Clearing Multi-Unit Markets with Simple Pricing

Ali Makhdoumi
Fuqua School of Business, Duke University, ali.makhdoumi@duke.edu

Saša Pekeč
Fuqua School of Business, Duke University, pekec@duke.edu

Xinchang Wang
Fuqua School of Business, Duke University, xinchang.wang@duke.edu

We study multi-unit markets in which a seller allocates indivisible units to buyers with heterogeneous valuations through one anonymous per-unit price. When valuations are submodular, a Walrasian equilibrium gives every buyer a utility-maximizing quantity and clears the market. Yet an arbitrarily small departure from submodularity can preclude equilibrium existence. We therefore introduce the τ-approximate Walrasian equilibrium, which preserves one price and market clearing while requiring each buyer's realized utility to be at least the maximum utility under τ-scaled valuations. Every such equilibrium attains at least a τ fraction of optimal social welfare. We study the largest supportable relaxation level τ^*, establish a tight lower bound in terms of relative distance from submodularity, and show that exact computation is NP-hard. Our main algorithmic result is an exact-clearing fully polynomial-time approximation scheme (FPTAS) for arbitrary monotone valuations.

1. Introduction

A central challenge in many markets is to allocate indivisible units among agents with quantity-dependent values. Pricing is a primary instrument for implementing an allocation, and in full generality an efficient market-clearing allocation can always be supported by prices, provided the pricing scheme is complex enough, potentially requiring quantity-specific and personalized prices. While theoretically possible, such schemes can be practically unimplementable due to the communication burden and the difficulty of administering them, and in settings such as regulated tariffs or public procurement, where equal treatment is required, they may not be permitted at all. Consequently, even though it may be suboptimal, a single anonymous per-unit price is widely used in practice, including in commodity spot markets, wholesale electricity, and treasury auctions. This

paper studies the power and limits of simple pricing: how to compute the best anonymous per-unit price for a given market, and how to measure how far that market-clearing outcome induced by that price falls short of the efficient outcome.

The question of when a single price supports an efficient market-clearing outcome has a classical answer: *Walrasian Equilibrium (WE)* formalizes a market equilibrium that can be implemented through a simple pricing mechanism. In particular, a WE consists of a single per-unit price and an integral market-clearing allocation such that each buyer receives a quantity in her *demand correspondence*, defined as the set of quantities that maximize her utility at that price. When buyers have submodular valuations, a WE is guaranteed to exist, and every Walrasian allocation maximizes social welfare. When valuations are not submodular, however, a WE may fail to exist: there may be no single price at which every buyer receives a quantity in her demand correspondence while the aggregate demand clears the market. (Such valuations arise naturally when units are inputs to a job rather than consumption goods. For example, a cloud platform allocating a fixed pool of GPUs faces buyers whose training jobs deliver little value until they receive enough GPUs to meet a device-memory requirement. The challenge is to use a single price to direct limited GPU capacity toward high-value uses while fully utilizing the available capacity, even when buyers' preferred quantities do not align perfectly with the available supply.)

This raises a natural question: how can we relax the notion of WE so that a market-clearing outcome supported by a single price exists for a broader class of valuations, while retaining a meaningful welfare guarantee? Any relaxation of WE must give up one of its three defining conditions: a single anonymous price, market clearing, or buyer optimality, i.e., each buyer receiving a quantity in her demand correspondence. Relaxing only market clearing does not work (as we illustrate in Section 3.2): even for a market with an arbitrarily small deviation from submodularity, the best allocation drawn from buyers' demand correspondences at a single price can achieve an arbitrarily small fraction of optimal welfare. We instead preserve both a single price and exact market clearing, and relax buyer optimality.

Formally, we introduce a relaxation parameter $\tau \in (0, 1]$ and define the τ-relaxed demand correspondence as the set of quantities whose utility under the original valuation is at least the utility of every other quantity under the τ-scaled valuation. As we show in Section 3, the relaxation follows the structure of the welfare problem: the dual constraints are precisely the demand correspondence conditions, so scaling the valuations in those constraints by τ is natural, and the corresponding primal scales its objective by τ. Building on this, a τ-Walrasian equilibrium (τ-WE) is defined as a single price and a market-clearing allocation in which each buyer receives a quantity in her τ-relaxed demand correspondence. When $\tau = 1$, this notion coincides with an exact WE. As τ decreases, the relaxed demand correspondences expand, which equivalently means that the buyers' optimality further relaxes. Thus, despite the name, a τ-WE is not an equilibrium in the usual sense: a buyer's allocation is a τ-approximate best response to the price. Finally, by construction, every τ-WE allocation achieves at least a τ fraction of optimal welfare

In Section 4, we address the problem of finding the tightest relaxation that a given market admits and let τ^* denote the largest value of τ for which a τ-WE exists in the market. To quantify how departures from submodularity affect τ^*, we introduce the market's *relative distance from submodularity*, defined as the largest relative gap, across buyers and quantities, between a valuation and its concave envelope. We then establish a tight lower bound on τ^* as a function of this distance measure. However, the bound could be conservative for a particular market, motivating the market-specific computation of τ^*. We show that computing τ^* is NP-hard via a reduction from the Subset Sum problem.

This computational difficulty motivates our main algorithmic result (Section 5): a fully polynomial-time approximation scheme (FPTAS) that finds the tightest relaxation for arbitrary monotone valuations. For any precision parameter, the algorithm returns a $\widehat{\tau}$-WE with a supporting price and a market-clearing allocation, where $\widehat{\tau}$ approximates τ^* from below within the desired multiplicative factor. The main algorithmic challenge is to jointly search over the relaxation level and price, satisfy the integral market-clearing constraint, and retain fully polynomial running time.

Our approach has two main components. First, we develop a bicriteria FPTAS that truncates and rounds valuations and quantities to construct rounded relaxed demand correspondences (rounded RDCs), each represented by a small number of integer intervals. This lets us replace the continuous search over prices with a finite set of candidate prices and use a Multiple-Choice Subset Sum (MCSS) approximation algorithm to find an allocation that nearly clears the market at a relaxation level close to τ^*. We then turn this near-clearing solution into one that clears the market exactly. When the supply is small relative to the number of buyers, we calibrate the clearing tolerance so that integrality implies exact clearing. When the supply is large, we fill any remaining shortfall using a multiplicative completion step that scales both the relaxation level and the price.

Our results suggest that the failure of a WE need not imply the failure of simple pricing. With a controlled relaxation of buyer optimality, a single price can still support exact market clearing while providing a provable welfare guarantee. The Tightest Relaxation problem captures the market-specific limits of simple pricing, distinguishing markets in which only a modest relaxation is needed from those in which substantial relaxation is unavoidable. Moreover, when the simplicity of a single price is to be retained, our FPTAS provides an efficient way to implement a near-best single-price outcome.

The rest of the paper proceeds as follows. We next describe the related literature. Section 2 introduces the model and Walrasian equilibrium. Section 3 develops the τ-WE framework and derives its welfare guarantee. Section 4 formulates the Tightest Relaxation problem and analyzes its computational complexity and structure. Section 5 presents an exact-clearing FPTAS for that problem, built on a bicriteria FPTAS subroutine. Section 6 concludes the paper. Omitted proofs and additional analysis are deferred to the appendix.

1.1. Related Work

We briefly review several strands of literature in economics, computer science, and operations research that are most relevant to our paper. A broad literature evaluates simple pricing rules relative to richer pricing policies. In revenue management, static prices can be asymptotically optimal

or attain constant-factor guarantees relative to optimal dynamic policies in inventory and reusable-resource problems (e.g., Gallego and van Ryzin 1994, Bitran and Caldentey 2003, Talluri and van Ryzin 2004, Besbes et al. 2022). Research on price discrimination and fairness constraints quantifies the gains from differentiated prices and their revenue and welfare consequences (Elmachtoub et al. 2021, Cohen et al. 2022). In algorithmic market design, posted-price mechanisms provide approximation guarantees for revenue or welfare, often under online arrivals (e.g., Balcan et al. 2008, Feldman et al. 2014, Dutting et al. 2020, Dütting et al. 2020, Correa et al. 2022), while static prices can face capacity-related limits in online Fisher markets (Jalota and Ye 2025). In online multi-unit markets with sequentially arriving buyers, Ezra et al. (2020) compare uniform and nonuniform prices across valuation classes and show that, for general valuations, even nonuniform prices can guarantee only a $1/m$ fraction of optimal welfare. Our work complements this literature by studying the effectiveness of simple pricing in a static multi-unit market with indivisible supply and arbitrary monotone valuations.

Our work is also closely related to the literature on Walrasian equilibrium existence and computation. Classical results guarantee the existence of a WE when buyers' valuations satisfy gross substitutes (GS) (Kelso Jr and Crawford 1982, Gul and Stacchetti 1999). Multi-unit markets are a special case of combinatorial markets in which buyers value only the number of identical units they receive; in this setting, GS is equivalent to submodularity (see, e.g., Paes Leme 2017). Under GS, a WE can also be computed in polynomial time in standard oracle models (Paes Leme and Wong 2020). Research beyond GS identifies broader substitute, near-substitute, and discrete-convex conditions that preserve equilibrium existence (e.g., Sun and Yang 2006, Baldwin and Klemperer 2019, Candogan et al. 2015, Candogan and Pekeč 2018, Baldwin et al. 2023, Gul et al. 2019). Related work studies incentive compatibility, demand reduction, efficiency loss, and learning in multi-unit auctions (e.g., Ausubel 2004, Ausubel et al. 2014, Markakis and Telelis 2015, Birmpas et al. 2019, Brânzei et al. 2023a, Galgana and Golrezaei 2025). We take valuations as given and focus on the optimization questions that arise when equilibrium existence is no longer guaranteed.

When valuations do not satisfy GS, a WE may fail to exist, and prior work develops approximate or alternative equilibria along several dimensions. Large-market approaches make indivisibilities negligible relative to market size through aggregation, continuum arguments, or randomized implementation (Starr 1969, Mas-Colell 1977, Azevedo et al. 2013, Akbarpour and Nikzad 2020). For finite markets, Nguyen and Vohra (2024) parameterize departures from substitutability to bound good-by-good excess demand, whereas the market's *relative distance from submodularity* γ bounds the relaxation of buyer optimality sufficient for market clearing. A different relaxation preserves buyer optimality at an anonymous price while allowing units to remain unsold in multi-unit auctions with budgets (Brânzei et al. 2023b). Nearby-economy and convexification results compare the original market with a perturbed or convexified economy in which an exact equilibrium exists, including models with small budget perturbations (Postlewaite and Schmeidler 1981, Budish 2011, Hübner 2025). Other work permits bundling or more flexible prices and transfers (Feldman et al. 2013, 2022, O'Neill et al. 2005, Milgrom and Watt 2026). Our approach preserves the original market primitives, market clearing, and a single anonymous per-unit price, and relaxes only buyer optimality.

Several precedents also relax buyer optimality, but differ in how the relaxation is structured. Conditional equilibrium structures the relaxation by the direction of deviations. It retains individual rationality and rules out deviations that add items, but permits profitable deviations that drop allocated items without placing a quantitative restriction on their gains beyond individual rationality (Fu et al. 2012). Related relaxations in single-minded combinatorial auctions structure the relaxation through the set of bidders required to satisfy exact buyer optimality. They maximize or guarantee the fraction of bidders whose buyer-optimality condition holds, while the remaining bidders need not satisfy such a condition (Chen et al. 2004, Huang et al. 2005, Chen and Rudra 2008). Yquilibrium takes an aggregate approach by selecting allocations and anonymous linear prices that maximize a potential based on the differences between realized utilities and the indirect utilities implied by prices and endowments. Because this potential sums these differences

across buyers, it does not impose a common buyer-specific bound (Goeree 2023). Our framework instead uses a common, multiplicative, and endogenous relaxation. A single level serves both as the optimality condition imposed on every buyer and as the fraction of optimal social welfare the outcome guarantees. The largest supportable level τ^* is a property of the market: it equals one if and only if an exact WE exists, and it measures how far a given market falls short of supporting one. Our Exact-Clearing FPTAS approximates this measure to any prescribed accuracy for arbitrary monotone valuations, so no structural assumption is required and the measure remains efficiently computable precisely where complementarities preclude an exact WE.

Finally, our relaxation is conceptually related to approximate incentive compatibility and strategy-proofness (e.g., Hartline and Lucier 2015, Balseiro et al. 2024, Lubin and Parkes 2012, Azevedo and Budish 2019), as well as to approximate fairness notions such as EF1 and EFX (e.g., Lipton et al. 2004, Caragiannis et al. 2019, Amanatidis et al. 2023, Aziz et al. 2022). These papers relax incentive or interpersonal fairness requirements, whereas we relax buyer optimality under a single market-clearing price.

2. Model

We consider a market with $m \in \mathbb{Z}_{\geq 2}$ identical units of an indivisible good and $n \in \mathbb{Z}_{\geq 2}$ buyers. Let $[n] \coloneqq \{1, \ldots, n\}$ and $[m] \coloneqq \{1, \ldots, m\}$ denote the set of buyers and units of the good, respectively. Buyer i's valuation is a function $f_i : \{0, \ldots, m\} \to \mathbb{R}_{\geq 0}$, where $f_i(k)$ is her value for k units. Each valuation is normalized and monotone, i.e., $f_i(0) = 0$ and $f_i(k) \leq f_i(k')$ whenever $k \leq k'$. Specifically, a valuation f_i is *submodular* if $f_i(k) - f_i(k-1)$ is weakly decreasing in k, that is, the valuation exhibits diminishing marginal values. It is *single-minded* if $f_i(k) = v_i \mathbb{1}\{k \geq q_i\}$ for some quantity $q_i \in [m]$ and value $v_i > 0$. Throughout, we assume that at least one buyer assigns strictly positive value to some positive quantity, i.e., $\max_{i,k} f_i(k) > 0$.

2.1. Welfare Maximization and Walrasian Equilibirum

An integral allocation is a vector $\boldsymbol{x} \in \{0, \ldots, m\}^n$ satisfying $\sum_{i=1}^n x_i \leq m$. An integral allocation is *market-clearing* if $\sum_{i=1}^n x_i = m$. The objective is to maximize the social welfare, formalized as

follows:

$$\max_{\boldsymbol{x}} \sum_{i=1}^{n} f_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^{n} x_i \le m, \quad x_i \in \{0,1,\cdots,m\}\ \forall i \in [n]. \tag{SWM}$$

Monotonicity of valuation functions implies that (SWM) has a market-clearing optimal solution.

