where the nodes have a power law degree distribution.

Viqueira, Greenwald, and Naroditskiy [47] considered a relaxation of envy-free pricing where only the winners are envy-free and there is a reserve price p so that all the unallocated items cost at least p. Anshelevich and Sekar [48] formulate two general techniques, called price doubling and item halving, for combinatorial markets with item pricing. By applying these methods, they obtain, e.g., mechanisms with good approximations for the revenue objective when the buyers have XOS valuations. Colini-Baldeschi, Leonardi, and Zhang [49] considered envy-free pricing mechanisms in matching markets where there are m items and n buyers with budgets and each buyer is interested in a subset of the items on sale. While computing an envy-free pricing allocation that maximizes revenue is computationally hard in this setting, the paper shows that in natural special cases such as where each buyer has a budget that is greater than her single-value valuation, it is possible to obtain a $1/4$ approximation to the optimal revenue.

1.3. Roadmap to the paper

The details of the model, utilities, demand and pricing are in Section 2. The main mechanism we design is in Section 3, while impossibility results are in Section 4. Finally, we give improved results for monotone auctions in Section 5 and a discussion in Section 6.

2. Preliminaries

Auction model. Consider a linear multi-unit auction with budgets, where a seller brings m indivisible units of a good for sale to a set $[n] = \{1, \ldots, n\}$ of buyers. Each buyer i has a valuation $v_i > 0$ and a budget $B_i > 0$. The valuation v_i indicates the value of the buyer for one unit of the good. In our setting the valuations v_i are *elicited* by the mechanism, whereas the budgets B_i are *known* information.[2] Let $\mathbf{v} = (v_1, v_2, \ldots, v_n)$ be the *valuation profile.*

An *allocation* is an assignment of units to the buyers and is denoted by a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}^n$, such that $x_i \geq 0$ 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 sets a price p per unit, such that the price of purchasing ℓ units is $p \cdot \ell$ for each buyer. The interests of the buyers at a given price are captured by the demand function.

Discrete domain. We assume the input parameters and the possible prices are drawn from a grid $\mathbb{D}$ of rational numbers. That is, $\mathbb{D} = \{\epsilon \cdot j \mid j \in \mathbb{N}\}$, for sufficiently small ϵ. In particular, ϵ will be such that $\mathbb{D}$ contains the valuations, all the numbers B_i/k for each $i \in [n]$, and each $k \in [m]$. This suffices for all of our results to hold, but in general, since the discrete domain is meant to be a good approximation of the space $\mathbb{R}_+$ of positive real numbers, we would like ϵ to be quite small, to be useful for real-world applications. A concrete example of such a domain would be $\mathbb{D} = \{0, 0.01, 0.02, 0.03, \ldots\}$, which should suffice for most such applications.

We later show in Section 3, where the main mechanism is defined, that this comes at no expense in the runtime of the mechanism as well as explain why a discrete domain is needed in the first place.

Definition 1 (Demand). *The* demand *of buyer i at a price p is a set consisting of all the possible bundle sizes (i.e. number of units) that the buyer would like to purchase at this price:*

[2] While the assumption of known budgets is necessary for the positive results to hold, obviously all the lower bounds and impossibility results become stronger under this assumption.

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

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

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

(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 (e.g. 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 1 (Non-existence of envy-free clearing prices). *Consider an auction with buyers $N = \{1, 2\}$, $m = 3$ units, valuations $v_1 = v_2 = 1.1$, and budgets $B_1 = B_2 = 1$. At each 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* obtained at an envy-free pricing. The *social welfare* at an envy-free pricing $(\mathbf{x}, p)$ is the sum of valuations of the buyers for the goods allocated: $\mathcal{SW}(\mathbf{x}, p) = \sum_{i=1}^{n} v_i \cdot x_i$. The *revenue* is the total payment received by the seller: $\mathcal{REV}(\mathbf{x}, p) = \sum_{i=1}^{n} x_i \cdot p$.

Mechanisms. The goal of the seller is to obtain money in exchange for the goods; however, it can only do that if the buyers are interested in purchasing them. The challenge for the seller is to obtain accurate information about the preferences of the buyers, which 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 [50].

An auction *mechanism* is a function $M : \mathbb{D}^n \to \mathbb{D} \times \mathbb{Z}_+^n$ that maps the valuations reported by the buyers to a price $p \in \mathbb{D}$ and an allocation $\mathbf{x} \in \mathbb{Z}_+^n$ of non-negative quantities of units. Truthful mechanisms incentivize buyers to reveal their true inputs and for this reason are often desirable, as the buyers have a simple strategy when participating in the mechanism.

Definition 3 (Truthful Mechanism). *A mechanism M is* truthful *if for every buyer $i \in [n]$ with valuation v_i, each fixed bid vector $\mathbf{v}_{-i}$ of the other bidders, and each alternative bid $s_i \in \mathbb{D}$ of bidder i, we have:*

$$u_i(M(\mathbf{v})) \geq u_i(M(s_i, v_{-i}))\,.$$

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

The next definitions and lemmas will be used extensively in the paper.

Hungry, Semi-hungry, and Essentially Hungry Buyers. Buyer i is

- *hungry* at price p if $v_i > p$, and
- *semi-hungry* at price p 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. That is, buyer i is essentially hungry if its value per unit is at least as high as the price per unit and buyer i also receives the largest non-empty bundle in its demand set.

Candidate Prices. Consider the following set:

$$\mathcal{P} = \left\{ v_i \mid \forall i \in [n] \right\} \cup \left\{ \frac{B_i}{k} \mid \forall i \in [n], \forall k \in [m] \right\}. \tag{1}$$

We call $\mathcal{P}$ the set of candidate prices; these prices are either equal to some valuation or have the property that one of the buyers could exhaust its budget by purchasing all the units it can afford. We will later show that for both revenue and welfare, w.l.o.g. the optimal solution uses a price from $\mathcal{P}$.

2.1. Useful Lemmas

In this section we show several lemmas that will be used throughout.

Lemma 1. *If p is an envy-free price, then any price $p' > p$ is also envy-free. On the other hand, if p is not an envy-free price, then any price $p' < p$ is not envy-free either.*

Proof. This follows from the fact that for every buyer i, the number of demanded units is non-increasing in the price. If at price p there are enough units to satisfy all demands, then the same holds at any price $p' > p$. Similarly, if at some price p there are not enough units to satisfy all demands, this is also the case for any price $p' < p$. □

Before we state the next lemma, we define some terminology. We will say that a price if *welfare-optimal* (respectively *revenue-optimal*) if the maximum possible social welfare (respectively revenue) at any envy-free price can be attained at this price, i.e., there is an allocation $\mathbf{x}$ such that $(\mathbf{x}, p)$ is an envy-free pricing that achieves the maximum social welfare (respectively revenue) among all possible envy-free pricings.

