Walrasian Pricing in Multi-unit Auctions*

Simina Brânzei[1], Aris Filos-Ratsikas[2], Peter Bro Miltersen[3], and Yulong Zeng[4]

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 Univesity, 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.2017.23

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

* 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). The authors also aknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation and from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. A part of this work was done when Simina was visiting the Simons Institute for the Theory of Computing.

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.

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 [15, 6, 17, 16, 18].

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 lead to a body of literature focused on achieving fair pricing [23, 20, 13, 21, 34].

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 [35] —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., [27, 22]) and the induced mechanism – the Walrasian mechanism [3, 12] – 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 [23], 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 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* [24], which recently proposed item pricing [4, 21] as a way of designing simpler auctions while at the same time avoiding the ill effects of discriminatory pricing [20, 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.

[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 [23] which applies to unit-pricing, due to its attractive fairness properties [20].

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. Our main contributions can be summarized as follows.

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

- *There exists no envy-free mechanism that is both truthful and non-wasteful.*
- *There exists a truthful envy-free and prior-free auction, which attains at least half of the optimal revenue and at least* $1 - s^*$ *of the optimal welfare in competitive auctions, i.e. where the market share* s^* *is at most* 50%*. This mechanism is optimal for both the revenue and welfare objectives, and its approximation for welfare converges to* 1 *as the auction becomes fully competitive.*

A market share of at most 50% roughly means that no buyer can purchase more than half of the resources when competition is maximal. The impossibility theorem implies in particular that truthfulness is incompatible with Pareto efficiency. Our positive results are for *known* budgets, similarly to [15]. 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., [33]).

1.2 Related Work

The multi-unit setting has been studied in a large body of literature on auctions ([15, 6, 17, 16, 18]), 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 [15], 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 [15] 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 [16] and [6], and without budgets ([18, 5, 17]); 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, 11, 29, 30]). 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 ([23]) and ([25]) for *pair envy-freeness*, a different notion which dictates that no buyer would want to swap its allocation with that of any other buyer [28]. 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 [20, 14].

It is worth mentioning that the good approximations achieved by our truthful mechanism are a *prior-free setting* ([26]), i.e. we don't require any assumptions on prior distributions from which the input valuations are drawn. 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* ([9, 10, 26, 31]).

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 1** (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 2** (Utility). The *utility* of buyer i given a price p and an allocation $\mathbf{x}$ is

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

Envy-free Pricing. An allocation and price $(\mathbf{x}, p)$ represent an *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 3** (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 is 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 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 [32].

An auction *mechanism* is a function $M : \mathbb{V}^n \to \mathbb{O} \times \mathbb{Z}_+^n$ 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 drawn[2], and an allocation vector $\mathbf{x} \in \mathbb{Z}_+^n$.

▶ **Definition 4** (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.

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.

▶ **Theorem 5.** *There exists a truthful prior-free and envy-free auction, such that its approximation ratio*

- *or revenue is at most* 2*, whenever* $s^* \leq 1/2$*,*
- *for welfare is at most* $1/(1 - s^*)$*, whenever* $s^* < 1$*,*

where $0 \leq s^* \leq 1$ *is the market share.*

Moreover, this mechanism is optimal among all truthful mechanisms for both objectives whenever the market share $s^* \leq 1/2$*. The approximation for welfare converges to* 1 *as the auction becomes fully competitive.*

Consider the following mechanism.

[2] In principle the spaces $\mathbb{V}$ and $\mathbb{O}$ can be the same but in some scenarios, 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.

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.

▶ Remark. *In order for the mechanism to be well-defined, we need to assume that the price domain $\mathbb{O}$ is discrete; in fact when the domain is $\mathbb{R}$ (i.e. the price can be any real number), the minimum envy-free price does not necessarily exist, as we show next.*[3] *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.* ◀

Both the valuation v_i and budget B_i of each player will be assumed to be rational numbers, so that the denominator and numerator of each valuation and budget are represented with t bits of precision. The price output by the mechanism will be defined on a slightly finer grid, with s bits of precision, where $s \geq 4t + 4 \log m$. The All-or-Nothing mechanism runs in time polynomial in t and s, since the minimum envy-free price can be found using binary search, 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 in Appendix B.