We are interested in implementing the social welfare-maximizing (integral) allocation with a single price and also guarantee that each buyer does not have a deviation from the allocation given to her to further maximize her utility. We next formalize this by introducing *Walrasian Equilibrium (WE)*, a benchmark solution concept.

Under a single anonymous per-unit price $p \ge 0$, buyer i's utility from k units is

$$U_i(p,k) \coloneqq f_i(k) - pk.$$

Next, we define the demand correspondence, which characterizes each buyer's utility-maximizing quantities and Walrasian equilibrium.

Definition 1 (Demand correspondence). Buyer i's demand correspondence at price p is

$$D_i(p) \coloneqq \operatorname*{arg\,max}_{k \in \{0,\ldots,m\}} U_i(p,k). \tag{1}$$

That is, the set of quantities maximizing her utility given a price p.

Definition 2 (Walrasian equilibrium). A price-allocation pair $(p, \boldsymbol{x})$ is a *Walrasian equilibrium* (WE) if

(a) $x_i \in D_i(p)$ for every $i \in [n]$ (*buyer optimality*); and

(b) $\sum_{i=1}^{n} x_i = m$ (*market clearing*).

In other words, a WE pairs a single price with a market-clearing allocation in which every buyer receives a utility-maximizing quantity at that price. When every valuation is submodular, a WE exists (see, e.g., Kelso Jr and Crawford 1982, Dress and Terhalle 1995, Gul and Stacchetti 1999).

2.2. Primal-Dual Interpretation and Welfare Guarantee of WE

Here, we establish that WE (if it exists) implements the social welfare-maximizing solution, formalized next.

PROPOSITION 1. *Every allocation* $\mathbf{x}$ *in a WE maximizes social welfare.*

Proof of Proposition 1: We prove this by using LP Relaxation and Duality. The integer program (SWM) has the following LP relaxation and dual. The primal replaces each buyer's integral quantity by fractional weights $x_i(k)$ on positive quantities. The dual uses a nonnegative per-unit price p and a nonnegative utility variable u_i for each buyer.

$$\begin{aligned} \max_{\boldsymbol{x}} \quad & \sum_{i=1}^{n}\sum_{k=1}^{m} x_i(k)\cdot f_i(k) \\ \text{s.t.} \quad & \sum_{i=1}^{n}\sum_{k=1}^{m} x_i(k)\cdot k \le m; \\ & \sum_{k=1}^{m} x_i(k) \le 1,\ \forall\ i\in[n]; \\ & x_i(k)\ge 0,\ \forall\ i\in[n],\ k\in[m]. \end{aligned} \qquad \text{(LP)}$$

$$\begin{aligned} \min_{p,\boldsymbol{u}} \quad & m\cdot p + \sum_{i=1}^{n} u_i \\ \text{s.t.} \quad & u_i \ge f_i(k) - k\cdot p,\ \forall\ i\in[n],\ k\in[m]; \\ & p\ge 0,\ u_i\ge 0,\ \forall\ i\in[n]. \end{aligned} \qquad \text{(Dual)}$$

Every feasible integral allocation of (SWM) induces a feasible solution of (LP) with the same welfare, so (LP) is a relaxation of (SWM). Dual feasibility requires $u_i \ge f_i(k) - kp$ for all $k \in \{0,\dots,m\}$. At a fixed price p, the smallest dual-feasible utilities are the indirect utilities

$$u_i = \max_{k\in\{0,\dots,m\}} \{f_i(k) - kp\}.$$

If $(p,\boldsymbol{x})$ is a WE, these utilities equal the realized utilities $u_i = f_i(x_i) - x_i p$, so $(p,\boldsymbol{u})$ is dual-feasible, and market clearing gives

$$mp + \sum_i u_i = \sum_i f_i(x_i).$$

Weak duality then yields

$$\sum_i f_i(x_i) \le \mathrm{OPT(SWM)} \le \mathrm{OPT(LP)} \le mp + \sum_i u_i = \sum_i f_i(x_i).$$

Thus, every WE allocation maximizes social welfare. □

It is worth noting that a WE may fail to exist, which raises the question of what constitutes the "right" relaxation while still achieving a "reasonable" approximation to the optimal social welfare. We formalize these notions in the next section by introducing an approximate notion of Walrasian equilibrium that relaxes buyer optimality while preserving exact market clearing and a single price. Motivated by the proof above, we also derive a welfare guarantee using a scaled LP duality argument.

Before that, since WE always exists for submodular valuations, it is useful to introduce a notion of distance from submodularity to capture the extent to which the notion of WE must be relaxed.

2.3. Relative Distance from Submodularity

For each buyer i and $k \in \{0, \dots, m\}$, define

$$\overline{f}_i(k) \coloneqq \max \left\{ \sum_{q=0}^{m} \lambda_q f_i(q) : \sum_{q=0}^{m} \lambda_q = 1, \quad \sum_{q=0}^{m} q\lambda_q = k, \quad \lambda_q \geq 0 \right\}. \tag{2}$$

We call $\overline{f}_i$ the *concave envelope* of f_i. Equivalently, $\overline{f}_i$ is the pointwise smallest submodular valuation that weakly dominates f_i. The following lemma shows some basic properties:

LEMMA 1 **(Concave envelope equivalency)**. *If $f : \{0, \dots, m\} \to \mathbb{R}_{\geq 0}$ is normalized and monotone, then its concave envelope $\overline{f}$ is normalized, nonnegative, monotone, and submodular. Moreover, $\overline{f} = f$ if and only if f is submodular.*

For each buyer i and each quantity k with $\overline{f}_i(k) > 0$, the *relative gap* at k is $1 - f_i(k)/\overline{f}_i(k)$. Buyer i's *relative distance from submodularity* is the largest such gap,

$$\gamma_i \coloneqq \max_{k:\, \overline{f}_i(k) > 0} \left(1 - \frac{f_i(k)}{\overline{f}_i(k)} \right),$$

with the convention that an empty maximum is zero (See Figure 1 for an illustration). The market's *relative distance from submodularity* is

$$\gamma \coloneqq \max_{i \in [n]} \gamma_i. \tag{3}$$

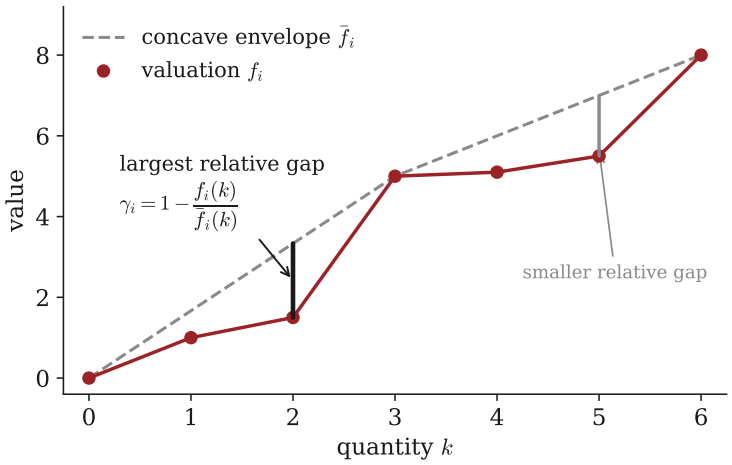


Figure 1 Buyer i's relative distance from submodularity. The largest relative gap between f_i and its concave envelope $\overline{f}_i$ is γ_i.

Since $0 \le f_i(k) \le \overline{f}_i(k)$ for every i and k, every relative gap lies in $[0,1]$ and therefore $\gamma \in [0,1]$. By definition of γ,

$$f_i(k) \ge (1-\gamma)\overline{f}_i(k), \qquad i \in [n], \quad k \in \{0,\dots,m\}. \tag{4}$$

The equality $\gamma = 0$ holds exactly when every valuation is submodular.

3. The τ-Approximate Walrasian Equilibrium

A Walrasian equilibrium requires three conditions: buyer optimality, market clearing, and a single price (Definition 2). In this section, we introduce the τ-approximate Walrasian equilibrium, which preserves the latter two requirements and relaxes buyer optimality through a scaled-valuation benchmark. The relaxed equilibrium expands the set of quantities each buyer accepts at a given price. In what follows, we first formally define the τ-approximate Walrasian equilibrium framework. Then, we use a two-buyer market to show why classical WE with a single price fails and how our framework restores market clearing. We then develop a primal-dual representation of the framework and derive its welfare guarantee.

3.1. Definitions and Basic Properties

Fix a relaxation parameter $\tau \in (0,1]$. For buyer i, define the utility benchmark with scaled valuation

$$U_i^\tau(p,k) := \tau f_i(k) - pk.$$

We then define τ-relaxed demand correspondence through the above benchmark:

Definition 3 (τ-Relaxed Demand Correspondence). For any relaxation parameter $\tau \in (0,1]$ and price $p \geq 0$, the *τ-relaxed demand correspondence* (τ-RDC) of buyer i is

$$D_i^\tau(p) \coloneqq \{k \in \{0,1,\ldots,m\} : U_i(p,k) \geq U_i^\tau(p,k'),\ \forall\, k' \in \{0,1,\ldots,m\}\}. \tag{5}$$

A quantity k lies in $D_i^\tau(p)$ when the realized utility $U_i(p,k)$ is at least the utility of every alternative under the τ-scaled valuation. At $\tau = 1$, $D_i^\tau(p)$ recovers the classical demand correspondence (Definition 1), i.e., $D_i^1(p) = D_i(p)$. For $0 < \tau' \leq \tau \leq 1$, nonnegativity of valuations gives $U_i^{\tau'}(p,k) \leq U_i^\tau(p,k)$ for every k, and hence

$$D_i^\tau(p) \subseteq D_i^{\tau'}(p).$$

In other words, τ indexes the degree of relaxation: smaller values of τ expand the set of acceptable quantities for the buyers. For a fixed pair (τ, p), we call an allocation $\boldsymbol{x}$ *demand-feasible* if $x_i \in D_i^\tau(p)$ for every $i \in [n]$.

The above τ-RDC then induces an approximate notion of Walrasian equilibrium:

Definition 4 (τ-Approximate Walrasian Equilibrium). A price-allocation pair $(p, \boldsymbol{x})$ forms a *τ-approximate Walrasian equilibrium* (τ-WE) if

(a) $x_i \in D_i^\tau(p)$ for every $i \in [n]$ (relaxed buyer optimality), and

(b) $\sum_{i=1}^n x_i = m$ (market clearing).

The next example shows that even an arbitrarily small violation of submodularity can preclude an exact WE, while a τ-WE with only a correspondingly small relaxation still clears the market.

3.2. An Illustrative Example: Equilibrium Failure and Recovery

Consider two buyers and $m \geq 2$ identical units. For a parameter $\xi \in (0,1)$, let

$$f_1(k) = 2 \quad (k \geq 1), \qquad f_2(k) = 1 + (k-1)(1+\xi) \quad (k \geq 1). \tag{Market 1}$$

Buyer 1 values the first unit at 2 and has zero marginal value for additional units, so her valuation is unit-demand. Buyer 2 values the first unit at 1 and each additional unit at $1+\xi$. Her marginal

values are nondecreasing, so her valuation is supermodular. The only violation of submodularity is the increase of ξ between the first and second marginal values. Since $2 > 1 + \xi$, the welfare-maximizing allocation gives one unit to buyer 1 and the remaining $m-1$ units to buyer 2, yielding

$$\mathrm{OPT} = m + 1 + (m-2)\xi.$$

Failure of exact WE. Let

$$p^* := \frac{f_2(m)}{m} = 1 + \xi - \frac{\xi}{m}.$$

This is buyer 2's break-even price between accepting zero and all m units. At a fixed price p, her marginal utilities are $1-p$ for the first unit and $1+\xi-p$ for each additional unit. These marginal utilities are nondecreasing, so $U_2(p,\cdot)$ is convex on the quantity grid and attains its maximum at 0 or m. Comparing the endpoint utilities gives

$$D_2(p) = \begin{cases} \{m\}, & p < p^*, \\ \{0, m\}, & p = p^*, \\ \{0\}, & p > p^*. \end{cases}$$

Since $p^* \in (1,2)$, buyer 1 demands at least one unit for prices below p^*, exactly one unit at p^*, and at most one unit for prices above p^*. Aggregate demand is therefore at least $m+1$ below p^*, can equal only 1 or $m+1$ at p^*, and is at most 1 above p^*. It never equals the supply m, so no Walrasian equilibrium exists. Even if the seller may leave units unsold, every allocation that assigns each buyer a quantity from her demand correspondence and uses at most m units has total quantity at most 1 and welfare at most 2, only an $O(1/m)$ fraction of OPT.

Recovery under relaxed buyer optimality. Set

$$\overline{\tau} := \frac{m + (m-2)\xi}{m + (m-1)\xi} = 1 - \frac{\xi}{m + (m-1)\xi}.$$

At price $p = 1$, consider the allocation $(1, m-1)$. For buyer 1, we have that

$$\max_{k \in \{0,\ldots,m\}} U_1^{\overline{\tau}}(1,k) = \max\{0, 2\overline{\tau} - 1\} \le 1 = U_1(1,1).$$

For buyer 2, $\overline{\tau}(1+\xi) > 1$, so the term $U_2^{\overline{\tau}}(1,k)$ is increasing in k and attains a maximum at $k = m$. The definition of $\overline{\tau}$ gives

$$\max_{k\in\{0,\dots,m\}} U_2^{\overline{\tau}}(1,k) = U_2^{\overline{\tau}}(1,m) = (m-2)\xi = U_2(1,m-1).$$

Hence $1 \in D_1^{\overline{\tau}}(1)$ and $m-1 \in D_2^{\overline{\tau}}(1)$. Since these quantities sum to m, the pair $(1,(1,m-1))$ is a $\overline{\tau}$-WE. Its allocation coincides with the welfare-maximizing allocation identified above.