Lemma 2. *There always exists a welfare-optimal and a revenue-optimal envy-free price in $\mathcal{P}$. This holds even when we consider prices from the continuous domain, rather than the domain $\mathbb{D}$.*

Proof. Let q be any envy-free price; we will argue that there is an envy-free price $p \geq q$ in $\mathcal{P}$ such that the same allocation of items to the same buyers is possible at both q and p. That will imply that if q is a welfare-optimal envy-free price, then so is p, since the social welfare at q is the same as the social welfare at p. Similarly, if q is a revenue-optimal envy-free price, then it must be the case that $q \in \mathcal{P}$, as otherwise the seller could sell exactly the same amount of items at a higher price, thus obtaining more revenue.

Let p be the smallest element in $\mathcal{P}$ that is at least as large as q, i.e., $p = \min\{\mathcal{P} \mid p \geq q\}$. Obviously, if $q \in \mathcal{P}$ we are done, so assume that $q \notin \mathcal{P}$, which clearly implies that $p \neq q$ and that for any $p' \in (q, p)$, $p' \neq \mathcal{P}$. This implies that there are no semi-hungry buyers at q, as $v_i \in \mathcal{P}$ for all i, and hence $v_i \neq q$ for all i. In other words, the set of buyers that can possibly receive items at prices p and q (since their valuations are not smaller than the price) are the same. It remains to argue that each of these buyers can actually receive the same number of items in p and in q. Consider any such buyer i and let k_q and k_p be the maximum number of items that it can receive at prices q and p respectively. Since $p > q$, by the fact that the number of demanded units is non-increasing in the price, it must be the case that $k_p \leq k_q$. Assume that $k_p \neq k_q$. Then buyer i can afford k_q items at price q, which implies that $q \leq B_i/k_q$. By the assumption that $q \neq \mathcal{P}$, it holds that $q < B_i/k_q$. At the same time, the buyer but cannot afford k_q items at price k_p, which implies that $p > B_i/k_q$. That means that in (q, p), there exists a price $p' = B_i/k_q$ that is in $\mathcal{P}$, contradicting the fact that p is chosen to be the smallest element of $\mathcal{P}$ that is at least as large as q. This implies that $k_q = k_p$ and concludes the proof.

For the second part of the statement of the theorem, we remark that the arguments above did not use anywhere that $q \in \mathcal{D}$, which means that they also hold when the price q comes from the continuous domain $\mathbb{R}$. □

Remark 1. *We remark that while our domain is indeed the discrete domain $\mathbb{D}$, the second part of the statement in Lemma 2 establishes that the guarantees of our mechanism also hold when compared against the best possible social welfare and revenue attained at any envy-free price, even if that comes from the continuous domain, thus making them stronger. This is because by definition, the set $\mathcal{P}$ is contained in $\mathbb{D}$.*

We conclude the section with a simple lemma that explains how to find a revenue- or welfare-maximizing allocation at a given price in polynomial time.

Lemma 3. *For a linear multi-unit market, given an envy-free price p, a revenue- or welfare-maximizing allocation at p can be found in polynomial time in n and $\log(m)$.*

Proof. First, given the valuation functions of the hungry buyers, we can compute their demands at price p. Note these demands are singletons and so the allocation for these buyers is uniquely determined. For the non-hungry buyers (if any), we assign the remaining units (if any) in a greedy fashion: Fix an arbitrary order of buyers and assign them units according to that order, until all of them exhaust their budgets or we run out of units. All these operations can be done in polynomial time. □

3. A best possible envy-free and truthful mechanism

In this section, we present our main contribution, an envy-free and truthful mechanism, which is best-possible 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.

Theorem 1. *There exists a truthful (Walrasian) envy-free mechanism, which attains a fraction of at least*

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

on any market. This mechanism is best possible 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, which we refer to as All-or-Nothing. We first describe the mechanism at a high level and then provide the pseudocode in Algorithm 1.

All-or-Nothing Mechanism: Given as input the valuations of the buyers, let p be the minimum envy-free price. The price can be found using binary search on the set $\mathcal{P}$ rather than the whole of $\mathbb{D}$, and hence in time polynomial in $\log|P| \leq \log(2nm)$; we explain how in Lemma 4. Let $\mathbf{x}$ 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.

Input: Valuations $\mathbf{v} = (v_1, v_2, \ldots v_n)$ provided by the buyers and known budgets $\mathbf{B} = (B_1, \ldots, B_n)$.
Output: Price $p \in \mathbb{D}$ and allocation $\mathbf{x}$ such that $(\mathbf{x}, p)$ is an envy-free pricing.

Initialize $x_i = 0$ for all $i \in [n]$.
Let $\mathrm{H}(p)$ and $\mathrm{SH}(p)$ be the sets of hungry and semi-hungry buyers at price p respectively.
Compute the minimum envy-free price $p_{\min} \in \mathbb{D}$ ▷ *This can be done using binary search on $\mathcal{P}$, see Lemma 4.*
for $i \in \mathrm{H}(p_{\min})$ **do**
    Set $x_i = D_i(p_{\min})$. ▷ *Allocate the buyer its demand.*
end for
for $i \in \mathrm{SH}(p_{\min})$ in lexicographic order **do**
    if $\sum_{j \in [n]} x_j + \lfloor B_i/p \rfloor \leq m$ **then** ▷ *If it is possible to allocate the maximum element in the buyer's demand.*
        $x_i = \lfloor B_i/p \rfloor$. ▷ *Allocate that many items to the buyer (*ALL*).*
    else
        Set $x_i = 0$. ▷ *Allocate* 0 *items to the buyer (*NOTHING*).*
    end if
end for
for $i \in [n]$ such that $v_i < p$ **do**
    Set $x_i = 0$. ▷ *Allocate* 0 *items to buyers that are not hungry or semi-hungry.*
end for

Algorithm 1: ALL-OR-NOTHING

To see why a discrete domain is needed, consider the next example, which shows that on a continuous domain the minimum envy-free price might not exist.

Example 2. (The minimum envy-free price is not guaranteed to exist when the price domain is $\mathbb{R}$).

Suppose the price domain is $\mathbb{R}$. Consider an auction with $n = 2$ buyers, $m = 2$ units, valuations $v_1 = v_2 = 3$, and budgets $B_1 = B_2 = 2$. Then:

- *At each 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 each 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.*

However, by making the price domain discrete (e.g. with 0.1 *increments starting from zero), the minimum envy-free price is* 1.01*. At this price each buyer purchases* 1 *unit.*

Theorem 4 establishes the approximation guaranteed of All-or-Nothing for the discrete domain $\mathbb{D}$. Recall from Remark 1 however that Lemma 2 implies that the same guarantees hold even if we compare against the optimal social welfare and revenue possible on the continuous domain.

As mentioned above, computing the minimum envy-free price on $\mathbb{D}$ can be done via binary search. However, it can be computed much faster by observing that the minimum envy-free price (even when considered on the continuous domain) is close to a point in the set of candidate prices $\mathcal{P}$ from (1). Thus it suffices to perform binary search on $\mathcal{P}$ instead, and then check the neighboring points.