▶ **Theorem 6.** *When the valuation and budget of each buyer are rational numbers, where both the numerator and denominator are integers specified with t bits of precision, and the price is specified with $s \geq 4t + 4 \log m$ bits, 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 $O(t + s)$.*

Thus by setting $s = 4t + 4 \log m$, we simultaneously obtain a runtime that is polynomial in the input, $O(t + \log m)$, and that there is no loss from restricting All-or-Nothing to output a price on the grid, compared to the revenue and welfare obtained when the price can

[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.

be an arbitrary real number (or the supremum value if the minimum envy-free price does not exist).

Truthfulness of the All-or-Nothing Mechanism

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

▶ **Theorem 7.** *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'. To see this, first consider the case when $p = p'$. Since buyer i was essentially hungry at price p, and p is envy-free, it received the maximal set of items possible by the budget restriction and d_i be that set. Since the price has not changed, the agent can not receive more than d_i and hence its utility is not larger that before. Now consider the case when $p < p'$. Since p is an envy-free price, buyer i receives a maximal d_i 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 $x_j' > x_j$. 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 the 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. This contradicts the assumption that buyer i is not essentially hungry at price 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.

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

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", 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., [19], 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 8.** *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 $\tilde{x}_\ell$ 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 $\sum_{i=1}^n x_i$, 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. ◀

▶ Remark. From the proof of Theorem 8, 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 that the approximation ratio for welfare is also constant.

▶ **Theorem 9.** *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 auction becomes fully competitive.*

Proof. For social welfare we have, similarly to Theorem 8, that

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

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 10.** *Let* M *be any truthful mechanism. 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 $i = 3, 4, \ldots, n$. 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 $\mathbf{v}, \mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \ldots, \mathbf{v}^{(k)}$ 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 $\frac{m}{(m/2)+1} = 2 - \frac{4}{m+2}$.

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. ◀

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 11.** *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 included in the appendix. 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 12.** *There is no in-range, non-wasteful and truthful mechanism that always outputs an envy-free pricing scheme, even when the budgets are known.*

The full proof is in the appendix. 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 and example where the mechanism must leave some items unallocated in order to satisfy truthfulness.

5 Conclusion

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 appendix, 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.

References

1 E. Anderson and D. Simester. Price stickiness and customer antagonism. *Available at SSRN 1273647*, 2008.

2 Lawrence 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. *arXiv preprint arXiv:1511.06954*, 2015.

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 Y. Cai, C. Daskalakis, and M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *FOCS*, pages 130–139, 2012.

10 Y. Cai, C. Daskalakis, and M. Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *SODA*, pages 578–595, 2013.

11 N. Chen, X. Deng, and J. Zhang. How profitable are strategic behaviors in a market? In *ESA*, pages 106–118. Springer, 2011.

12 M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *FOCS*, pages 35–44, 2008.

13 V. Cohen-Addad, A. Eden, M. Feldman, and A. Fiat. The invisible hand of dynamic market pricing. In *ACM EC*, pages 383–400, 2016.

14 R. Colini-Baldeschi, S. Leonardi, P. Sankowski, and Q. Zhang. Revenue maximizing envy-free fixed-price auctions with budgets. In *WINE*, pages 233–246. 2014.

15 S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *GEB*, 74(2):486–503, 2012.

16 S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *ICALP*, pages 392–404. Springer, 2014.

17 S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *ACM EC*, pages 346–351, 2007.

18 S. Dobzinski and N. Nisan. Multi-unit auctions: beyond roberts. *JET*, 156:14–44, 2015.

19 P. Farris, N. Bendle, P. Pfeifer, and D. Reibstein. *Marketing metrics: The definitive guide to measuring marketing performance*. Pearson Education, 2010.

20 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.

21 M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *SODA*, pages 123–135, 2015.

22 F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.

23 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.