3.3. Primal-Dual Interpretation and Welfare Guarantee of Approximate WE

In the previous subsection, we illustrated that a τ-WE can restore market clearing when an exact WE fails. We now ask what welfare guarantee an arbitrary τ-WE provides, the analogue of Proposition 1 for approximate WE:

PROPOSITION 2. *Every allocation $\boldsymbol{x}$ in a τ-WE satisfies*

$$\sum_{i=1}^{n} f_i(x_i) \geq \tau \cdot \mathrm{OPT(SWM)}.$$

Proof of Proposition 2: In the proof of Proposition 1, we used duality to show the welfare of an exact WE is optimal. Scaling the valuation coefficients in that LP by τ produces the corresponding primal-dual pair:

$$\begin{aligned}
\max_{\boldsymbol{x}} \quad & \tau \sum_{i=1}^{n}\sum_{k=1}^{m} x_i(k) f_i(k) \\
\text{s.t.} \quad & \sum_{i=1}^{n}\sum_{k=1}^{m} x_i(k)k \leq m; \\
& \sum_{k=1}^{m} x_i(k) \leq 1, \ \forall\, i \in [n]; \\
& x_i(k) \geq 0, \ \forall\, i \in [n], \ k \in [m].
\end{aligned} \qquad (\tau\text{-LP})$$

$$\begin{aligned}
\min_{p,\boldsymbol{u}} \quad & mp + \sum_{i=1}^{n} u_i \\
\text{s.t.} \quad & u_i \geq \tau f_i(k) - kp, \ \forall\, i \in [n], \ k \in [m]; \\
& p \geq 0, \ u_i \geq 0, \ \forall\, i \in [n].
\end{aligned} \qquad (\tau\text{-Dual})$$

This scaling leaves the feasible allocations, dual variables, and dual objective unchanged compared to (LP)-(Dual). Every valuation function in the primal objective and in the dual constraints is scaled by τ. For an integral quantity x_i and realized utility $u_i = f_i(x_i) - px_i$, the dual inequalities indexed by $k \in [m]$ become

$$f_i(x_i) - px_i \geq \tau f_i(k) - pk \quad \text{for every } k \in [m].$$

The remaining constraint $u_i \geq 0$ is the same comparison at $k = 0$ because $f_i(0) = 0$. Together, these conditions are equivalent to $x_i \in D_i^\tau(p)$. The dual, therefore, characterizes relaxed buyer optimality at a single price.

Now, let $(p, \boldsymbol{x})$ be a τ-WE, and write $u_i = f_i(x_i) - x_i p$ for each $i \in [n]$. Since $x_i \in D_i^\tau(p)$, we have $u_i \geq \tau f_i(k) - kp$ for all $k \in \{0, 1, \ldots, m\}$, so $(p, \boldsymbol{u})$ is feasible for (τ-Dual). Market clearing gives

$$\sum_{i=1}^{n} f_i(x_i) = mp + \sum_{i=1}^{n} u_i.$$

Weak duality then yields

$$\sum_{i=1}^{n} f_i(x_i) \geq \text{OPT}(\tau\text{-Dual}) = \text{OPT}(\tau\text{-LP}) = \tau \cdot \text{OPT}(\text{LP}) \geq \tau \cdot \text{OPT}(\text{SWM}).$$

Strong duality and the factor τ on the primal objective give the two equalities, and the last inequality follows because (LP) relaxes (SWM). □

As a consequence of Proposition 2, the closer τ is to one, the larger the guaranteed fraction of optimal welfare. Rearranging the same dual comparison gives, for any τ-WE $(p, \boldsymbol{x})$, buyer i, and quantity $q \in \{0, 1, \ldots, m\}$,

$$U_i(p, q) - U_i(p, x_i) \leq (1 - \tau) f_i(q).$$

Taking $q = x_j$, this gap is buyer i's envy of buyer j's quantity at the common price p, and it is at most $(1 - \tau) f_i(x_j)$, controlled by the same parameter τ.

4. The Tightest Relaxation Problem

So far, we have established that an exact WE may fail to exist, while a τ-WE can exist for some $\tau < 1$. Since a larger τ corresponds to a smaller relaxation of buyer optimality and guarantees a larger fraction of the optimal social welfare, it is natural to ask how large τ can be for a given market.

In this section, we formalize this market-specific question through the Tightest Relaxation problem. We first show that computing the largest feasible τ exactly is NP-hard. We then derive a lower bound on the tightest relaxation in terms of the market's relative distance from submodularity, which quantifies how far the buyers' valuations are from submodular functions.

4.1. Formulation and Basic Properties

We define the *Tightest Relaxation* problem (TR) as follows:

$$\begin{aligned} \max_{\tau\in(0,1],\, p\geq 0,\, \boldsymbol{x}} \quad & \tau \\ \text{s.t.} \quad & x_i \in D_i^{\tau}(p)\ \forall i \in [n], \quad \textstyle\sum_{i=1}^{n} x_i = m, \end{aligned} \tag{TR}$$

with the optimal value denoted by τ^*. Equivalently, τ^* is the largest value of τ for which a τ-WE exists.

The next lemma establishes that (TR) is well-defined.

Lemma 2. *Every market satisfies $\tau^* \geq 1/m$, and the maximum in* (TR) *is attained.*

The lemma implies that $\tau^* = 1$ if and only if a Walrasian equilibrium exists. More generally, it ensures that τ^* is attained by a feasible solution to (TR). However, computing τ^* is NP-hard, as established next.

Proposition 3 **(NP-Hardness of the Tightest Relaxation)**. *Computing τ^* is NP-hard. In particular, deciding whether $\tau^* = 1$, and hence whether a Walrasian equilibrium exists, is NP-hard.*

4.2. A Tight Lower Bound

In this subsection, we use the market's relative distance from submodularity γ to derive a lower bound on the optimal value τ^* of (TR). As defined in Section 2, γ measures the relative gaps between buyers' valuations and their concave envelopes. The main result of this subsection is presented next.

Theorem 1. *For every market with relative distance from submodularity γ, the optimal value τ^* of* (TR) *satisfies*

$$\tau^* \geq 1 - \gamma.$$

The bound is tight as a function of γ: for every $\gamma \in [0, 1)$, there exists a market whose relative distance from submodularity is exactly γ and whose tightest relaxation equals $1 - \gamma$.

Proof of Theorem 1. Write $\beta := 1-\gamma$. If $\beta = 0$, then $\tau^* \geq 0 = 1-\gamma$, so it suffices to consider $\beta > 0$. By Lemma 1, each concave envelope $\overline{f}_i$ is normalized, nonnegative, monotone, and submodular. The market with valuations $\{\overline{f}_i\}_{i\in[n]}$ therefore admits a Walrasian equilibrium supported by a single price. Let $\widehat{p} \geq 0$ and $\boldsymbol{x}$ denote the equilibrium's price and allocation. Thus $\sum_i x_i = m$ and

$$\overline{f}_i(x_i) - \widehat{p}x_i \geq \overline{f}_i(k) - \widehat{p}k$$

for every buyer i and quantity k. Set $p := \beta\widehat{p}$. Using (4) and $\overline{f}_i(k) \geq f_i(k)$, and multiplying the preceding inequality by β, we obtain

$$f_i(x_i) - p\,x_i \;\geq\; \beta\,\overline{f}_i(x_i) - \beta\,\widehat{p}\,x_i \;\geq\; \beta\,\overline{f}_i(k) - \beta\,\widehat{p}\,k \;\geq\; \beta\,f_i(k) - p\,k.$$

Hence $x_i \in D_i^{\beta}(p)$ for every buyer. The allocation clears the market, so $(p, \boldsymbol{x})$ is a β-WE and $\tau^* \geq \beta = 1-\gamma$.

To establish tightness, consider the market

$$f_1(0) = 0, \quad f_1(k) = \tfrac{1}{\beta} \;\; (k \geq 1); \qquad f_2(k) = \beta k \;\; (0 \leq k < m), \quad f_2(m) = m. \tag{Market 2}$$

For every integer $m \geq 1 + (1-\beta)/\beta^2$, this market has relative distance from submodularity γ and tightest relaxation $\tau^* = 1-\gamma$. Thus the lower bound is attained for every $\gamma \in [0,1)$. We defer the detailed calculation to Appendix C. □

Example 1. Theorem 1 implies that every market with $\gamma < 1$ admits a $(1-\gamma)$-WE, and $\gamma = 0$ recovers an exact WE. For Market 1 introduced earlier in Subsection 3.2, we have

$$\gamma = \frac{(m-1)\xi}{m + (m-1)\xi} < \xi,$$

so $\tau^* \geq 1-\gamma > 1-\xi$ even though no exact WE exists. For $m \geq 3$, this first comparison is strict, i.e., $\tau^* > 1-\gamma$. For Market 2 defined above, however, the $\tau^* = 1-\gamma$ bound is attained for every $\gamma \in [0,1)$. Thus, $1-\gamma$ is the best possible worst-case lower bound, but it need not be exact for a given market. These two instances leave open whether a large relative distance from submodularity

necessarily implies a small τ^*. The following parametric example shows that no such implication holds.

$$f_1(0)=0, \quad f_1(k)=1+\xi \ \ (k\geq 1); \qquad f_2(k)=0 \ \ (0\leq k<m), \quad f_2(m)=m, \qquad \xi\in(0,1). \tag{Market 3}$$

For $m\geq(1+\xi)^2$, its lower bound and tightest relaxation parameter satisfy

$$\gamma=1, \qquad 1-\gamma=0, \qquad \tau^*=\frac{1}{1+\xi}>0.$$

Here, $\gamma=1$ because $f_2(k)=0$ at every intermediate quantity while $\overline{f}_2(k)=k>0$. Thus, the lower bound is identically zero, while $\tau^*=1/(1+\xi)$ approaches one as ξ approaches zero. We defer the detailed analysis of Markets 1–3 to Appendix C.

Market	Relative distance γ	Comparison with $1-\gamma$	lower bound
Market 1	$\gamma\in(0,1)$	$\tau^*>1-\gamma$	Conservative
Market 2	$\gamma\in[0,1)$	$\tau^*=1-\gamma$	Tight
Market 3	$\gamma=1$	$1-\gamma=0<\tau^*=\frac{1}{1+\xi}$	Uninformative

Table 1 Evaluation of the lower bound across Markets 1-3. The rows for Markets 1 and 3 assume $m\geq 3$ and $m\geq(1+\xi)^2$, respectively. For Market 2, $\beta=1-\gamma$ and $m\geq 1+(1-\beta)/\beta^2$.

The preceding examples clarify the role of Theorem 1. When valuations depart only slightly from submodularity, a single price can still support exact market clearing with only a correspondingly small relaxation of buyer optimality and loss of social welfare, and this dependence on γ is the best possible among guarantees based on γ alone. As an estimate of a given market's τ^*, however, $1-\gamma$ can be conservative, as in Market 1, or uninformative, as in Market 3, and can therefore substantially understate the tightest relaxation.

5. An Exact-Clearing FPTAS for the Tightest Relaxation Problem

Section 4 established that computing τ^* exactly is NP-hard, and derived a lower bound based on the market's relative distance from submodularity γ. That bound, however, does not resolve the computational problem. It may be conservative or even uninformative for a particular market, and computing γ itself requires evaluating every valuation at all quantities and constructing the corresponding concave envelopes, with a cost that grows linearly with m. We therefore seek to approximate a τ^*-WE directly in a computationally efficient way.

We measure computational complexity in the standard value-oracle model (see, e.g., Dobzinski and Nisan 2007, Dobzinski and Dughmi 2013). A query specifies a buyer i and a quantity $k \in \{0, \dots, m\}$ and returns $f_i(k)$. Each value query, arithmetic operation, and comparison has unit cost, and the running time is measured by the number of those queries and operations. Our target is an approximation scheme that approximates a τ^*-WE within ε precision, whose computational complexity is polynomial in n, $\log m$, and $1/\varepsilon$.

Our main algorithmic result is the following exact-clearing FPTAS:

THEOREM 2 **(Exact-clearing FPTAS)**. *For every $\varepsilon \in (0,1)$ and arbitrary monotone valuations $\{f_i\}_{i\in[n]}$ with $f_i(0) = 0$, there exists an FPTAS (Algorithm 1) that finds a $(1-\varepsilon)\tau^*$-WE in time polynomial in n, $\log m$, and $1/\varepsilon$.*

The algorithm returns a market-clearing outcome under a single price that comes arbitrarily close to the tightest relaxation the market admits. Computing τ^* exactly is NP-hard, and an FPTAS is then the best efficient algorithm one can hope for. The running time scales with the logarithm of the supply and is independent of how widely buyers' values are spread. Therefore the algorithm remains efficient when the supply is large and the market is complicated.

Roadmap. We now give an overview of the challenges and the construction of our algorithm. Computing the tightest relaxation requires a search over a relaxation level, a price, and an integral allocation that satisfies relaxed buyer optimality and market clearing. At a fixed relaxation level and price, buyers' RDCs determine the quantities available, but deciding whether one quantity

can be chosen for each buyer to achieve market clearing is a Multiple-Choice Subset Sum (MCSS) problem, which is NP-hard (see, e.g., Kellerer et al. 2004). Moreover, the relaxation level and the price both range over continua, and each pair determines a different collection of quantities available for the selection. A direct search over this space does not yield an algorithm whose running time is polynomial in $\log m$.

We resolve the challenges by introducing a bicriteria FPTAS Bicriteria(ε,δ) as a subroutine, which returns a relaxation level of at least $(1-\varepsilon)\tau^*$ and a price-allocation pair satisfying relaxed buyer optimality, with at most δm units unsold. Section 5.1 restores market clearing from two different calls to this subroutine. When the supply is small, we choose δ so that the shortfall is less than one unit, and integrality then forces market clearing. When the supply is large, the subroutine is called with calibrated precisions, and a multiplicative completion step assigns the remaining units while scaling the relaxation level and the price together to convert the allocation to a market-clearing one.

To design an efficient bicriteria subroutine, we construct rounded RDCs that can be represented as unions of a small number of integer intervals. Controlled truncation and rounding of valuations, together with rounding of quantities, keep both the number of intervals and their possible endpoints small. These interval representations induce a discretized search over finitely many relaxation levels and prices. An MCSS approximation algorithm then selects one quantity from each union so that the allocation nearly clears the market within the desired precision. The discretization is constructed so that an optimal market-clearing solution can be mapped, after controlled scaling and rounding, to a near-clearing candidate in the discretized search. Conversely, every candidate accepted by the discretized search is demand-feasible in the original market after a controlled reduction in the relaxation level. Together, these guarantees allow the discretized search to return a near-optimal (within ε precision) relaxation and a near-clearing (with up to δm unallocated units) outcome that is demand-feasible in the original market. Under this construction, the bicriteria subroutine runs in time polynomial in n, $\log m$, $1/\varepsilon$, and $1/\delta$. Section 5.2 constructs the rounded RDCs and derives their interval representations, Section 5.3 uses them to construct the bicriteria subroutine. Appendix A provides the full construction details and proofs.