In more detail, the algorithm first performs binary search on $\mathcal{P}$ to find the smallest envy-free price $p^{\mathcal{P}}_{\min} \in \mathcal{P}$. Then, it first checks the point of the domain $\mathbb{D}$ that lies directly to the left of $p^{\mathcal{P}}_{\min} \in \mathcal{P}$; if it is not envy-free, then $p^{\mathcal{P}}_{\min} \in \mathcal{P}$ is the minimum envy-free price on $\mathbb{D}$. Otherwise, the minimum envy-free price in $\mathbb{D}$ is obviously in the interval $(q, p^{\mathcal{P}}_{\min})$, where q is the point that lies directly to the left of $p^{\mathcal{P}}_{\min}$ in $\mathcal{P}$ (not in $\mathbb{D}$). The proof of Lemma 4 establishes that the point $q + \epsilon$ that lies directly to the right of q in $\mathbb{D}$ would then be the minimum envy-free price. Since $|\mathcal{P}| = (n + m!)$, binary search on $\mathcal{P}$ has runtime $O(\log(n) + m\log(m))$.

Lemma 4. *There is an algorithm that computes the minimum envy-free price on* $\mathbb{D}$ *and has runtime* $O(\log(n) + m\log(m))$*.*

Proof. The algorithm works as follows:

- Compute via binary search a price $p^{\mathcal{P}}_{\min} \in \mathcal{P}$ such that $p^{\mathcal{P}}_{\min}$ is the smallest envy-free price in $\mathcal{P}$.

  1. If $p^{\mathcal{P}}_{\min} - \epsilon$ is not envy-free, then output $p^{\mathcal{P}}_{\min}$.
  2. Else, compute $q = \max_{p \in \mathcal{P}}\{p < p^{\mathcal{P}}_{\min}\}$ and output $q + \epsilon$.

Next we show that the algorithm correctly computes the minimum envy-free price $p_{\min}$ in $\mathbb{D}$. Assume by contradiction this is not the case. This is equivalent to $p_{\min} \notin \mathcal{P}$ and $(p_{\min} - \epsilon) \notin \mathcal{P}$.

Define $p^- = p_{\min} - \epsilon$. By definition, p^- is not envy-free, thus the total demand is

$$D(p^-) = \sum_{i \in [n]} D_i(p^-) > m\,.$$

At the same time, since $p_{\min}$ is envy-free, we have

$$D(p_{\min}) = \sum_{i \in [n]} D_i(p_{\min}) \leq m\,.$$

The demand reduction from p^- to $p_{\min}$ can only happen for two reasons:

(a) there exist buyers that are hungry at p^- but become semi-hungry at $p_{\min}$ ($v_i = p_{\min}$) or not interested ($v_i < p_{\min}$), or

(b) some buyers at p^- can no longer buy the same number of units at some price in $(p^-, p_{\min}]$.

Case (a) implies that for at least one buyer i, it holds that $v_i \in (p^-, p_{\min})$, which is not possible since there is no point in $\mathcal{P}$ in the interval $(p^-, p_{\min}]$.

Similarly, in case (b) this implies that for some buyer i and some number of units k, it holds that $B_i/k \in (p^-, p_{\min}]$, which is not possible for the same reason. This completes the proof. □

Now we move on to the guarantees of All-or-Nothing. We first show the truthfulness of the mechanism, and then prove its approximation guarantees.

3.1. Truthfulness of the All-or-Nothing *Mechanism.*

The following theorem establishes the truthfulness of All-or-Nothing.

Theorem 2. *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 $\mathbf{v} = (v_1, \ldots, v_n)$, 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 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, we have $p' < p$ 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. □

3.2. Performance of the All-or-Nothing *Mechanism.*

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

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

where $\mathcal{OBJ} \in \{\mathcal{SW}, \mathcal{REV}\}$ 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". A trivial profile is an input profile for which at any envy-free price, none of the buyers that are hungry or semi-hungry can afford to buy a single unit of the good, and thus no items are allocated in total. 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., [51], 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. We first state the approximation ratio of the mechanism with respect to the revenue objective.

Theorem 3. *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 $\mathbf{z}$, 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 $\mathbf{x}$. 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. The set of buyers that get allocated at p_{min} represent 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 can be bounded as follows

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

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, $\mathbf{z}$. However there are multiple ways of assigning the semi-hungry buyers to achieve an optimal allocation. Recall that $\mathbf{x}$ is the allocation computed by ALL-OR-NOTHING. Without loss of generality, we can assume that $\mathbf{z}$ is an optimal allocation with the property that $\mathbf{z}$ is a superset of $\mathbf{x}$ and the following condition holds:

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

This is because obviously $\mathbf{z}$ cannot allocate fewer items at the revenue-optimal price, by virtue of being an optimal allocation, and since the revenue only depends on how many items are allocated at the given price, not to which particular buyers.

We argue that $\mathbf{z}$ allocates units to at most one buyer more compared to $\mathbf{x}$. Assume by contradiction that there are at least two semi-hungry buyers i and j, such that $0 < z_i < \lfloor \alpha_i \rfloor$ and $0 < z_j < \lfloor \alpha_j \rfloor$. Then we can progressively take units from buyer j and transfer them to buyer i, until either buyer i receives $z'_i = \lfloor \alpha_i \rfloor$, or buyer j receives $z'_j = 0$. Hence we can assume that the set of semi-hungry buyers that receive non-zero, non-maximal allocations in the optimal solution $\mathbf{z}$ 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 $\tilde{z}_\ell$ 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 $\sum_{i=1}^{n} z_i$, but $z_\ell \leq \tilde{z}_\ell$, we get:

$$\frac{z_\ell}{\sum_{i=1}^{n} z_i} \leq \frac{\tilde{z}_\ell}{\sum_{i=1}^{n} z_i} = s_i^*.$$

Then we have:

$$
\begin{aligned}
\frac{\mathcal{REV}(AON)}{OPT} &= \frac{OPT - z_\ell \cdot p_{min}}{OPT} = 1 - \frac{z_\ell \cdot p_{min}}{OPT} \\
&\geq 1 - \frac{\tilde{z}_\ell \cdot p_{min}}{OPT} = 1 - \frac{\tilde{z}_\ell \cdot p_{min}}{\sum_{i=1}^n z_i \cdot p_{min}} = 1 - \frac{\tilde{z}_\ell}{\sum_{i=1}^n z_i} = 1 - s_i^* \\
&\geq 1 - s^* \\
&\geq \frac{1}{2}\,.
\end{aligned}
$$

The last inequality holds since the market share s^* is at most $1/2$. Combining the two cases, the bound follows. This completes the proof. □

Theorem 3 has the following corollary.

Corollary 1. *The approximation ratio 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 3, 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. □

The next theorem establishes the approximation ratio for the social welfare objective.

Theorem 4. *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. The proof is similar to that of Theorem 3; in fact it is easier, because of the fact that the minimum envy-free price $p_{\min}$ is a welfare-optimal price. To see this, let y be the number of units that we can allocate to the buyers at price $p_{\min}$, and let N_p be the set of buyers that receive at least one unit. At any price $p > p_{\min}$, since the demands are non-increasing in the price, we can allocate at most y units, and the only agents that can receive at least on unit are those in N_p. Since the social welfare only depends on the valuations and number of allocated units but not the price, the social welfare at $p_{\min}$ is at least that at p.