24 J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *ACM EC*, pages 225–234. ACM, 2009.

25 J. Hartline and Q. Yan. Envy, truth, and profit. In *ACM EC*, pages 243–252, 2011.

26 J. D. Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

27 A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.

28 E. Markakis and O. Telelis. Envy-free revenue approximation for asymmetric buyers with budgets. In *SAGT*, 2016.

29 R. Mehta and M. Sohoni. Exchange markets: Strategy meets supply-awareness. In *WINE*, pages 361–362. Springer, 2013.

30 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.

31 R. B. Myerson. Optimal auction design. *MOR*, 6(1):58–73, 1981.

32 N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge Univ. Press, (editors) 2007.

33 M. Petersen. Information: Hard and soft. Working paper, 7 2004.

34 Shreyas Sekar. Posted pricing sans discrimination. *arXiv preprint arXiv:1609.06844*, 2016.

35 L. Walras. Elements d'economie 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.

A Impossibility Results

▶ Theorem 11. 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. ◀

Next, we will provide the proof for Theorem 12, which we restate below.

▶ Theorem 12. There is no in-range, non-wasteful and truthful mechanism that always outputs an envy-free pricing scheme, even when the budgets are known.

Before we prove the theorem, we will prove a necessary condition for the price outputted by mechanisms that are in-range, non-wasteful and truthful.

▶ **Lemma 13.** *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 $p_{\min} = \min\{p \in \mathbb{O} : p \text{ is envy-free on } \mathbf{v}\}$ and γ is the distance between two consecutive elements of $\mathbb{O}$.*

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 $p_1 = p_{\min}$. In that case, 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 the same amount of 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 $J := \{j : v_j \geq p_1 : j \text{ is allocated a non-zero number of units}\}$. 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 13 can be strengthened so that M can only output the minimum envy-free price, whenever it exists. Using Lemma 13, we can now prove the theorem.

Proof. (of Theorem 12) Assume by contradiction that such an in-range, non-wasteful and truthful mechanism M exists. We will consider three different instances[4] with $n = 2$ and $m = 3$, denoted (v_1, v_2) where v_1 denotes the 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 13, 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 11, 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 11). By Lemma 13 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 11, $v_2' = 2.5$ 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

[4] 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.

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 13, 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 $\bar{x}_1 = 2$. ◀

B Performance with respect to the continuous domain

In this section, we will prove that if the grid is fine enough, then the loss in either objective compared to the continuous domain is zero. For each player i, let $v_i = u_i/w_i$, $B_i = c_i/d_i$, where u_i, w_i, c_i, d_i are integers, each represented with t bits. Also recall the precision of the output (i.e. of the price computed) is s bits, where $s \geq 2t + 2\log m$.

For both revenue and welfare, the optimal solution can be found in a set of *candidate prices*:

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

These prices are either equal to some valuation or have the property that some buyer could exhaust its budget by purchasing all the units it can afford.

▶ **Lemma 14.** *For both the revenue and social welfare objectives, there is an optimal envy-free price $p \in \mathcal{P}$.*

Proof. Let p' be a welfare maximizing envy-free price and let $\mathbf{x}$ be the corresponding allocation. If $p' \in \mathcal{P}$ then the statement of the lemma is trivially true. Otherwise, $p' \neq \mathcal{P}$. Then, the price p' can be increased until either a budget B_i is exhausted or the price becomes equal to some valuation v_i. Until that happens, the demand sets of all buyers remain constant and so the same allocation $\mathbf{x}$ can be supported at some price $p \in \mathcal{P}$. Since the social welfare only depends on the allocation and not the price, the lemma follows. The proof for revenue follows from the observation that increasing the price is beneficial for the seller as long as it continues to sell the same number of units. The discontinuities only happen at points where the price matches the valuation at some buyer (and so increasing the price above that value can result in losing the buyer) or when the number of units decreases because a buyer can no longer afford to purchase as many units. ◀

▶ Theorem 6. When the valuation and budget of each buyer are rational numbers, where both the numerator and denominator are integers specified with t bits of precision, and the price is specified with $s \geq 4t + 4\log m$ bits, 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 $O(t+s)$.