5.1. The Exact-Clearing FPTAS

This subsection constructs the exact-clearing FPTAS from the bicriteria FPTAS subroutine. The Bicriteria FPTAS is the core to our algorithmic design, and we first states the guarantees of it, and defer the detailed construction in later sections.

PROPOSITION 4 **(Bicriteria FPTAS)**. *There exists an algorithm* BICRITERIA(ε,δ) *(Algorithm 2) such that, for every* $\varepsilon,\delta\in(0,1)$ *and arbitrary monotone valuations* $\{f_i\}_{i\in[n]}$ *with* $f_i(0)=0$*, it returns* $\widehat{\tau}\in(0,1]$*,* $\widehat{p}\geq 0$*, and* $\widehat{\boldsymbol{x}}\in\{0,\dots,m\}^n$ *satisfying*

(i) $\widehat{\tau}\geq(1-\varepsilon)\tau^*$*;*

(ii) $\widehat{x}_i\in D_i^{\widehat{\tau}}(\widehat{p})$ *for every* i*;*

(iii) $\sum_i\widehat{x}_i\in[(1-\delta)m,m]$.

The running time is polynomial in n*,* $\log m$*,* $1/\varepsilon$*, and* $1/\delta$.

The subroutine takes two multiplicative precision parameters: ε controls the approximation to the tightest relaxation, and δ controls the clearing slack. The returned outcome need not be feasible for (TR), because it need not clear the market. However, the subroutine returns a controlled approximation to a τ^*-WE, and can be implemented efficiently. We apply the subroutine differently in two supply regimes to construct the Exact-Clearing FPTAS, with the corresponding parameter choices specified in Algorithm 1.

Small supply. When the supply m is at most on the order of n/ε, a running time polynomial in m remains polynomial in n and $1/\varepsilon$. In such a case, we can choose a small clearing slack to force market clearing, and still maintain an efficient running time.

COROLLARY 1 **(Exact-clearing specialization)**. *Fix* $\varepsilon\in(0,1)$*. Then* BICRITERIA$(\varepsilon,1/(2m))$ *returns a* $\widehat{\tau}$*-WE satisfying* $(1-\varepsilon)\tau^*\leq\widehat{\tau}\leq\tau^*$*, in time* $\mathrm{poly}(n,m,1/\varepsilon)$.

Setting $\delta=1/(2m)$ leaves less than one unit unsold, so integrality forces the returned allocation to be market-clearing. This makes the returned outcome a $(1-\varepsilon)\tau^*$-WE. The subroutine runs in time polynomial in m, which is bounded by a polynomial in n and $1/\varepsilon$ in the small supply regime and thus remains efficient.

Large supply. When the supply m is larger than on the order of n/ε, a running time polynomial in m is no longer polynomial in n and $1/\varepsilon$. Choosing the clearing slack to keep the subroutine computationally efficient may leave more than one unit unsold. The following lemma restores market clearing from such an outcome of the bicriteria subroutine.

LEMMA 3 **(Multiplicative Completion)**. *Let each f_i be monotone with $f_i(0) = 0$. Fix $\tau \in (0,1]$, $p \geq 0$, and an integral allocation $\boldsymbol{x}$ with $x_i \in D_i^{\tau}(p)$ for every i. Write $s \coloneqq \sum_i x_i$, $q \coloneqq |\{i : x_i > 0\}|$, and $c \coloneqq s/(m+q-1)$, and assume $1 \leq s < m$. In $O(n)$ operations one can construct a market-clearing allocation $\boldsymbol{y}$ such that $y_i \in D_i^{c\tau}(cp)$ for every i. If $s \geq (1-\delta)m$ for some $\delta \in [0,1)$, then $c \geq (1-\delta)/(1+(q-1)/m)$.*

Intuitively, the lemma assigns the unsold units to buyers with positive allocations while scaling the price and relaxation level together to preserve relaxed buyer optimality. In the large-supply regime, both the initial shortfall and the number of such buyers are small relative to m, so the scaling loss is small. The lemma therefore restores exact market clearing with only a controlled multiplicative loss in the relaxation guarantee.

Algorithm 1 combines the two supply regimes described above, and we defer the full analysis to Appendix B. The remainder of this section constructs the bicriteria FPTAS subroutine used in both branches of Algorithm 1.

5.2. Preprocessing and Rounded RDCs

Before constructing the bicriteria subroutine, we preprocess buyers' valuations and quantities. Truncation and value rounding limit the number of distinct valuation levels, while quantity rounding restricts the possible RDC endpoints at each fixed (τ, p). Together, these steps produce rounded RDCs, each of which is a union of at most one integer interval for every rounded valuation level. The compact interval representations of rounded RDCs become useful in designing our bicriteria subroutine. This subsection develops the preprocessing and establishes its properties.

Algorithm 1: Exact-Clearing FPTAS for the Tightest Relaxation (TR).

Input: Valuations $\{f_i\}_{i\in[n]}$ via value oracle, supply m, precision $\varepsilon \in (0,1)$.

Output: Relaxation $\widehat{\tau}$, price $\widehat{p}$, allocation $\widehat{\boldsymbol{x}}$.

if $m \le 4(n-1)/\varepsilon$ **then** `// Small supply regime`
  $(\tau, p, \boldsymbol{x}) \leftarrow$ Bicriteria(ε,$1/(2m)$); **return** $(\tau, p, \boldsymbol{x})$
$(\tau, p, \boldsymbol{x}) \leftarrow$ Bicriteria($\varepsilon/4$,$\varepsilon/4$) `// Large supply regime`
$s \leftarrow \sum_i x_i$
if $s = m$ **then**
  return $(\tau, p, \boldsymbol{x})$
$q \leftarrow \big|\{i : x_i > 0\}\big|$; $\quad c \leftarrow s/(m+q-1)$; $\quad \boldsymbol{y} \leftarrow \boldsymbol{x}$; $\quad r \leftarrow m - s$
foreach $i \in [n]$ *with* $x_i > 0$ **do** `// Multiplicative completion`
  $\Delta \leftarrow \min\big\{r, \min\{m, \lfloor x_i/c \rfloor\} - y_i\big\}$; $\quad y_i \leftarrow y_i + \Delta$; $\quad r \leftarrow r - \Delta$
return $(c\tau, cp, \boldsymbol{y})$

Value truncation. We begin with a buyer-specific truncation at the threshold $f_i(m)/m^2$:

$$f_i^{\mathrm{tr}}(k) := \begin{cases} 0, & f_i(k) < f_i(m)/m^2, \\ f_i(k), & \text{otherwise.} \end{cases} \tag{6}$$

Let $D^{\tau}_{i,\mathrm{tr}}(p)$ denote the τ-RDC under f_i^{tr}. The following lemma characterizes how this truncation affects relaxed demand correspondences.

Lemma 4 **(Demand preservation under truncation)**. *For every buyer i, price $p \ge 0$, and $\tau \in (0,1]$, $D^{\tau}_{i,\mathrm{tr}}(p) \subseteq D^{\tau}_i(p)$, with equality for every $\tau \in [1/m, 1]$.*

The containment $D^{\tau}_{i,\mathrm{tr}}(p) \subseteq D^{\tau}_i(p)$ means that any quantity selected after truncation remains in the buyer's original relaxed demand correspondence at the same relaxation level and price. The equality for $\tau \ge 1/m$, together with $\tau^* \ge 1/m$ from Lemma 2, ensures that truncation does not remove an optimal solution of (TR). Moreover, truncation bounds the ratio between retained positive values by m^2, ensuring that, for a given rounding precision, the value grid introduced next grows only logarithmically with m and not with valuation magnitudes.

Value rounding. Let $\varepsilon_1 \in (0, 1/2]$, set $\alpha := 1 + \varepsilon_1$ and $J := 1 + \lceil 2 \log_\alpha m \rceil$, and consider a buyer $i \in [n]$ with $f_i(m) > 0$. Define the buyer-specific grid and rounded valuation by

$$\mathcal{V}_i := \{0\} \bigcup \{f_i(m)\alpha^{-j} : 0 \le j \le J-1\}, \qquad \widetilde{f}_i(k) := \max\{v \in \mathcal{V}_i : v \le f_i^{\mathrm{tr}}(k)\}. \tag{7}$$

For a buyer with $f_i(m) = 0$, set $\widetilde{f}_i \equiv 0$. The rounding satisfies $f_i^{\mathrm{tr}}(k)/\alpha \le \widetilde{f}_i(k) \le f_i^{\mathrm{tr}}(k)$, so it changes the truncated values by at most a factor of α.

The rounded valuation also admits a compact interval representation. Because f_i^{tr} is monotone and rounding down to $\mathcal{V}_i$ preserves monotonicity, $\widetilde{f}_i$ is also monotone. Moreover, it takes values only in $\mathcal{V}_i$, so its domain is partitioned into at most $J+1$ maximal constant intervals, which we call plateaus. Denote these intervals by $I_{i,j} = [a_{i,j}, b_{i,j}] \bigcap \mathbb{Z}$ for $j = 0, \dots, L_i$ with $L_i \le J$, and let $v_{i,j}$ be their common values. To construct the local quantity grids used to form the rounded RDCs below, the algorithm locates these fixed plateau boundaries efficiently by monotone binary search without enumerating all $k \in \{0, \dots, m\}$.

Quantity grids and rounded demand correspondences. For fixed $\tau \in (0, 1]$ and $p \ge 0$, write $\widetilde{D}_i^\tau(p)$ for the τ-RDC under $\widetilde{f}_i$. That is, $k \in \widetilde{D}_i^\tau(p)$ if and only if $\widetilde{f}_i(k) - kp \ge \tau \widetilde{f}_i(k') - k'p$ for every $k' \in \{0, \dots, m\}$. Because $\widetilde{f}_i$ is constant on each plateau, utility is nonincreasing in quantity there, so $\widetilde{D}_i^\tau(p) \bigcap I_{i,j}$ is either empty or contains all integers from $a_{i,j}$ to some right endpoint in $I_{i,j}$. Although $a_{i,j}$ and $b_{i,j}$ are fixed, this right endpoint varies with (τ, p) and may take $\Theta(m)$ values on a long plateau. We therefore retain only a sparse set of candidate demand endpoints.

To restrict the possible right endpoints to a sparse set, let $\varepsilon_2 \in (0, 1/2]$, set $M_0 := \lceil 1/\varepsilon_2 \rceil$ and define the quantity grid by

$$\mathcal{G} := \big([0, \min\{m, M_0\}] \bigcap \mathbb{Z}\big) \bigcup \Big\{ \lceil (1+\varepsilon_2)^t \rceil : t \in \mathbb{Z}_{\ge 0},\, M_0 < \lceil (1+\varepsilon_2)^t \rceil < m \Big\} \bigcup \{m\}. \tag{8}$$

The first set retains small quantities exactly, while the second provides geometrically spaced grid points for larger quantities. This keeps the candidate set sparse while controlling the clearing slack introduced when demand endpoints are rounded down. The grid contains $|\mathcal{G}| = O(\varepsilon_2^{-1}(1 + \log m))$ points, so its size depends on the supply only through $\log m$.

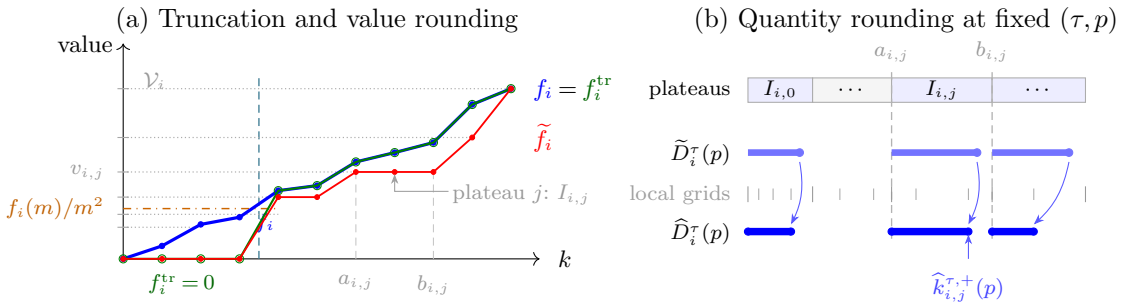


Figure 2 Illustration of the Preprocessing step. Panel (a) first truncates f_i below $f_i(m)/m^2$ to obtain f_i^{tr}, and then rounds its retained values down to $\mathcal{V}_i$ to obtain the plateau valuation $\widetilde{f}_i$. At fixed (τ, p), Panel (b) first computes the RDC $\widetilde{D}_i^\tau(p)$ under $\widetilde{f}_i$, and then rounds its right endpoint within each nonempty plateau down to the corresponding local grid to obtain the rounded RDC $\widehat{D}_i^\tau(p)$.

For each plateau $I_{i,j}$, define the local grid

$$\mathcal{G}_{i,j} := (\mathcal{G} \bigcap I_{i,j}) \bigcup \{a_{i,j}, b_{i,j}\}. \tag{9}$$

This local grid consists of the global grid points that lie in the plateau, together with its two endpoints. Including $a_{i,j}$ makes the rounded endpoint well-defined whenever $\widetilde{D}_i^\tau(p) \bigcap I_{i,j}$ is nonempty, while including $b_{i,j}$ retains the original endpoint when the entire plateau is demanded.

Let $J_i^\tau(p)$ be the set of plateau indices for which $\widetilde{D}_i^\tau(p) \bigcap I_{i,j}$ is nonempty. For each $j \in J_i^\tau(p)$, let $\widehat{k}_{i,j}^{\tau,+}(p)$ be the largest local-grid quantity in $\widetilde{D}_i^\tau(p)$. Equivalently, $\widehat{k}_{i,j}^{\tau,+}(p)$ is the right endpoint of $\widetilde{D}_i^\tau(p) \bigcap I_{i,j}$ rounded down to $\mathcal{G}_{i,j}$. Keeping all quantities from $a_{i,j}$ to this rounded endpoint defines the rounded demand correspondence

$$\widehat{D}_i^\tau(p) := \bigcup_{j \in J_i^\tau(p)} [a_{i,j}, \widehat{k}_{i,j}^{\tau,+}(p)] \bigcap \mathbb{Z}. \tag{10}$$

Therefore, $\widehat{D}_i^\tau(p)$ retains at most one integer interval from each plateau and is thus a union of at most $J+1$ such intervals. With this compact interval representation, the search need only consider local-grid points as possible right endpoints, and not enumerating all quantities in each plateau.