Let $\mathbf{z}$ be the social welfare-optimal allocation at $p_{\min}$, and let $\mathbf{x}$ be the allocation of the All-or-Nothing mechanism at this price. To obtain a bound on the approximation ratio of the mechanism, it suffices to bound the social welfare loss due to the semi-hungry buyers that could receive a non-zero number of units but receive zero units, due to the way the mechanism works. Again, we can assume without loss of generality that

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

Using a similar argument to that Case 2 of the proof of Theorem 3, we establish that $\mathbf{z}$ allocates units to at most one buyer compared to $\mathbf{x}$. If it allocates the same number of units to all buyers, then All-or-Nothing is optimal. Otherwise, similarly to before, let ℓ be that buyer; the welfare loss for this buyer is at most $z_\ell \cdot v_\ell$. Let $\hat{z}_\ell$ be the maximum number of units that ℓ can receive in any envy-free allocation at $p_{\min}$. As before, we obtain that

$$\frac{z_\ell}{\sum_{i=1}^n z_i} \leq \frac{\tilde{z}_\ell}{\sum_{i=1}^n z_i} = s_i^*.$$

Then we have:

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

where OPT is now the optimal welfare, and we used the fact that $v_\ell \leq v_i$ for all buyers i that receive a non-zero number of units in $\mathbf{z}$. □

Finally, All-or-Nothing is best-possible among all truthful mechanisms for both objectives whenever the market share s^* is at most $1/2$.

Theorem 5. *Let M be any truthful mechanism that always outputs an envy-free pricing. 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 cases:

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 $\frac{m}{(m/2)+1} = 2 - \frac{4}{m+2}$.

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 $\mathbf{v}^{(\mathbf{1})}$; 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 $\mathbf{v}^{(\mathbf{i})}$, $i = 1, \ldots, k$ that we construct, the optimal envy-free price is equal to the second highest reported valuation, i.e. v_2 or $v_2^{(i)}$, $i = 1, \ldots, k$ 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. $\square$

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 [2] for a different setting.

Theorem 6. *There is no Pareto efficient, truthful mechanism that always outputs an envy-free pricing, even when the budgets are known.*

Proof. Assume by contradiction that a Pareto efficient and truthful mechanism that always outputs an envy-free price exists. Consider the following instance I_1 with $n = 2$ and $m = 3$ (the instance can be adapted to work for any number of buyers by adding many buyers with very small valuations and many items by scaling the budgets appropriately): $v_1 = v_2 = 3$ and $B_1 = B_2 = 6$. It is not hard to see that the only Pareto efficient envy-free outcome is to set $p = 3$ and allocate 2 items to one buyer (wlog buyer 1) and 1 item to the other buyer. Indeed, any price $p' < p$ would not be envy-free and any price $p' > p$ would sell 0 items, yielding a utility of 0 for both agents and the auctioneer. At the same time, any allocation that does not allocate all three items at price $p = 3$ is Pareto dominated by the above allocation, since the utilities of buyers 1 and 2 would be 0, but the utility of the auctioneer would be smaller.

Now consider a new instance I_2 where $v_1 = 3$, $v_2 = 2.5$ (and it still holds that $B_1 = B_2 = 6$). We claim that the only Pareto efficient envy-free outcome

(x, q) is to set the price $q = 2.5$, allocate $x_1 = 2$ items to buyer 1 and $x_2 = 1$ item to buyer 2. At $(x, 2.5)$, the utility of buyer 1 is $u_1(x, 2.5) = 6 - 5 = 1$, the utility of buyer 2 is $u_2(x, 2.5) = 2.5 - 2.5 = 0$ and the utility of the auctioneer is $u_a(x, 2.5) = 2.5 \cdot 3 = 7.5$. The only other possible allocation x' at price 2.5 would be $x'_1 = 2$ (since buyer 1 is hungry at price 2.5) and $x'_2 = 0$, which is Pareto dominated by $(x, 2.5)$. Therefore, for another Pareto efficient pair (x', q') to exist, it would have to hold that $q' \neq 2.5$.

Obviously, any choice $q' < 2.5$ is not envy-free and therefore we only need to consider the case when $q' > 2.5$. At any such price q', the utility of buyer 1 is at most 1, since the buyer can purchase $x'_1 \leq 2$ items at a price strictly higher than 2.5, the utility of buyer 2 is 0 since the price is higher than its valuation and hence it gets $x'_2 = 0$ items, and finally, the utility of the auctioneer is at most 6, since it can only sell at most two items at a price no higher than 3. This means that (x', q') is Pareto dominated by $(x, 2.5)$.

The paragraphs above establishes that on Instance I_1, buyer 2 receives one item at price 3 and on instance I_2, buyer 2 receives one item at price 2.5. But then, buyer 2 would have an incentive to misreport his valuation on instance I_1 as being $v'_2 = 2.5$ and receive the same number of items at a lower price, thus increasing its utility and contradicting truthfulness.

Since the proof only requires valuations and budgets to lie on points 2.5, 3 and 6, the theorem also holds for the discrete domain. □

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

For the proof of the theorem, we will need the following lemma, which restricts the price of any mechanism with the properties stated in Theorem 7 to be either the minimum envy-free price, or the next point on $\mathbb{D}$.

Lemma 5. *Let M be an in-range, non-wasteful and truthful mechanism. Then on any valuation profile $\mathbf{v}$ which is not trivial, M must output a price $p \in \{p_{\min}, p_{\min} + \gamma\}$, where $p_{\min} = \min\{p \in \mathbb{D} : p$ is envy-free on $\mathbf{v}\}$ and γ is the distance between two consecutive elements of $\mathbb{D}$.*

Proof. Assume by contradiction that M does not always output a price $p \in \{p_{\min}, p_{\min} + \gamma\}$. Let $\mathbf{v} = (v_1, \ldots, v_n)$ be any valuation profile that is not trivial and let p_v be the price outputted by M; by assumption, it holds that $p_v > p_{\min} + \gamma$. By the assumption that M is in-range, it holds that $v_j \geq p_v$ for some relevant buyer $j \in N$. Define

$$J = \{j : v_j \geq p_v : j \text{ is allocated a non-zero number of units}\}$$

as the set of all relevant buyers with valuations at least as high as the envy-free price chosen by M.

Now, consider an instance $\mathbf{v^1}$ such that $v_i^1 = v_i$ for all buyers $i \in [n] \backslash \{j_1\}$ and $v_{j_1}^1 = p_{\min} + \gamma$ for some buyer $j_1 \in J$, i.e. the instance obtained by $\mathbf{v}$ when some buyer $j_1 \in J$ reports a valuation equal to $p_{\min} + \gamma$. Let p_1 be the price outputted by M on input $\mathbf{v^1}$. Note that since on instance $\mathbf{v^1}$ buyer j_1's valuation is still higher than $p_{\min}$, it holds that $p_{\min}$ is still the minimum envy-free price in $\mathbb{O}$ on the profile $\mathbf{v^1}$.