Proof. When $s \geq 4t + 4\log m$, the entire set $\mathcal{P}$ is a possible output, and so the minimum envy-free price computed by the mechanism operating on the grid is either the envy-free price p^* that minimizes welfare among the prices in the set $\mathcal{P}$, or a price lower than p^*, which does not decrease welfare since welfare is non-increasing in the price. In both cases, the loss in welfare from using the grid is zero.

For the revenue objective, let p_S be the infimum of the set of envy-free prices on the real line and p_M the price output by the mechanism on the grid with precision s. If p_S is envy-free, then it is equal to some valuation (otherwise, by decreasing the price by an arbitrarily small amount $\epsilon > 0$, the set of buyers with valuations higher than or equal to the price remains the same, there are no semi-hungry buyers, and nobody's allocation can increase compared to p_S when ϵ is small enough, thus we would still obtain an envy-free price). Then p_S is part of the output grid and there is no loss in revenue.

Otherwise, consider the case when p_S is not envy-free. Then p_S cannot be equal to a valuation and there is overdemand at p_S, i.e. there exists buyer i that can afford an extra unit at p_S compared to what it could afford at $p_S + \epsilon$ for small enough $\epsilon > 0$. That is, $p_S = B_i/k$ for some buyer i and number of units $k \in [m]$. Let $\Delta = \min_{q,r\in\mathcal{P}} |q - r|$. Then $\Delta \geq 1/2^{2t+2\log m}$ by the choice of precision of the elements in $\mathcal{P}$.

Moreover, $p_M > p_S = B_i/k$ and $s \geq 4t + 4\log m$, thus $(p_S, p_M] \cap \mathcal{P} = \emptyset$. It follows there are no semi-hungry buyers at any price $p \in (p_S, p_M]$ and the set of hungry buyers is the same at all the prices in $(p_S, p_M]$. Next, we argue the demands of the hungry buyers are the same at any price in $(p_S, p_M]$. If this were not the case, there should be a price $p \in (p_S, p_M]$ at which the demand changes. I.e., there exists buyer i for which $\lfloor B_i/p' \rfloor < \lfloor B_i/p \rfloor$ for all $p' \in (p, p_M)$. This can only happen if there exists $j \in [m]$ such that $B_i/j = p$, i.e. the buyer exhausts its budget at p. By choice of the interval, this is a contradiction since $B_i/j \in \mathcal{P}$ but there is no element of $\mathcal{P}$ in $(p_S, p_M]$.

Since the set of hungry buyers is the same at p_M and p_S and the limit of the demand as the price approaches p_S from above is equal to the demand at p_M, the revenue at p_M is at least as high as the limit of the revenue at p_S. Thus there is no loss in revenue either, which completes the proof.

The runtime follows since the mechanism finds the minimum envy-free price using binary search. ◀

C Monotone Auctions

In the main text, we proved the approximation ratio guarantees of Mechanism All-or-Nothing, as a function of the market share. In this section, we will examine the case of common budgets and the more general class of monotone auctionss:

- The budgets are *common* when $B_i = B$ for all buyers $i \in N$.
- 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 for those cases, Mechanism All-or-Nothing is optimal 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 auctionin which a single buyer can afford to buy all the items at a very high price.

▶ **Definition 15.** 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.[5]

▶ **Theorem 16.** *When the auction is monotone, the approximation ratio of Mechanism* All-or-Nothing *is*

- *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 10 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 8, 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

[5] 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 $\text{EFO}^{(2)}$ benchmark for revenue, see [26], Chapter 6.

$x_\ell \cdot v_\ell \leq v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor$ 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:

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

- $p^* > p_{min}$. In that case, the argument is exactly the same as in Case 2 of the proof of Theorem 8, 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 17.** *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.[6] Finally, for a given $\mathcal{B} > 1$ let 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 $v < v_1' < p^*$ and 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 devation, contradicting the truthfulness of M. ◀

[6] 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.