The rounded RDCs also retain the nesting property: for fixed $p \geq 0$ and $0 < \tau' \leq \tau \leq 1$, $\widehat{D}_i^\tau(p) \subseteq \widehat{D}_i^{\tau'}(p)$. Thus, an allocation feasible at one relaxation level remains available when the level is lowered.

Finally, the rounded-RDC construction must also connect the discretized search to the original market in both directions. In the following lemma, part (a) ensures that every quantity selected from a rounded RDC remains feasible for the original demand correspondence after scaling the relaxation level. Part (b) ensures that an optimal market-clearing allocation survives the discretization with only a controlled quantity loss.

LEMMA 5 **(Guarantees of rounding)**. *We have the following guarantees:*

(a) (Original-demand feasibility) *For any* $i \in [n]$, $\tau \in (0,1]$, $p \geq 0$, *and* $\widehat{x}_i \in \widehat{D}_i^{\tau}(p)$, *we have* $\widehat{x}_i \in D_i^{\tau/\alpha}(p)$.

(b) (Rounding an optimal solution) *Let* $(\tau^*, p^*, \boldsymbol{x}^*)$ *be an optimal solution to* (TR), *and set* $\widetilde{\tau} \coloneqq \tau^*/\alpha$ *and* $\widetilde{p} \coloneqq p^*/\alpha$. *Then there exists an allocation* $\widehat{\boldsymbol{x}}$ *such that* $\widehat{x}_i \in \widehat{D}_i^{\widetilde{\tau}}(\widetilde{p})$ *for every* $i \in [n]$ *and*

$$(1-\varepsilon_2)\, m \;\leq\; \sum_{i=1}^{n} \widehat{x}_i \;\leq\; m.$$

To summarize, the preprocessing starts from the original RDC $D_i^{\tau}(p)$, truncates and rounds its underlying valuation, and then rounds the corresponding RDC endpoints on local quantity grids to obtain a rounded RDC $\widehat{D}_i^{\tau}(p)$. Figure 2 illustrates this progression. In particular, $J_i^{\tau}(p)$ identifies the plateaus on which $\widetilde{D}_i^{\tau}(p)$ is nonempty, while $\widehat{k}_{i,j}^{\tau,+}(p)$ gives the rounded right endpoint on each such plateau; the corresponding left endpoint $a_{i,j}$ is fixed by the preprocessing. Hence, only $O(J)$ plateau indices and interval endpoints need to be tracked. In the next subsection, we build the Bicriteria FPTAS based on the preprocessing and the rounded RDCs.

5.3. Construction of the Bicriteria Subroutine

The preprocessing and rounded RDCs constructed in Section 5.2 make it possible to construct an efficient bicriteria subroutine. For each fixed relaxation level τ, we reduce the continuous price space to a finite candidate-price list. Once a candidate price p is fixed, selecting one quantity from each rounded RDC $\widehat{D}_i^{\tau}(p)$ to maximize the total allocation without exceeding the supply is a Multiple-Choice Subset Sum (MCSS) problem. We develop an MCSS approximation algorithm for this problem and combine it with the candidate-price list and a geometric search over relaxation levels to construct Algorithm 2.

Candidate prices. We nowreduce the continuous price space at a given relaxation level τ to a finite candidate-price list. The compact interval representations of the rounded RDCs constructed in Section 5.2 make this reduction possible by allowing us to identify all prices at which some $\widehat{D}_i^\tau(p)$ can change. More explicitly, for every plateau j and quantity $k \in I_{i,j}$, Equation (10) gives

$$k \in \widehat{D}_i^\tau(p) \quad \Leftrightarrow \quad \exists g \in \mathcal{G}_{i,j} :\ k \le g,\ v_{i,j} - gp \ge \max_\ell \{\tau v_{i,\ell} - a_{i,\ell} p\}.$$

As p varies, the maximum on the right-hand side of the displayed inequality is the upper envelope of the affine functions $\tau v_{i,\ell} - a_{i,\ell} p$. Prices at which two such functions are equal partition the price space into intervals on which the same function determines the maximum. Within each interval, the rounded RDC can change only when the displayed inequality becomes tight at a local-grid quantity g, or equivalently, when $v_{i,j} - gp = \tau v_{i,\ell} - a_{i,\ell} p$ for an ℓ that attains the maximum. It therefore suffices to consider the following candidate prices:

$$\begin{aligned} \widetilde{\mathcal{L}}^\tau := \{0\} \bigcup \Big\{ p \ge 0 : \tau v_{i,j} - a_{i,j} p = \tau v_{i,\ell} - a_{i,\ell} p,\ i \in [n],\ 0 \le j < \ell \le L_i \Big\} \\ \bigcup \Big\{ p \ge 0 : v_{i,j} - gp = \tau v_{i,\ell} - a_{i,\ell} p,\ i \in [n],\ j, \ell \in \{0, \dots, L_i\},\ g \in \mathcal{G}_{i,j},\ g \ne a_{i,\ell} \Big\}. \end{aligned} \tag{11}$$

The first set contains prices at which the affine function that attains the maximum may switch, and the second contains prices at which a local-grid inequality becomes tight. In the second set, $g \ne a_{i,\ell}$ gives the two sides of each equality different slopes, so the equality yields at most one price. On every open interval between consecutive prices in $\widetilde{\mathcal{L}}^\tau$, the same affine function attains the maximum and the status of every local-grid inequality remains unchanged. Thus, each rounded RDC is unchanged throughout the interval. Consequently, we attain the following lemma:

LEMMA 6 **(Finite-price reduction)**. *Fix $\tau \in (0, 1]$.*

(a) *The list $\widetilde{\mathcal{L}}^\tau$ has $O(nJ^2 + nJ|\mathcal{G}|)$ entries and can be generated in the same time by direct enumeration.*

(b) *If an allocation $\boldsymbol{x}$ satisfies $x_i \in \widehat{D}_i^\tau(p)$ for every i at some price $p \ge 0$, then there exists $p' \in \widetilde{\mathcal{L}}^\tau$ such that $x_i \in \widehat{D}_i^\tau(p')$ for every i.*

Lemma 6 shows that it suffices to check $\widetilde{\mathcal{L}}^\tau$ rather than the continuous price space, without excluding any feasible allocation. Since $J = 1 + \lceil 2 \log_{1+\varepsilon_1} m \rceil$ and $|\mathcal{G}| = O(\varepsilon_2^{-1}(1 + \log m))$, the list size depends on the supply only through $\log m$.

MCSS Approximation Algorithm. For each fixed relaxation level τ and candidate price p, the remaining question is to determine the maximum total quantity that can be allocated by selecting one quantity from each $\widehat{D}_i^{\tau}(p)$ without exceeding the supply m. This is an instance of the Multiple-Choice Subset Sum problem.

Definition 5 (Multiple-Choice Subset Sum). Given sets $S_1, \ldots, S_n \subseteq \{0, 1, \ldots, m\}$ and a target $T \in \{0, 1, \ldots, m\}$, the *Multiple-Choice Subset Sum* (MCSS) problem is to compute

$$M^* := \max\Big\{ \textstyle\sum_{i=1}^n x_i : \boldsymbol{x} \in \prod_{i=1}^n S_i,\ \sum_{i=1}^n x_i \le T \Big\}$$

and return an allocation attaining M^* when the displayed feasible set is nonempty. Otherwise the instance is infeasible.

In our problems, $S_i = \widehat{D}_i^{\tau}(p)$ and $T = m$. Thus, M^* is the maximum total allocation at (τ, p), and $m - M^*$ is the minimum clearing shortfall. The interval representation of each S_i from Section 5.2 enables the following MCSS approximation algorithm.

Proposition 5 **(MCSS Approximation Algorithm)**. *Suppose that R is a positive integer and each $S_i \subseteq \{0, 1, \ldots, m\}$ is a union of at most R disjoint integer intervals. For any target $T \in \{0, 1, \ldots, m\}$ and $\delta_{\mathrm{mcss}} \in (0, 1)$, there is an algorithm that returns* Infeasible *if and only if the feasible set in Definition 5 is empty. Otherwise, it returns a feasible allocation $\boldsymbol{x}$ satisfying*

$$M^* - \delta_{\mathrm{mcss}} m \le \sum_i x_i \le M^* \le T.$$

Its running time is $O\big((n^2 R/\delta_{\mathrm{mcss}}) \log(nR/\delta_{\mathrm{mcss}})\big)$. If $\delta_{\mathrm{mcss}} < 1/m$, the returned sum equals M^*.

The algorithm adapts the standard dynamic-programming approach for MCSS (see e.g., Kellerer et al. (2004)) and takes advantage of the compact interval representations constructed in Section 5.2. In particular, each rounded RDC contains at most one interval per plateau, so $R \le J + 1$, where $J = O(1 + \varepsilon_1^{-1} \log m)$. Consequently, the running time of applying the MCSS Approximation Algorithm depends on the supply only through $\log m$.

Algorithm 2: BICRITERIA(ε,δ) for the Tightest Relaxation (TR).

Input: Valuations $\{f_i\}_{i\in[n]}$ via value oracle, supply m, precision $\varepsilon \in (0,1)$, clearing slack $\delta \in (0,1)$.

Output: Relaxation $\widehat{\tau}$, price $\widehat{p}$, allocation $\widehat{\boldsymbol{x}}$.

Parameter calibration. Set $(\varepsilon_1, \varepsilon_2, \varepsilon_{\mathrm{bs}}, \delta_{\mathrm{mcss}}) \leftarrow (\varepsilon/6, \delta/2, \varepsilon/3, \delta/2)$ and $\alpha \leftarrow 1+\varepsilon_1$ as in (12)

Preprocessing. Set $J \leftarrow 1 + \lceil 2\log_\alpha m \rceil$. For each buyer, construct f_i^{tr} and $\widetilde{f}_i$ using (6)–(7) with value-rounding parameter ε_1 and $\alpha = 1+\varepsilon_1$, taking $\widetilde{f}_i \equiv 0$ when $f_i(m) = 0$. Construct $\mathcal{G}$ and $\{\mathcal{G}_{i,j}\}$ using (8) and (9) with quantity-rounding parameter ε_2

Relaxation-level grid. Construct $\{\lambda_t\}_{t=0}^{T_\lambda}$ using (13) with parameter $\varepsilon_{\mathrm{bs}}$

Function `ClearingTest(`λ`)`**:**

  Candidate prices. Generate $\widetilde{\mathcal{L}}^\lambda$ using (11) at relaxation level λ

  foreach $p \in \widetilde{\mathcal{L}}^\lambda$ **do**

    Allocation at price p**.** Form $S_i \leftarrow \widehat{D}_i^\lambda(p)$ using (10) for each $i \in [n]$. Run the MCSS approximation algorithm on $(\{S_i\}_{i\in[n]}, m)$ with parameter δ_{mcss}, and let $\boldsymbol{x}'$ denote its output

    if $\boldsymbol{x}' \neq$ INFEASIBLE **and** $\sum_i x_i' \geq (1-\delta)\, m$ **then return** $(p, \boldsymbol{x}')$

  return REJECT

if `ClearingTest(`λ_0`)` *returns some* $(p, \boldsymbol{x}')$ **then return** $(\lambda_0/\alpha, p, \boldsymbol{x}')$

$(p, \boldsymbol{x}') \leftarrow$ `ClearingTest(`λ_{T_λ}`)`;

$(\tau^{\mathrm{acc}}, \widehat{p}, \widehat{\boldsymbol{x}}) \leftarrow (\lambda_{T_\lambda}, p, \boldsymbol{x}')$

$(t_-, t_+) \leftarrow (0, T_\lambda)$ `// rejected and accepted endpoints`

while $t_+ - t_- > 1$ **do**

  $t \leftarrow \lfloor (t_- + t_+)/2 \rfloor$

  if `ClearingTest(`λ_t`)` *returns some* $(p, \boldsymbol{x}')$ **then**

    $(t_+, \tau^{\mathrm{acc}}, \widehat{p}, \widehat{\boldsymbol{x}}) \leftarrow (t, \lambda_t, p, \boldsymbol{x}')$

  else

    $t_- \leftarrow t$

return $(\tau^{\mathrm{acc}}/\alpha, \widehat{p}, \widehat{\boldsymbol{x}})$

Construction of the Bicriteria FPTAS. We now construct the Bicriteria FPTAS (Algorithm 2) by combining the preprocessing and rounded RDCs from Section 5.2 with the candidate-price list and the MCSS approximation algorithm developed above. We first calibrate the internal parameters:

$$(\varepsilon_1, \varepsilon_2, \varepsilon_{\mathrm{bs}}, \delta_{\mathrm{mcss}}) \coloneqq \left(\frac{\varepsilon}{6}, \frac{\delta}{2}, \frac{\varepsilon}{3}, \frac{\delta}{2}\right), \qquad \alpha \coloneqq 1+\varepsilon_1. \tag{12}$$

Here, ε_1 and ε_{bs} control the multiplicative loss, while ε_2 and δ_{mcss} control the clearing shortfall. With these parameters fixed, the algorithm first performs the value and quantity preprocessing and then constructs the following geometric grid of relaxation levels:

$$T_\lambda \coloneqq \left\lceil \log_{1+\varepsilon_{\text{bs}}} m \right\rceil + 1, \qquad \lambda_t \coloneqq (1+\varepsilon_{\text{bs}})^{-t}, \quad t = 0, \ldots, T_\lambda. \tag{13}$$

The algorithm uses the geometrically spaced grid of relaxation levels, ranging from $\lambda_0 = 1$ to a sufficiently small level that is guaranteed to be accepted under the parameter calibration. Adjacent grid levels differ by a factor of $1+\varepsilon_{\text{bs}}$, which controls the multiplicative loss from discretizing the relaxation level.

At each queried relaxation level λ_t, the algorithm iterates over the candidate-price list $\widetilde{\mathcal{L}}^{\lambda_t}$. For every candidate price, it forms the rounded RDCs and runs the MCSS approximation algorithm. The level is accepted if some candidate price yields a total allocation of at least $(1-\delta)m$; the corresponding price and allocation are then recorded. If no candidate price meets this threshold, the level is rejected.