- Assume first that $p_1 = p_{\min}$. Then buyer j_1 on input profile $\mathbf{v}$ would have an incentive to misreport its valuation as $v_{j_1}^1 = p_{\min} + \gamma$; that would lower the price and since B_{j_1} is fixed, the buyer would receive at least as many units at a lower price (since it still appears to be hungry at price $p_{\min}$). This would contradict the truthfulness of M.

- Now consider the case when $p_1 = v_{j_1}^1 = p_{\min} + \gamma$. Note that since $p_1 > p_{\min}$, it holds that $\lfloor B_{j_1}/p_1 \rfloor \leq \lfloor B_{j_1}/p_{\min} \rfloor$, i.e. buyer j_1 can not demand more units at price p_1 compared to $p_{\min}$. On profile $\mathbf{v}$, it would be possible to allocate $\lfloor B_{j_1}/p_{\min} \rfloor$ units to buyer j_1 at price $p_{\min}$, therefore on profile $\mathbf{v^1}$, it is possible to allocate $\lfloor B_{j_1}/p_1 \rfloor$ units to buyer j_1 at price $p_1 = p_{\min} + \gamma$. Buyer j_1 is semi-hungry at p_1 but since M is non-wasteful, it must allocate at least $\lfloor B_i/p_1 \rfloor \geq \lfloor B_i/p_v \rfloor$ units to buyer j_1 at a price $p_1 < p_v$, and buyer j_1 increases its utility by misreporting.

From the discussion above, it must hold that $p_1 > p_{\min}+\gamma$. For the valuation profile $\mathbf{v^1}$ (which can be seen as the different instance where buyer 1 has deviated from v_1 to $p_{\min}+\gamma$), update the set

$$J := \{j \mid \text{ buyer } j \text{ is allocated a non-zero number of units and } v_j \geq p_1\}.$$

If $J = \emptyset$, then Mechanism M is not in-range and we have obtained a contradiction. Otherwise, there must exist some other buyer $j_2 \in J$ with valuation higher than p_1.

Now, consider such a buyer $j_2 \in J$ and the instance $\mathbf{v^2}$ such that $v_i^2 = v_i^1$ for all buyers $i \in N \backslash \{j_2\}$ and $v_{j_2}^2 = p_{\min} + \gamma$ for buyer j_2, i.e. the instance obtained from $\mathbf{v^1}$ when some buyer j_2 in J misreports its value being between $p_{\min}+\gamma$. Note that for the same reasons explained above, $p_{\min}$ is the minimum envy-free price in $\mathbb{O}$ on profile $\mathbf{v}^2$ as well. Let p_2 be the price outputted by M on valuation profile $\mathbf{v^2}$. Using exactly the same arguments as we did before, we can argue that by truthfulness, it holds that $p_2 \notin \{p_{\min}, p_{\min}+\gamma\}$ and therefore it must hold that $p_2 > p_{\min}+\gamma$, as every other choice is not envy-free.

By iteratively considering sequences of valuations obtained in this manner, we eventually obtain an instance $\mathbf{v^{k-1}}$ such that $J = \{j_k\}$, i.e. there is only one buyer with a valuation higher than the envy-free price p_{k-1} output by M. Repeating the argument once more will result in a valuation profile $\mathbf{v^k}$ where the price p_k is higher than the reported valuation $v_{j_k}^k = p_{\min}+\gamma$ of buyer j_k and the set J will be empty, contradicting the fact that M is in-range.

Overall, this implies that M either violates truthfulness, non-wastefulness or in-range, contradicting our assumption. □

We remark here that in the continuous domain, Lemma 5 can be strengthened so that M can only output the minimum envy-free price, whenever it exists.

Using Lemma 5, we can now prove the theorem.

Proof. (of Theorem 7) Assume by contradiction that such an in-range, non-wasteful and truthful mechanism M exists. We will consider three different instances[3] with $n = 2$ and $m = 3$, denoted (v_1, v_2) where v_1 denotes the

[3] The instances can be extended to any number of buyers by simply adding buyers with very low valuations and to many items by scaling the valuations and budgets appropriately.

valuation of buyer 1 and v_2 denotes the valuation of buyer 2, with budgets $B_1 = B_2 = 6 + 2\gamma$.

First, consider the instance $(2.5, 2.5)$ and note that since the instance is not trivial and since the minimum envy-free price is 2.5, by Lemma 5, the price chosen by M for this instance must be either 2.5 or $2.5+\gamma$. Furthermore, since M is in-range, the price can not be $2.5+\gamma$, therefore the price chosen on $(2.5, 2.5)$ is 2.5. Since M is non-wasteful and each buyer can afford exactly 2 items at price 2.5 and there are 3 available items, one buyer (wlog buyer 1) gets allocated 2 items and the other buyer (wlog buyer 2) gets allocated 1 item at this price.

Now consider the instance $(3, 2.5)$ and note that since it is not trivial and since again, 2.5 is the minimum envy-free price, M must either output 2.5 or $2.5 + \gamma$ as the price. Assume first that M selects the price to be $2.5 + \gamma$. Since buyer 1 is hungry at this price and can afford to buy exactly 2 units, its allocation on instance $(3, 2.5)$ is 2 units at price $2.5 + \gamma$. But then, on instance $(3, 2.5)$ buyer 1 would have an incentive to misreport its valuation as being 2.5 since on the resulting instance, which is $(2.5, 2.5)$, it still receives 2 items at a lower price, increasing its utility. Note that if it was buyer 2 that received 2 items on instance $(2.5, 2.5)$, we could have made the same argument using instance $(2.5, 3)$ instead.

Finally, assume that on instance $(3, 2.5)$, M outputs 2.5 as the price. By non-wastefulness, buyer 2 receives exactly 1 unit at this price. But then, consider the instance $(3, 3)$, where, using the same arguments as in the case of instance $(2.5, 2.5)$, Mechanism M must output 3 as the price and allocate 2 units to one buyer and 1 unit to the other buyer. Crucially, both buyers have utility 0 on instance $(3, 3)$. But then, buyer 2 could misreport its valuation as being 2.5, resulting in instance $(3, 2.5)$ where it receives 1 unit at a price lower than its actual valuation, benefiting from the misreport. This contradicts truthfulness.

Assume by contradiction that such an in-range, non-wasteful and truthful mechanism M exists. Consider the same instance I_1 as the one used in the proof of Theorem 6, with $n = 2$, $m = 3$ and $v_1 = v_2 = 3$ and $B_1 = B_2 = 6+2\gamma$. (Again the proof can be generalized to many agents and units similarly to the proof of Theorem 6). By Lemma 5 and since I_1 is not trivial, M must either output $p = 3$ or $p = 3 + \gamma$ and by the fact that it is in-range, it must output $p = 3$. Since M is non-wasteful, it must allocate 2 units to one of the buyers with valuation 3 (wlog buyer 1) and 1 unit to the other buyer.

Now consider an instance I_{2a} where $v_1' = 3$ and $v_2' = 2.5$. Since 2.5 is now the minimum envy-free price and I_2 is again not trivial, M must output either $p' = 2.5$ or $p = 2.5 + \gamma$. We will obtain a contradiction for each case. Assume first that $p' = 2.5$; since buyer 1 is hungry, it must hold that $x'_1 = 2$ and by non-wastefulness, it must hold that $x'_1 = 1$. In that case however, for the same reason explained in the proof of Theorem 6, $v'_2 = 2.5$ could be a beneficial deviation of buyer 2 on instance I_1, violating truthfulness. Now we argue for the case when $p' = 2.5 + \gamma$. Consider the instance I_3 where $\bar{v}_1 = \bar{v}_2 = 2.5$. Since M is in-range and I_3 is not trivial, M must select price $\bar{p} = 2.5$, since every other price is either not envy-free, or higher than all the valuations. By non-wastefulness, one buyer must receive 2 units at $\bar{p}$ and the other agent must receive 1 unit (because each buyer can afford exactly 2 units and there are 3 units available). If buyer 1 receives 2 units, i.e. $\bar{x}_1 = 2$, misreporting its valuation on instance I_{2a} as 2.5 would give the buyer higher utility, since it gets allocated the same number of items at a lower price. It remains to deal with the case when on instance I_3, buyer 1 is allocated 1 item and buyer 2 is allocated 2 items, i.e. $\bar{x}_1 = 1$ and $\bar{x}_2 = 2$.

Now consider the instance I_{2b} where $\hat{v}_1 = 2.5$ and $\hat{v}_2 = 3$, i.e. instance I_{2b} is exactly the same as instance I_{2a} with the indices of the two buyers swapped. Again, since instance I_{2b} is not trivial, by Lemma 5, M must output a price $\hat{p} \in \{2.5, 2.5 + \gamma\}$. If $\hat{p} = 2.5 + \gamma$, then we consider again Instance I_3. Since on that instance $\bar{p} = 2.5$ and $\bar{x}_2 = 2$ by the assumption above, buyer 2 has an incentive to misreport its valuation on instance I_{2b} as being 2.5, contradicting truthfulness. Therefore, it must hold that $\hat{p} = 2.5$ on instance I_{2b}.

However, by non-wastefulness, buyer 1 receives one unit at price $\hat{p}$ on instance I_{2b}, i.e. $\hat{x}_1 = 1$. We will consider the 2.5 as a potential deviation of buyer 1 on instance I_1 (where its true valuation is $v_1 = 3$). The utility of the buyer before misreporting is 0 (since the chosen price on instance I_1 is $p = 3$) whereas the utility after misreporting is $3 - 2.5 = 0.5$, i.e. strictly positive. Therefore, buyer 1 has a beneficial deviation on instance I_1, violating the truthfulness of M.

By truthfulness, it must also hold that $\bar{p} \geq 2.5 + \gamma$, otherwise on instance I_2 buyer 1 would have an incentive to misreport its valuation as $2.5 + \gamma$ and still receive 2 items at a lower price (since at any price $p < 2.5 + \gamma$ buyer 1 on instance I_3 is hungry). From the discussion above, it must hold that $\bar{p} = 2.5 + \gamma$ and by non-wastefulness and since buyer 1 can afford two items at price $2.5 + \gamma$, it must hold that $\bar{x}_1 = 2$. $\square$

5. Monotone Auctions

In the previous sections, we proved the approximation ratio guarantees of Mechanism All-or-Nothing, as a function of the market share. In this section, we will examine two special cases. The first case of common budgets is as follows:

- The budgets are *common* when $B_i = B$ for all buyers $i \in N$.

The second, more general, case is the class of monotone auctions:

- The budgets are monotone in the valuations when $v_i \geq v_j \Leftrightarrow B_i \geq B_j$. We call such auctions *monotone*.

Note that the second case is more general than the first, where for the right-hand side we have $B_i = B_j$ for all $i, j \in N$.

We will prove that in these scenarios, Mechanism All-or-Nothing is best possible among all truthful mechanisms, for both the welfare and the revenue objective. For the welfare objective, the approximation ratio guarantee will be completely independent of the market share. For the revenue objective, the dependence will be rather weak; we prove that the bound holds in all auctions except monopsonies. A monopsony is an auction in which a single buyer can afford to buy all the items at a very high price.

Definition 4. *An auction is a* monopsony, *if the buyer with the highest valuation v_1 has enough budget B_1 to buy all the units at a price equal to the second highest valuation v_2.*

Note that when the market is not a monopsony, that implies that the market share s^* is less than 1.[4]

The proof of the following theorem follows along similar arguments as those of the results in previous sections, namely Theorem 3, Theorem 4, and Theorem 5. We provide a full proof for completeness.

Theorem 8. *The approximation ratio of Mechanism* All-or-Nothing *for monotone auctions is*

[4]Note that instead of ruling out monopsonies, another approach would be to consider a different benchmark, that does not include the case of an omnipotent buyer, like the $\mathrm{EFO}^{(2)}$ benchmark for revenue, see [37], Chapter 6.

- *at most 2 for the social welfare objective.*
- *at most 2 for the revenue objective when the auction is not a monopsony.*

Furthermore, no truthful mechanism can achieve an approximation ratio smaller than $2 - \frac{4}{m+2}$ *even in the case of common budgets.*

Proof. First, note that the profile constructed in Theorem 5 is one where the budgets are common and therefore the lower bound extends to both cases mentioned above. Therefore, it suffices to prove the approximation ratio of Mechanism All-or-Nothing for both objectives, when the auction is monotone.

We start from the social welfare objective and consider an arbitrary profile $\mathbf{v}$. Without loss of generality, we can assume that $\mathbf{v}$ is not trivial (otherwise the optimal allocation allocates 0 items in total) and note that the optimal envy-free price is $p^* = p_{min}$ and let $\mathbf{x}$ be the corresponding optimal allocation. Following the arguments in the proof of Theorem 3, we establish that the according to $\mathbf{x}$ at most one additional semi-hungry buyer is allocated a positive number of units, compared to the allocation of Mechanism All-or-Nothing; let ℓ be that buyer and let x_ℓ be its optimal allocation.