The algorithm first tests $\lambda_0 = 1$. If it is accepted, the algorithm returns λ_0/α together with the recorded price and allocation. Otherwise, it initializes the search with λ_0 as the rejected upper endpoint and the smallest grid level as the accepted lower endpoint. It then performs a binary search over the grid indices until the two endpoints are adjacent. Finally, it divides the accepted level by α and returns it together with the corresponding price and allocation.

6. Concluding Remarks

We study the performance of market clearing supported by an anonymous per-unit price. We use duality of the linear representation of welfare-maximizing market clearing to identify a natural relaxation of Walrasian equilibrium. Specifically, we relax the buyer optimality condition of WE by discounting only the valuations in the dual constraints corresponding to demand correspondences, yielding the notion of τ-WE. Such a relaxation of market equilibrium preserves a single price and exact market clearing, while allocating to buyers quantities that are acceptable under their τ-discounted valuations, which guarantees at least a τ fraction of the optimal welfare.

We go on to investigate the maximum τ, denoted τ^*, for which market clearing can be supported by a single price, and show that computing τ^* is NP-hard. Nevertheless, by combining duality-like price adjustments with combinatorial arguments for partitioning the supply, we develop an FPTAS that finds τ^* to any prescribed accuracy for arbitrary monotone valuations, together with a supporting price and a market-clearing allocation. Its running time depends on the supply only through its logarithm.

Our analysis takes valuations as given and measures performance by welfare, as does the classical Walrasian benchmark. This is what makes τ^* a property of the market rather than of a mechanism. A seller facing a specific market can quantify the welfare guarantee attainable with a single price and the degree of buyer-optimality relaxation required to support market clearing at that price, rather than relying on worst-case guarantees that hold across all markets. Our approach underscores that the failure of a Walrasian equilibrium does not imply the failure of simple pricing: an approximately optimal single price can be efficiently computed and implemented for a given market.

References

Akbarpour M, Nikzad A (2020) Approximate random allocation mechanisms. *The Review of Economic Studies* 87(6):2473–2510.

Amanatidis G, Aziz H, Birmpas G, Filos-Ratsikas A, Li B, Moulin H, Voudouris AA, Wu X (2023) Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.* 322:103965.

Ausubel LM (2004) An efficient ascending-bid auction for multiple objects. *American Economic Review* 94(5):1452–1475.

Ausubel LM, Cramton P, Pycia M, Rostek M, Weretka M (2014) Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies* 81(4):1366–1400.

Azevedo EM, Budish E (2019) Strategy-proofness in the large. *The Review of Economic Studies* 86(1):81–116.

Azevedo EM, Weyl EG, White A (2013) Walrasian equilibrium in large, quasi-linear markets. *Theoretical Economics* 8(2):281–290.

Aziz H, Li B, Moulin H, Wu X (2022) Algorithmic fair allocation of indivisible items: A survey and new questions. *ACM SIGecom Exchanges* 20(1):24–40.

Balcan MF, Blum A, Mansour Y (2008) Item pricing for revenue maximization. *Proceedings of the 9th ACM Conference on Electronic Commerce*, 50–59.

Baldwin E, Jagadeesan R, Klemperer P, Teytelboym A (2023) The equilibrium existence duality. *Journal of Political Economy* 131(6):1440–1476.

Baldwin E, Klemperer P (2019) Understanding preferences: "demand types", and the existence of equilibrium with indivisibilities. *Econometrica* 87(3):867–932.

Balseiro SR, Besbes O, Castro F (2024) Mechanism design under approximate incentive compatibility. *Operations Research* 72(1):355–372.

Besbes O, Elmachtoub AN, Sun Y (2022) Technical note—static pricing: Universal guarantees for reusable resources. *Operations Research* 70(2):1143–1152, URL `http://dx.doi.org/10.1287/opre.2020.2054`.

Birmpas G, Markakis E, Telelis O, Tsikiridis A (2019) Tight welfare guarantees for pure Nash equilibria of the uniform price auction. *Theory of Computing Systems* 63(7):1451–1482.

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* 5(3):203–229.

Brânzei S, Derakhshan M, Golrezaei N, Han Y (2023a) Learning and collusion in multi-unit auctions. *Advances in Neural Information Processing Systems* 36:22191–22225.

Brânzei S, Filos-Ratsikas A, Miltersen PB, Zeng Y (2023b) Walrasian pricing in multi-unit auctions. *Artificial Intelligence* 322:103961, URL `http://dx.doi.org/10.1016/j.artint.2023.103961`.

Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.

Candogan O, Ozdaglar A, Parrilo PA (2015) Iterative auction design for tree valuations. *Operations Research* 63(4):751–771.

Candogan O, Pekeč S (2018) Efficient allocation and pricing of multifeatured items. *Management Science* 64(12):5521–5543.

Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2019) The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* 7(3):1–32.

Chen N, Deng X, Sun X (2004) On complexity of single-minded auction. *Journal of Computer and System Sciences* 69(4):675–687, URL `http://dx.doi.org/10.1016/j.jcss.2004.04.012`.

Chen N, Rudra A (2008) Walrasian equilibrium: Hardness, approximations and tractable instances. *Algorithmica* 52(1):44–64, URL `http://dx.doi.org/10.1007/s00453-007-9103-9`.

Cohen MC, Elmachtoub AN, Lei X (2022) Price discrimination with fairness constraints. *Management Science* 68(12):8536–8552.

Correa J, Cristi A, Fielbaum A, Pollner T, Weinberg SM (2022) Optimal item pricing in online combinatorial auctions. *International Conference on Integer Programming and Combinatorial Optimization*, 126–139 (Springer).

Dobzinski S, Dughmi S (2013) On the power of randomization in algorithmic mechanism design. *SIAM Journal on Computing* 42(6):2287–2304, URL `http://dx.doi.org/10.1137/090780146`.

Dobzinski S, Nisan N (2007) Mechanisms for multi-unit auctions. *Proceedings of the 8th ACM conference on Electronic commerce*, 346–351.

Dress AW, Terhalle W (1995) Well-layered maps—a class of greedily optimizable set functions. *Applied Mathematics Letters* 8(5):77–80.

Dutting P, Feldman M, Kesselheim T, Lucier B (2020) Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing* 49(3):540–582.

Dütting P, Kesselheim T, Lucier B (2020) An $O(\log\log m)$ prophet inequality for subadditive combinatorial auctions. *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 306–317.

Elmachtoub AN, Gupta V, Hamilton ML (2021) The value of personalized pricing. *Management Science* 67(10):6055–6070.

Ezra T, Feldman M, Roughgarden T, Suksompong W (2020) Pricing multi-unit markets. *ACM Transactions on Economics and Computation (TEAC)* 7(4):1–29, URL `http://dx.doi.org/10.1145/3373715`.

Feldman M, Gravin N, Lucier B (2013) Combinatorial walrasian equilibrium. *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 61–70.

Feldman M, Gravin N, Lucier B (2014) Combinatorial auctions via posted prices. *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, 123–135 (SIAM).

Feldman M, Shabtai G, Wolfenfeld A (2022) Two-price equilibrium. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5008–5015.

Fu H, Kleinberg R, Lavi R (2012) Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 586, EC '12, ISBN 9781450314152, URL http://dx.doi.org/10.1145/2229012.2229055.

Galgana R, Golrezaei N (2025) Learning in repeated multiunit pay-as-bid auctions. *Manufacturing & Service Operations Management* 27(1):200–229, URL http://dx.doi.org/10.1287/msom.2023.0403.

Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* 40(8):999–1020.

Goeree JK (2023) Yquilibrium: A theory for (non-)convex economies. *arXiv preprint arXiv:2305.06256* URL http://dx.doi.org/10.48550/arXiv.2305.06256.

Gul F, Pesendorfer W, Zhang M (2019) Market design and Walrasian equilibrium, working paper, Department of Economics, Princeton University.

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

Hartline JD, Lucier B (2015) Non-optimal mechanism design. *American Economic Review* 105(10):3102–3124.

Huang LS, Li M, Zhang B (2005) Approximation of Walrasian equilibrium in single-minded auctions. *Theoretical Computer Science* 337(1–3):390–398, URL http://dx.doi.org/10.1016/j.tcs.2005.03.008.

Hübner T (2025) Approximate equilibria in nonconvex markets: Theory and evidence from european electricity auctions. *arXiv preprint arXiv:2503.02464* URL http://dx.doi.org/10.48550/arXiv.2503.02464.

Jalota D, Ye Y (2025) Stochastic online fisher markets: Static pricing limits and adaptive enhancements. *Operations Research* 73(2):798–818, URL http://dx.doi.org/10.1287/opre.2023.0636.

Karp RM (2009) Reducibility among combinatorial problems. *50 Years of Integer Programming 1958-2008: from the Early Years to the State-of-the-Art*, 219–241 (Springer).

Kellerer H, Pferschy U, Pisinger D (2004) *Knapsack Problems*. ISBN 978-3-540-40286-2, URL `http://dx.doi.org/10.1007/978-3-540-24777-7`.

Kelso Jr AS, Crawford VP (1982) Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* 1483–1504.

Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131 (ACM).

Lubin B, Parkes DC (2012) Approximate strategyproofness. *Current Science* 1021–1032.

Markakis E, Telelis O (2015) Uniform price auctions: Equilibria and efficiency. *Theory of Computing Systems* 57(3):549–575.

Mas-Colell A (1977) Indivisible commodities and general equilibrium theory. *Journal of Economic Theory* 16(2):443–456.

Milgrom P, Watt M (2026) A walrasian mechanism with markups for nonconvex markets. *The Review of Economic Studies* 93(3):1995–2020, URL `http://dx.doi.org/10.1093/restud/rdaf067`.

Nguyen T, Vohra R (2024) (near-) substitute preferences and equilibria with indivisibilities. *Journal of Political Economy* 132(12):4122–4154, URL `http://dx.doi.org/10.1086/731413`.

O'Neill RP, Sotkiewicz PM, Hobbs BF, Rothkopf MH, Stewart Jr WR (2005) Efficient market-clearing prices in markets with nonconvexities. *European Journal of Operational Research* 164(1):269–285, URL `http://dx.doi.org/10.1016/j.ejor.2003.12.011`.

Paes Leme R (2017) Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106:294–316.

Paes Leme R, Wong SCw (2020) Computing walrasian equilibria: Fast algorithms and structural properties. *Mathematical Programming* 179(1):343–384.

Postlewaite A, Schmeidler D (1981) Approximate walrasian equilibria and nearby economies. *International Economic Review* 105–111.

Starr RM (1969) Quasi-equilibria in markets with non-convex preferences. *Econometrica* 25–38.

Sun N, Yang Z (2006) Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74(5):1385–1402.

Talluri KT, van Ryzin GJ (2004) *The Theory and Practice of Revenue Management* (Springer).

Appendix A: The Bicriteria FPTAS: Full Construction and Analysis

The bicriteria FPTAS is the computational core of the exact-clearing algorithm in Section 5. For any approximation precision $\varepsilon \in (0,1)$ and clearing slack $\delta \in (0,1)$, it returns a relaxation level, a price, and an integral allocation that satisfy the three guarantees in Proposition 4: a multiplicative approximation to τ^*, relaxed-demand feasibility in the original market, and a one-sided clearing shortfall of at most δm. This appendix gives the full construction and analysis of that subroutine.

The presentation follows the dependency structure of Algorithm 2. Appendix A.1 constructs the rounded demand correspondences by truncating and rounding values and then rounding quantities on local grids, and proves the two rounding guarantees used in the final analysis. Appendix A.2 reduces the continuous price space to a finite candidate list and proves the MCSS guarantee used by the inner clearing test. Appendix A.3 combines these results with the geometric search over relaxation levels to prove Proposition 4. The algorithm itself is specified in Section 5.3 and is not repeated here. All complexity statements use the unit-cost value-oracle model specified in Section 5.

A.1. Preprocessing and Rounded Demand Correspondences

The preprocessing uses two internal parameters $\varepsilon_1, \varepsilon_2 \in (0, 1/2]$. Value truncation and value rounding compress each valuation into a monotone step function with a logarithmic number of plateaus. Quantity rounding places a sparse grid on those plateaus and restricts the possible endpoints of their demanded prefixes. The output is a compact representation of $\widehat{D}_i^\tau(p)$ as a union of $O(1+\varepsilon_1^{-1}\log m)$ integer intervals for every buyer and every fixed (τ, p).

Value truncation. For each buyer, query $F_i := f_i(m)$ and set

$$\rho_i := \frac{F_i}{m^2}.$$

If $F_i = 0$, monotonicity and nonnegativity imply $f_i \equiv 0$; set $f_i^{\mathrm{tr}} = f_i$. If $F_i > 0$, define the truncated valuation implicitly by

$$f_i^{\mathrm{tr}}(k) := \begin{cases} 0, & f_i(k) < \rho_i, \\ f_i(k), & f_i(k) \geq \rho_i. \end{cases}$$

The strict inequality retains every value equal to ρ_i. In both branches, f_i^{tr} is normalized and monotone, and the value at m is F_i.

The following restatement of Lemma 4 records the effect of truncation on relaxed demand. The proof establishes a slightly more general threshold result that will also make clear why the cutoff F_i/m^2 preserves demand at every relaxation level relevant to (TR).

LEMMA 4 **(Demand preservation under truncation)**. *For every buyer i, price $p \geq 0$, and $\tau \in (0,1]$, $D_{i,\mathrm{tr}}^\tau(p) \subseteq D_i^\tau(p)$, with equality for every $\tau \in [1/m, 1]$.*

Proof. Fix a normalized, nonnegative, monotone valuation $f : \{0, \dots, m\} \to \mathbb{R}_{\geq 0}$ and write $F \coloneqq f(m)$. For $\kappa \in (0, 1]$, set

$$\rho \coloneqq \frac{\kappa F}{m}$$

and define

$$f^{\mathrm{tr}}(k) \coloneqq \begin{cases} 0, & f(k) < \rho, \\ f(k), & f(k) \geq \rho. \end{cases}$$

When $F = 0$, set $f^{\mathrm{tr}} = f$. For $\tau \in (0, 1]$ and $p \geq 0$, let $D^{\tau}_{f}(p)$ and $D^{\tau}_{\mathrm{tr}}(p)$ denote the τ-RDCs under f and f^{tr}, respectively. We show

$$D^{\tau}_{\mathrm{tr}}(p) \subseteq D^{\tau}_{f}(p) \quad \text{for every } \tau \in (0, 1], \qquad D^{\tau}_{\mathrm{tr}}(p) = D^{\tau}_{f}(p) \quad \text{for } \tau \in [\kappa, 1].$$

If $F = 0$, monotonicity and nonnegativity give $f(k) = 0$ for every k, so both claims are immediate. Assume $F > 0$.