The social welfare loss of Mechanism All-or-Nothing is $x_\ell \cdot v_\ell \leq v_\ell \cdot \lfloor B_\ell/v_\ell \rfloor$, i.e. the contribution of the the semi-hungry buyer that receives 0 items by All-or-Nothing, in contrast to the optimal allocation. Since the profile $\mathbf{v}$ is not trivial, there exists at least on other buyer j that receives $\min\{m, \lfloor B_j/v_\ell \rfloor\}$ units in the optimal allocation $\mathbf{x}$. If it receives m units, then $x_\ell = 0$ and the ratio on the profile is 1. Otherwise, the contribution to the welfare (for both the optimal allocation and the allocation of All-or-Nothing) from buyer j is $v_j \cdot \lfloor B_j/v_\ell \rfloor\} \geq v_j \cdot \lfloor B_\ell/v_\ell \rfloor\}$, since $v_\ell \leq v_j \Leftrightarrow B_\ell \leq B_j$ by the monotonicity of the auction. Then we have:

$$\begin{aligned}
\frac{\mathcal{SW}(AON)}{OPT} &\geq \frac{OPT - v_\ell \cdot \lfloor B_\ell/v_\ell \rfloor}{OPT} = 1 - \frac{v_\ell \cdot \lfloor B_\ell/v_\ell \rfloor}{OPT} \\
&\geq 1 - \frac{v_\ell \cdot \lfloor B_\ell/v_\ell \rfloor}{(v_\ell + v_j) \cdot \lfloor B_\ell/v_\ell \rfloor} = 1 - \frac{v_\ell}{v_j + v_\ell} \geq \frac{1}{2}.
\end{aligned}$$

For the revenue objective, again let p^* be the optimal envy-free price and let $\mathbf{x}$ be the corresponding allocation. We consider two cases:

- $p^* = p_{min}$: The argument in this case is very similar to the one used above for the social welfare objective. In particular, since $p^* = p_{min} = v_\ell$, we now have that the loss in revenue from the semi-hungry buyer ℓ for Mechanism All-or-Nothing is at most $x_\ell \cdot v_\ell \leq v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor$ whereas the contribution from buyer j is $v_\ell \cdot \lfloor B_j / v_\ell \rfloor \}$, which is at most $v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor \}$ by the monotonicity of the auction. Therefore, we have that:

$$
\begin{aligned}
\frac{\mathcal{REV}(AON)}{OPT} &\geq \frac{OPT - v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} = 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} \\
&\geq 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{2 v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor} = 1 - \frac{v_\ell}{2 v_\ell} = \frac{1}{2}.
\end{aligned}
$$

- $p^* > p_{min}$. In that case, the argument is exactly the same as in Case 2 of the proof of Theorem 3, which holds when the market share is less than 1, i.e. when the auction is not a monopsony.

□

To complete the picture, we prove in the following that if the auction is a monopsony, the approximation ratio of any truthful mechanism is unbounded. This can be captured by the following theorem.

Theorem 9. *If the auction is a monopsony, the approximation ratio of any truthful mechanism for the revenue objective is at least $\mathcal{B}$ for any $\mathcal{B} > 1$, even if the budgets are public.*

Proof. Consider the following monopsony. Let $i_1 = \text{argmax}_i\, v_i$, for $i = 1, \ldots, n$ be a single buyer with the highest valuation and denote $v_{i_1} = v_1$ for ease of notation. Similarly, let $i_2 \in \text{argmax}_{i \in N \setminus \{i_1\}}\, v_i$ be one buyer with the second largest valuation and let $v_{i_2} = v_2$. Furthermore, let $v_1 > v_i$ for all $i \neq i_1$ and $B_{i_1} = p \cdot m$, for some $v_2 < p \leq v_1$ i.e. buyer i_1 can afford to buy all the units at some price $p > v_2$. Additionally, let $B_{i_2} \geq v_2$, i.e. buyer i_2 can afford to buy at least one unit at price v_2.[5] Finally, for a given $\mathcal{B} > 1$ let

[5] Note that setting $B_{i_2} = B_{i_1}$ satisfies this constraint and creates an auction with identical budgets, so the proof goes through for that case as well.

v_2 and p be such that $\mathcal{B} = p/v_2$. Note that the revenue-maximizing envy-free price for the instance $\mathbf{v}$ is at least p and the maximum revenue is at least $p \cdot m$.

Assume for contradiction that there exists a truthful mechanism M with approximation ratio smaller than $\mathcal{B}$ and let p^* be the envy-free price output by M on $\mathbf{v}$. Since p^* is envy-free and $B_{i_1} > v_1 \cdot m$ and $B_{i_2} \geq v_2$, it can not be the case that $p^* < v$, otherwise there would be over-demand for the units. Furthermore, by assumption it can not be the case that $p^* = v_2$ as otherwise the ratio would be $\mathcal{B}$ and therefore it must hold that $p^* > v_2$.

Now let $\mathbf{v}'$ be the instance where all buyers have the same valuation as in $\mathbf{v}$ except for buyer i_1 that has value v_1' such that $v < v_1' < p^*$; let $\tilde{p}$ be the envy-free price that M outputs on input $\mathbf{v}'$.

- If $\tilde{p} > v_1'$, then the ratio of M on the instance $\mathbf{v}'$ is infinite, a contradiction.

- If $\tilde{p} \leq v_1'$ and since $\tilde{p}$ is envy-free, it holds that $v_2 \leq \tilde{p} < p^*$. In that case however, on instance $\mathbf{v}$, buyer i_1 would have an incentive to misreport its valuation as v_1' and reduce the price. The buyer still receives all the units at a lower price and hence its utility increases as a result of the deviation, contradicting the truthfulness of M.

This completes the proof of the theorem. □

6. 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.

Another interesting direction is to consider the case of private budgets; for this case, it is not very difficult to see that a simple class of mechanisms based on ordered statistics (i.e., mechanisms that select the agent with the k-th smallest valuation and set the price to be that valuation) are truthful, but their welfare or revenue guarantees might be rather poor. Whether truthful mechanisms with good approximations for either objective exist for private budgets is an interesting open question. Interestingly, showing general lower bounds for settings where the market share is bounded by e.g., 50% seems to be challenging. It would also be interesting to obtain a complete characterization of truthfulness in the case of private or known budgets.

References

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

[2] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, A. Saberi, Multi-unit auctions with budget-constrained bidders, in: Proceedings of the 6th ACM Conference on Electronic Commerce, 2005, pp. 44–51.

[3] S. Dobzinski, N. Nisan, Mechanisms for multi-unit auctions, Journal of Artificial Intelligence Research 37 (2010) 85–98.

[4] S. Dobzinski, R. P. Leme, Efficiency guarantees in auctions with budgets, in: International Colloquium on Automata, Languages, and Programming, Springer, 2014, pp. 392–404.

[5] S. Dobzinski, N. Nisan, Multi-unit auctions: beyond roberts, Journal of Economic Theory 156 (2015) 14–44.

[6] E. T. Anderson, D. I. Simester, Price stickiness and customer antagonism, The quarterly journal of economics 125 (2) (2010) 729–765.

[7] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, F. McSherry, On profit-maximizing envy-free pricing, in: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, 2005, pp. 1164–1173.

[8] M. Feldman, A. Fiat, S. Leonardi, P. Sankowski, Revenue maximizing envy-free multi-unit auctions with budgets, in: Proceedings of the 13th ACM Conference on Electronic Commerce, 2012, pp. 532–549.

[9] V. Cohen-Addad, A. Eden, M. Feldman, A. Fiat, The invisible hand of dynamic market pricing, in: Proceedings of the 2016 ACM Conference on Economics and Computation, 2016, pp. 383–400.

[10] M. Feldman, N. Gravin, B. Lucier, Combinatorial auctions via posted prices, in: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete algorithms, SIAM, 2014, pp. 123–135.

[11] S. Sekar, Posted pricing sans discrimination, in: Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017, pp. 388–394.

[12] L. Walras, Elements d'economie politique pure, ou theorie de la richesse sociale (in french)English translation: Elements of pure economics; or, the theory of social wealth. American Economic Association and the Royal Economic Society, 1954 (1874).