Removed benchmarks. Consider $k \geq 1$ with $f(k) < \rho$. Quantity m is retained because $\rho \leq F$, and the strict truncation rule gives $k \leq m - 1$. For every $\tau \in (0, 1]$ and $p \geq 0$,

$$\tau f(k) - kp \leq \max\{0, \tau F - mp\}. \tag{14}$$

If $p \geq \tau f(k)/k$, the left-hand side is nonpositive. Otherwise,

$$\begin{aligned} (\tau F - mp) - (\tau f(k) - kp) &= \tau(F - f(k)) - (m - k)p \\ &> \tau(F - f(k)) - \frac{(m - k)\tau f(k)}{k} \\ &= \frac{\tau}{k}\big(kF - mf(k)\big) > 0. \end{aligned}$$

The final inequality follows from $mf(k) < \kappa F \leq F \leq kF$, proving (14).

Unconditional inclusion. Fix $x \in D^{\tau}_{\mathrm{tr}}(p)$. We first claim that

$$f^{\mathrm{tr}}(x) = f(x). \tag{15}$$

The claim holds at $x = 0$. If $x > 0$ and $f^{\mathrm{tr}}(x) = 0 < f(x)$, comparison with quantity 0 in the truncated correspondence gives $-xp \geq 0$, hence $p = 0$. Comparison with quantity m then gives $0 \geq \tau F$, a contradiction. If $f(x) = 0$, truncation leaves its value unchanged. Thus (15) holds, including at $p = 0$.

Fix a benchmark k. If $f^{\mathrm{tr}}(k) = f(k)$, its defining inequality transfers directly. If truncation changes its value, then $0 < f(k) < \rho$. Comparisons with quantities 0 and m, followed by (14), give

$$\begin{aligned} f(x) - xp = f^{\mathrm{tr}}(x) - xp \\ &\geq \max\{0, \tau F - mp\} \\ &\geq \tau f(k) - kp. \end{aligned}$$

This holds for every k, so $x \in D_f^\tau(p)$.

Reverse inclusion for $\tau \geq \kappa$. Let $\tau \in [\kappa, 1]$ and $x \in D_f^\tau(p)$. If $x > 0$, comparison with quantity 0 gives $p \leq f(x)/x$, while comparison with quantity m gives

$$f(x) + (m - x)p \geq \tau F.$$

Consequently,

$$\tau F \leq f(x) + (m - x)p \leq \frac{m}{x} f(x),$$

and therefore

$$f(x) \geq \frac{\tau x}{m} F \geq \frac{\kappa F}{m} = \rho.$$

Every positive quantity demanded under f is retained, and quantity 0 is unchanged. Since $f^{\mathrm{tr}}(k) \leq f(k)$ for every benchmark k,

$$f^{\mathrm{tr}}(x) - xp = f(x) - xp \geq \tau f(k) - kp \geq \tau f^{\mathrm{tr}}(k) - kp.$$

Hence $x \in D_{\mathrm{tr}}^\tau(p)$. Together with the unconditional inclusion, this proves equality for $\tau \in [\kappa, 1]$. Setting $\kappa = 1/m$ gives $\rho = F/m^2$ and proves the lemma. □

Value rounding and plateau representation. Let $\alpha := 1 + \varepsilon_1$ and

$$J := 1 + \lceil 2 \log_\alpha m \rceil = O(1 + \varepsilon_1^{-1} \log m).$$

For $F_i > 0$, define

$$\mathcal{V}_i := \{0\} \bigcup \{F_i \alpha^{-j} : 0 \leq j \leq J - 1\};$$

for $F_i = 0$, set $\mathcal{V}_i = \{0\}$. Round the truncated valuation down on this buyer-specific grid:

$$\widetilde{f}_i(k) := \max\{v \in \mathcal{V}_i : v \leq f_i^{\mathrm{tr}}(k)\}, \qquad i \in [n], \quad k \in \{0, \ldots, m\}. \tag{16}$$

The rounded valuations are normalized and monotone, and they satisfy

$$\frac{f_i^{\mathrm{tr}}(k)}{\alpha} \leq \widetilde{f}_i(k) \leq f_i^{\mathrm{tr}}(k) \qquad i \in [n], \quad k \in \{0, \ldots, m\}. \tag{17}$$

To verify the lower bound, suppose $f_i^{\mathrm{tr}}(k) > 0$ and define

$$j^* := \left\lceil \log_\alpha \frac{F_i}{f_i^{\mathrm{tr}}(k)} \right\rceil.$$

Every retained positive value lies in $[\rho_i, F_i]$, so $j^* \leq \lceil 2 \log_\alpha m \rceil = J - 1$. The grid value $F_i \alpha^{-j^*}$ therefore appears in $\mathcal{V}_i$ and lies between $f_i^{\mathrm{tr}}(k)/\alpha$ and $f_i^{\mathrm{tr}}(k)$. When the truncated value is zero, both sides of the lower bound are zero.

The graph of each $\widetilde{f}_i$ is a step function with at most $J+1$ *plateaus* that tile $\{0,\dots,m\}$ by closed integer intervals: there are integers

$$0 = a_{i,0} \le b_{i,0} < a_{i,1} \le b_{i,1} < \cdots < a_{i,L_i} \le b_{i,L_i} = m, \qquad a_{i,j+1} = b_{i,j} + 1, \qquad L_i \le J,$$

such that $\widetilde{f}_i$ takes the constant value $v_{i,j} \coloneqq \widetilde{f}_i(a_{i,j})$ on $I_{i,j} \coloneqq [a_{i,j}, b_{i,j}] \bigcap \mathbb{Z}$. If $F_i = 0$, record

$$L_i = 0, \qquad I_{i,0} = [0,m] \bigcap \mathbb{Z}, \qquad v_{i,0} = 0,$$

and perform no boundary search.

For $F_i > 0$, a binary search on the monotone predicate $f_i(k) < \rho_i$ finds the right endpoint of the zero plateau. For every positive grid level $v \in \mathcal{V}_i$, a binary search on $f_i(k) < \alpha v$ finds the associated right endpoint, after which repeated endpoints are discarded and m is recorded as the final endpoint. The raw-oracle predicate is valid for every positive level. At the lowest level,

$$\alpha v = F_i \alpha^{-(J-2)} > \frac{F_i}{m^2} = \rho_i,$$

and the inequality also holds at every higher level. Thus truncation does not change whether a value is below αv. Charging every searched grid level, rather than only the distinct plateaus found, gives $O(J \log m)$ boundary queries per positive buyer. Including the n full-supply queries, the preprocessing uses

$$O\big(n(1 + J \log m)\big)$$

value queries.

Quantity grids. Let $\varepsilon_2 \in (0, \frac{1}{2}]$ and $M_0 \coloneqq \lceil 1/\varepsilon_2 \rceil$. The quantity grid is

$$\mathcal{G} \coloneqq \{0,1,\dots,\min\{m, M_0\}\} \bigcup \big\{ \lceil (1+\varepsilon_2)^t \rceil : t \in \mathbb{Z}_{\ge 0},\ M_0 < \lceil (1+\varepsilon_2)^t \rceil < m \big\} \bigcup \{m\}, \tag{18}$$

combining a dense prefix with a geometric tail; the dense prefix avoids creating empty integer intervals at small quantities. Note $0, m \in \mathcal{G}$ and

$$|\mathcal{G}| = O\big(\varepsilon_2^{-1}(1 + \log m)\big)\,.$$

For each buyer i and plateau index $j \in \{0,1,\dots,L_i\}$, the *local grid* is

$$\mathcal{G}_{i,j} \coloneqq \big(\mathcal{G} \bigcap I_{i,j}\big) \bigcup \{a_{i,j},\, b_{i,j}\}, \tag{19}$$

i.e., the restriction of $\mathcal{G}$ to the plateau, augmented by the plateau endpoints; in particular $m = b_{i,L_i} \in \mathcal{G}_{i,L_i}$. The local grids satisfy $\sum_j |\mathcal{G}_{i,j}| = O(|\mathcal{G}| + J)$ for every i. A useful consequence of the geometric tail is the following gap bound.

LEMMA 7 **(Consecutive-gap of quantity grid)**. *For any two consecutive elements $g < g'$ of $\mathcal{G}$, $g' - g \le \varepsilon_2 g + 1$. The same bound holds within every local grid $\mathcal{G}_{i,j}$.*

Proof. We prove the bound for the global grid $\mathcal{G}$; the local-grid case follows because every $\mathcal{G}_{i,j}$ either inherits a pair of consecutive elements directly from $\mathcal{G}$, or one of the two is a plateau endpoint $a_{i,j}$ or $b_{i,j}$, in which case the bound follows by sandwiching against the surrounding $\mathcal{G}$-element, as detailed at the end of this proof.

For $\mathcal{G}$: fix consecutive $g < g'$. If $g' = g + 1$, the bound $g' - g \le \varepsilon_2 g + 1$ is immediate. So assume $g' \ge g + 2$. Since $\{0, 1, \ldots, \min(m, M_0)\} \subseteq \mathcal{G}$ consists of consecutive integers, $g \ge M_0$ (otherwise $g + 1 \in \mathcal{G}$ contradicts $g' \ge g + 2$). In particular $\varepsilon_2 g \ge \varepsilon_2 M_0 \ge 1$. Recall $g, g' \le m$ by (18).

Let $s \coloneqq \min\{t \ge 0 : \lceil (1+\varepsilon_2)^t \rceil > g\}$. Since $\lceil (1+\varepsilon_2)^0 \rceil = 1 \le M_0 \le g$, $s \ge 1$, so $(1+\varepsilon_2)^{s-1} \le \lceil (1+\varepsilon_2)^{s-1} \rceil \le g$ by minimality. Two cases.

Case A: $\lceil (1+\varepsilon_2)^s \rceil \le m$. Then $\lceil (1+\varepsilon_2)^s \rceil \in \mathcal{G}$ and exceeds g, so $g' \le \lceil (1+\varepsilon_2)^s \rceil \le (1+\varepsilon_2)^s + 1 \le (1+\varepsilon_2)g + 1$, i.e., $g' - g \le \varepsilon_2 g + 1$.

Case B: $\lceil (1+\varepsilon_2)^s \rceil > m$. Then $(1+\varepsilon_2)^s > m$, and by the minimality of s, no geometric grid point lies in $(g, m]$. Since $m \in \mathcal{G}$ and $m > g$, $g' = m < (1+\varepsilon_2)^s = (1+\varepsilon_2)(1+\varepsilon_2)^{s-1} \le (1+\varepsilon_2)g$, so $g' - g < \varepsilon_2 g$.

For the local case: fix i, j and consecutive $g < g'$ in $\mathcal{G}_{i,j}$. By construction, every element of $\mathcal{G}_{i,j}$ either lies in $\mathcal{G}$ or is one of the plateau endpoints $\{a_{i,j}, b_{i,j}\}$. Set $c \coloneqq \max\{q \in \mathcal{G} : q \le g\}$ and $d \coloneqq \min\{q \in \mathcal{G} : q \ge g'\}$; both exist because $0, m \in \mathcal{G}$. Then $c \le g < g' \le d$, so $c < d$. We claim c, d are consecutive in $\mathcal{G}$: any $q \in \mathcal{G}$ with $c < q < d$ must satisfy $q > g$ (else it contradicts the maximality of c) and $q < g'$ (else it contradicts the minimality of d); but $q \in \mathcal{G} \bigcap (g, g') \subseteq \mathcal{G} \bigcap I_{i,j} \subseteq \mathcal{G}_{i,j}$, contradicting that g, g' are consecutive in $\mathcal{G}_{i,j}$. Applying the bound for $\mathcal{G}$ to the consecutive pair c, d, $g' - g \le d - c \le \varepsilon_2 c + 1 \le \varepsilon_2 g + 1$. □

Rounded demand correspondences. The construction has three stages. It truncates f_i to f_i^{tr}, rounds the retained values to the plateau function $\widetilde{f}_i$, and then rounds the right endpoint of each demanded plateau interval down to the local grid $\mathcal{G}_{i,j}$. Fix $\tau \in (0, 1]$ and $p \ge 0$, and write

$$\widetilde{D}_i^{\tau}(p) \coloneqq \left\{k \in \{0, 1, \ldots, m\} : \widetilde{f}_i(k) - kp \ge \tau\, \widetilde{f}_i(k') - k'p,\ \forall\ k' \in \{0, 1, \ldots, m\}\right\} \tag{20}$$

for the *value-rounded demand correspondence*, the τ-RDC under $\widetilde{f}_i$. Because $\widetilde{f}_i$ is constant on each plateau, $\widetilde{D}_i^{\tau}(p)$ is a union of contiguous integer intervals indexed by the set of *accepted plateaus*

$$J_i^{\tau}(p) \coloneqq \left\{j \in \{0, 1, \ldots, L_i\} : I_{i,j} \bigcap \widetilde{D}_i^{\tau}(p) \ne \emptyset\right\}. \tag{21}$$

For each accepted plateau $j \in J_i^{\tau}(p)$, round the right endpoint of the plateau interval down to the local grid by setting

$$\widehat{k}_{i,j}^{\tau,+}(p) \coloneqq \max\left\{g \in \mathcal{G}_{i,j} : g \in \widetilde{D}_i^{\tau}(p)\right\}. \tag{22}$$

The maximum exists and lies in $\mathcal{G}_{i,j}$ because $a_{i,j} \in \mathcal{G}_{i,j} \bigcap \widetilde{D}_i^\tau(p)$ for every $j \in J_i^\tau(p)$ (on a plateau the rounded value is constant, so utility is non-increasing in the quantity and the left endpoint $a_{i,j}$ attains the plateau's maximum utility). The *rounded relaxed demand correspondence* is then

$$\widehat{D}_i^\tau(p) \coloneqq \bigcup_{j \in J_i^\tau(p)} \big[a_{i,j},\ \widehat{k}_{i,j}^{\tau,+}(p)\big] \bigcap \mathbb{Z}. \tag{23}$$

By construction, $\widehat{D}_i^\tau(p)$ is a union of at most $L_i + 1 = O(J)$ disjoint integer intervals, one per accepted plateau, with both endpoints lying in $\mathcal{G}_{i,j}$; because $m = b_{i,L_i} \in \mathcal{G}_{i,L_i}$, it can represent full supply. Truncation and the two rounding steps reduce the description at any (τ, p) to $O(J)$ integer intervals rather than potentially $\Theta(m)$ individual integers. The construction is illustrated in Figure 2.