[13] A. S. Kelso, V. P. Crawford, Job matching, coalition formation, and gross substitutes, Econometrica 50 (1982) 1483–1504.

[14] F. Gul, E. Stacchetti, Walrasian equilibrium with gross substitutes, Journal of Economic Theory 87 (1) (1999) 95–124.

[15] M. Babaioff, B. Lucier, N. Nisan, R. Paes Leme, On the efficiency of the walrasian mechanism, in: Proceedings of the 15th ACM Conference on Economics and Computation, 2014, pp. 783–800.

[16] M. Cheung, C. Swamy, Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply, in: Proceeings of the 49th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2008, pp. 35–44.

[17] J. D. Hartline, T. Roughgarden, Simple versus optimal mechanisms, in: Proceedings of the 10th ACM Conference on Electronic Commerce, 2009, pp. 225–234.

[18] M.-F. Balcan, A. Blum, Y. Mansour, Item pricing for revenue maximization, in: Proceedings of the 9th ACM Conference on Electronic Commerce, 2008, pp. 50–59.

[19] J. M. Liberti, M. A. Petersen, Information: Hard and soft, Review of Corporate Finance Studies 8 (1) (2019) 1–41.

[20] L. M. Ausubel, An efficient ascending-bid auction for multiple objects, The American Economic Review 94 (5) (2004) 1452–1475.

[21] S. Bhattacharya, V. Conitzer, K. Munagala, L. Xia, Incentive compatible budget elicitation in multi-unit auctions, in: Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms, SIAM, 2010, pp. 554–572.

[22] Y. Bartal, R. Gonen, N. Nisan, Incentive compatible multi unit combinatorial auctions, in: Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge, 2003, pp. 72–87.

[23] Z. Abrams, Revenue maximization when bidders have budgets, in: Proceedings of the 17th Annual ACM-SIAM symposium on Discrete Algorithms, 2006, pp. 1074–1082.

[24] S. Brânzei, A. Filos-Ratsikas, Walrasian dynamics in multi-unit markets, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 1812–1819.

[25] A. Borodin, O. Lev, T. Strangway, Budgetary effects on pricing equilibrium in online markets, in: 15th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2016, International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2016, pp. 95–103.

[26] S. Brânzei, Y. Chen, X. Deng, A. Filos-Ratsikas, S. Frederiksen, J. Zhang, The fisher market game: Equilibrium and welfare, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 28, 2014.

[27] N. Chen, X. Deng, J. Zhang, How profitable are strategic behaviors in a market?, in: Proceedings of the 19th European Symposium on Algorithms, Springer, 2011, pp. 106–118.

[28] R. Mehta, M. Sohoni, Exchange markets: Strategy meets supply-awareness, in: Proceedings of the 9th International Conference on Web and Internet Economics, 2013, pp. 361–362.

[29] R. Mehta, N. Thain, L. Végh, A. Vetta, To save or not to save: The fisher game, in: Proceedings of the 10th International Conference on Web and Internet Economics, Springer, 2014, pp. 294–307.

[30] J. Hartline, Q. Yan, Envy, truth, and profit, in: Proceedings of the 12th ACM Conference on Electronic Commerce, 2011, pp. 243–252.

[31] E. Markakis, A. Ntokos, O. Telelis, On envy-free revenue approximation for combinatorial buyers with budgets, Theory of Computing Systems 64 (8) (2020) 1338–1361.

[32] R. Colini-Baldeschi, S. Leonardi, P. Sankowski, Q. Zhang, Revenue maximizing envy-free fixed-price auctions with budgets, in: Proceedings of the 10th International Conference on Web and Internet Economics, 2014, pp. 233–246.

[33] Y. Azar, M. Feldman, N. Gravin, A. Roytman, Liquid price of anarchy, in: Proceedings of the 10th International Symposium on Algorithmic Game Theory, Springer, 2017, pp. 3–15.

[34] I. Caragiannis, A. Voudouris, The efficiency of resource allocation mechanisms for budget-constrained users, Mathematics of Operations Research 46 (2) (2021) 503–523.

[35] A. A. Voudouris, Simple combinatorial auctions with budget constraints, Theoretical Computer Science 842 (2020) 6–17.

[36] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, in: Annual symposium on theoretical aspects of computer science, Springer, 1999, pp. 404–413.

[37] J. D. Hartline, Mechanism design and approximation, Book draft. October 122 (2013).

[38] N. R. Devanur, S. M. Weinberg, The optimal mechanism for selling to a budget constrained buyer: The general case, in: Proceedings of the 18th ACM Conference on Economics and Computation, 2017, pp. 39–40.

[39] J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, T. Vredeveld, Posted price mechanisms for a random stream of customers, in: Proceedings of the 18th ACM Conference on Economics and Computation, 2017, pp. 169–186.

[40] N. R. Devanur, N. Haghpanah, A. Psomas, Optimal multi-unit mechanisms with private demands, Games and Economic Behavior 121 (2020) 482–505.

[41] Y. Cai, C. Daskalakis, S. M. Weinberg, Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization, in: 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, IEEE, 2012, pp. 130–139.

[42] Y. Cai, C. Daskalakis, S. M. Weinberg, Reducing revenue to welfare maximization: Approximation algorithms and other generalizations, in: Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2013, pp. 578–595.

[43] R. B. Myerson, Optimal auction design, Mathematics of Operations Research 6 (1) (1981) 58–73.

[44] S. Brânzei, A. Filos-Ratsikas, P. B. Miltersen, Y. Zeng, Walrasian pricing in multi-unit auctions, in: 42nd International Symposium on Mathematical Foundations of Computer Science, MFCS, Vol. 83 of LIPIcs, 2017, pp. 80:1–80:14.

[45] M. Flammini, M. Mauro, M. Tonelli, On fair price discrimination in multi-unit markets, Artificial Intelligence 290 (2021) 103388.

[46] M. Flammini, M. Mauro, M. Tonelli, C. Vinci, Inequity aversion pricing in multi-unit markets, in: Proceedings of the 24th European Conference on Artificial Intelligence, Vol. 325 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2020, pp. 91–98.

[47] E. Areyan Viqueira, A. Greenwald, V. Naroditskiy, On approximate welfare- and revenue-maximizing equilibria for size-interchangeable bidders, in: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2017, p. 1466–1468.

[48] E. Anshelevich, S. Sekar, Price doubling and item halving: Robust revenue guarantees for item pricing, in: Proceedings of the 18th ACM Conference on Economics and Computation, Association for Computing Machinery, 2017, p. 325–342.

[49] R. Colini-Baldeschi, S. Leonardi, Q. Zhang, Revenue maximizing envy-free pricing in matching markets with budgets, in: Proceedings of the 12th International Conference on Web and Internet Economics, Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 207–220.

[50] N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, Algorithmic Game Theory, Cambridge Univ. Press, (editors) 2007.

[51] P. Farris, N. Bendle, P. Pfeifer, D. Reibstein, Marketing metrics: The definitive guide to measuring marketing performance, Pearson Education, 2010.