The definition (23) of $\widehat{D}_i^\tau(p)$ uses only the value-rounded demand correspondence $\widetilde{D}_i^\tau(p)$ and the rounded endpoint $\widehat{k}_{i,j}^{\tau,+}(p)$. The proofs below additionally use the following two auxiliary quantities. For $\tau \in (0,1]$, $p \geq 0$, $i \in [n]$, and $j \in \{0, 1, \ldots, L_i\}$, set

$$\widetilde{F}_i^\tau(p) \coloneqq \max_{j' \in \{0,\ldots,L_i\}} \big(\tau v_{i,j'} - a_{i,j'}\, p\big), \tag{24}$$

where $\widetilde{F}_i^\tau(p)$ is distinct from the full-supply value $F_i = f_i(m)$ used in the truncation threshold, and, for $j \in J_i^\tau(p)$ and $p > 0$, define

$$k_{i,j}^{\tau,+}(p) \coloneqq \min\Big(b_{i,j},\ \big\lfloor (v_{i,j} - \widetilde{F}_i^\tau(p))/p \big\rfloor\Big). \tag{25}$$

At $p = 0$, set $\widetilde{F}_i^\tau(0) \coloneqq \tau v_{i,L_i}$ and $k_{i,j}^{\tau,+}(0) \coloneqq b_{i,j}$. With these definitions, the value-rounded demand correspondence $\widetilde{D}_i^\tau(p)$ defined in (20) admits the equivalent representation

$$\widetilde{D}_i^\tau(p) \;=\; \bigcup_{j \in J_i^\tau(p)} \big[a_{i,j},\, k_{i,j}^{\tau,+}(p)\big] \bigcap \mathbb{Z}, \tag{26}$$

and a plateau j belongs to $J_i^\tau(p)$ if and only if $v_{i,j} - a_{i,j}\, p \geq \widetilde{F}_i^\tau(p)$. The rounded endpoint defined in (22) is then equivalently $\widehat{k}_{i,j}^{\tau,+}(p) = \max\{g \in \mathcal{G}_{i,j} : g \leq k_{i,j}^{\tau,+}(p)\}$.

The relaxation-level search uses the following monotonicity property of the rounded correspondences.

LEMMA 8 **(Nesting of rounded RDCs).** *For every $i \in [n]$, $p \geq 0$, and $0 < \tau' \leq \tau \leq 1$,*

$$\widehat{D}_i^\tau(p) \subseteq \widehat{D}_i^{\tau'}(p).$$

Proof. Nonnegativity of $\widetilde{f}_i$ implies $\widetilde{D}_i^\tau(p) \subseteq \widetilde{D}_i^{\tau'}(p)$: lowering the relaxation level weakens every inequality that defines the value-rounded demand correspondence. Hence every plateau accepted at τ is accepted at τ', and $\widehat{k}_{i,j}^{\tau,+}(p) \leq \widehat{k}_{i,j}^{\tau',+}(p)$ for each $j \in J_i^\tau(p)$. Each interval in $\widehat{D}_i^\tau(p)$ is therefore contained in the corresponding interval in $\widehat{D}_i^{\tau'}(p)$. □

Guarantees of rounding. Part (a) below shows that an allocation selected from the rounded demand correspondences is feasible for the original demand correspondences after scaling τ by $1/\alpha$. Part (b) rounds an optimal solution of (TR) while losing at most an ε_2 fraction of the supply.

LEMMA 5 **(Guarantees of rounding)**. *We have the following guarantees:*

(a) (Original-demand feasibility) *For any $i \in [n]$, $\tau \in (0,1]$, $p \geq 0$, and $\widehat{x}_i \in \widehat{D}_i^{\tau}(p)$, we have $\widehat{x}_i \in D_i^{\tau/\alpha}(p)$.*

(b) (Rounding an optimal solution) *Let $(\tau^*, p^*, \boldsymbol{x}^*)$ be an optimal solution to* (TR)*, and set $\widetilde{\tau} \coloneqq \tau^*/\alpha$ and $\widetilde{p} \coloneqq p^*/\alpha$. Then there exists an allocation $\widehat{\boldsymbol{x}}$ such that $\widehat{x}_i \in \widehat{D}_i^{\widetilde{\tau}}(\widetilde{p})$ for every $i \in [n]$ and*

$$(1-\varepsilon_2)\, m \;\leq\; \sum_{i=1}^{n} \widehat{x}_i \;\leq\; m.$$

Proof. *Part (a).* Fix i, τ, p, and $\widehat{x}_i \in \widehat{D}_i^{\tau}(p)$. Let j' be the plateau containing an arbitrary benchmark $k \in \{0, 1, \ldots, m\}$, so $\widetilde{f}_i(k) = v_{i,j'}$. There is an accepted plateau $j \in J_i^{\tau}(p)$ such that $\widehat{x}_i \in [a_{i,j}, \widehat{k}_{i,j}^{\tau,+}(p)]$, and hence $\widetilde{f}_i(\widehat{x}_i) = v_{i,j}$.

Suppose first that $p > 0$. The definitions of $k_{i,j}^{\tau,+}(p)$ and $\widehat{k}_{i,j}^{\tau,+}(p)$ give

$$\widehat{x}_i \leq \widehat{k}_{i,j}^{\tau,+}(p) \leq k_{i,j}^{\tau,+}(p) \leq \frac{v_{i,j} - \widetilde{F}_i^{\tau}(p)}{p}.$$

Therefore,

$$\begin{aligned} \widetilde{f}_i(\widehat{x}_i) - \widehat{x}_i p &= v_{i,j} - \widehat{x}_i p \\ &\geq \widetilde{F}_i^{\tau}(p) \\ &\geq \tau v_{i,j'} - a_{i,j'} p \\ &\geq \tau v_{i,j'} - kp = \tau \widetilde{f}_i(k) - kp, \end{aligned} \tag{27}$$

where the last inequality uses $k \geq a_{i,j'}$. If $p = 0$, the boundary convention and $j \in J_i^{\tau}(0)$ give

$$\widetilde{f}_i(\widehat{x}_i) = v_{i,j} \geq \tau v_{i,L_i} \geq \tau v_{i,j'} = \tau \widetilde{f}_i(k),$$

so the same value-rounded demand inequality holds.

Using (17), for every benchmark k we obtain

$$\begin{aligned} f_i^{\mathrm{tr}}(\widehat{x}_i) - \widehat{x}_i p &\geq \widetilde{f}_i(\widehat{x}_i) - \widehat{x}_i p \\ &\geq \tau \widetilde{f}_i(k) - kp \\ &\geq \frac{\tau}{\alpha} f_i^{\mathrm{tr}}(k) - kp. \end{aligned}$$

Thus $\widehat{x}_i \in D_{i,\mathrm{tr}}^{\tau/\alpha}(p)$. The unconditional inclusion in Lemma 4 then yields $\widehat{x}_i \in D_i^{\tau/\alpha}(p)$.

Part (b). Set $\widetilde{\tau} \coloneqq \tau^*/\alpha$ and $\widetilde{p} \coloneqq p^*/\alpha$. For each i, let $j(i)$ denote the plateau index of x_i^*. Fix i. Lemma 2 gives $\tau^* \geq 1/m$, so Lemma 4 gives

$$D_i^{\tau^*}(p^*) = D_{i,\mathrm{tr}}^{\tau^*}(p^*).$$

Consequently,

$$f_i^{\mathrm{tr}}(x_i^*) - x_i^* p^* \geq \tau^* f_i^{\mathrm{tr}}(k) - k p^* \qquad \text{for every } k. \tag{28}$$

If $x_i^* > 0$, the reverse-inclusion argument in the proof of Lemma 4 also gives $f_i(x_i^*) \geq \rho_i$. Thus every positive optimal quantity is retained by the truncation.

We first place x_i^* in the value-rounded correspondence. Dividing Equation (28) by α and using the value bounds gives, for every k,

$$\begin{aligned} \widetilde{f}_i(x_i^*) - x_i^* \widetilde{p} &\geq \frac{f_i^{\mathrm{tr}}(x_i^*)}{\alpha} - x_i^* \widetilde{p} \\ &\geq \widetilde{\tau} f_i^{\mathrm{tr}}(k) - k\widetilde{p} \\ &\geq \widetilde{\tau} \widetilde{f}_i(k) - k\widetilde{p}. \end{aligned} \tag{29}$$

Because $\widetilde{f}_i(x_i^*) = v_{i,j(i)}$, $x_i^* \geq a_{i,j(i)}$, and $\widetilde{p} \geq 0$, specializing to $k = a_{i,j'}$ yields

$$\begin{aligned} v_{i,j(i)} - a_{i,j(i)}\widetilde{p} &\geq v_{i,j(i)} - x_i^* \widetilde{p} \\ &\geq \widetilde{\tau} v_{i,j'} - a_{i,j'}\widetilde{p} \qquad \text{for every } j'. \end{aligned} \tag{30}$$

Maximizing the right-hand side over j' shows that $j(i) \in J_i^{\widetilde{\tau}}(\widetilde{p})$. Equation (29) also gives

$$x_i^* \widetilde{p} \leq v_{i,j(i)} - \widetilde{F}_i^{\widetilde{\tau}}(\widetilde{p}).$$

For $\widetilde{p} > 0$, integrality and $x_i^* \leq b_{i,j(i)}$ imply $x_i^* \leq k_{i,j(i)}^{\widetilde{\tau},+}(\widetilde{p})$. For $\widetilde{p} = 0$, this conclusion follows from the convention $k_{i,j(i)}^{\widetilde{\tau},+}(0) = b_{i,j(i)}$. Hence

$$x_i^* \in [a_{i,j(i)}, k_{i,j(i)}^{\widetilde{\tau},+}(\widetilde{p})] \subseteq \widetilde{D}_i^{\widetilde{\tau}}(\widetilde{p}).$$

Since $a_{i,j(i)} \in \mathcal{G}_{i,j(i)}$, the rounded endpoint is at least $a_{i,j(i)}$. Define $\widehat{x}_i := \min\big(x_i^*, \widehat{k}_{i,j(i)}^{\widetilde{\tau},+}(\widetilde{p})\big)$. Then

$$\widehat{x}_i \in [a_{i,j(i)}, \widehat{k}_{i,j(i)}^{\widetilde{\tau},+}(\widetilde{p})] \subseteq \widehat{D}_i^{\widetilde{\tau}}(\widetilde{p}).$$

It remains to bound the rounding loss. If $\widehat{x}_i = x_i^*$, the loss is zero. Otherwise, let g' be the smallest element of $\mathcal{G}_{i,j(i)}$ strictly larger than $\widehat{x}_i$. Such an element exists because $x_i^* \leq b_{i,j(i)} = \max \mathcal{G}_{i,j(i)}$. The maximality of the rounded endpoint gives

$$g' > k_{i,j(i)}^{\widetilde{\tau},+}(\widetilde{p}) \geq x_i^*,$$

so integrality implies $x_i^* \leq g' - 1$. Lemma 7 then yields

$$x_i^* - \widehat{x}_i \leq (g' - \widehat{x}_i) - 1 \leq \varepsilon_2 \widehat{x}_i \leq \varepsilon_2 x_i^*.$$

Summing the buyer-level losses gives

$$\sum_i (x_i^* - \widehat{x}_i) \leq \varepsilon_2 \sum_i x_i^* = \varepsilon_2 m.$$

Together with $\widehat{x}_i \le x_i^*$ and $\sum_i x_i^* = m$, this proves

$$(1-\varepsilon_2)m \le \sum_i \widehat{x}_i \le m.$$

□

The preprocessing therefore replaces each original valuation by a compact plateau representation and restricts every rounded demand endpoint to a local grid. Once (τ, p) is fixed, each $\widehat{D}_i^\tau(p)$ is described by $O(J)$ integer intervals and can be formed from the stored plateau and grid data without further value queries. The next subsection uses this representation to identify the finitely many prices at which a rounded demand correspondence can change and to solve the resulting allocation problem.

A.2. Finite-Price Reduction and the MCSS Approximation Algorithm

Fix a relaxation level $\tau \in (0,1]$. The construction in Appendix A.1 gives a compact interval representation of every rounded demand correspondence $\widehat{D}_i^\tau(p)$. The inner problem now has two parts. We first replace the continuous price space by a finite list without excluding any allocation feasible for the rounded correspondences. We then solve, at each listed price, the problem of selecting one quantity per buyer while keeping the total allocation below the supply.

Finite-price reduction. For $k \in I_{i,j}$, the equivalent representation in (26) gives

$$k \in \widehat{D}_i^\tau(p) \quad \Longleftrightarrow \quad \exists g \in \mathcal{G}_{i,j} :\ k \le g \text{ and } v_{i,j} - gp \ge \max_\ell \{\tau v_{i,\ell} - a_{i,\ell} p\}. \tag{31}$$

Thus, a rounded demand correspondence can change only when the active affine function in the maximum changes or when one of the local-grid inequalities becomes tight. This yields exactly the list introduced in (11).

Definition 6 (Candidate-price list). For fixed $\tau \in (0,1]$, define

$$\begin{aligned}\widetilde{\mathcal{L}}^\tau := \{0\} &\bigcup \Big\{ p \ge 0 : \tau v_{i,j} - a_{i,j} p = \tau v_{i,\ell} - a_{i,\ell} p,\ i \in [n],\ 0 \le j < \ell \le L_i \Big\} \\ &\bigcup \Big\{ p \ge 0 : v_{i,j} - gp = \tau v_{i,\ell} - a_{i,\ell} p,\ i \in [n],\ j,\ell \in \{0,\dots,L_i\},\ g \in \mathcal{G}_{i,j},\ g \ne a_{i,\ell} \Big\}.\end{aligned} \tag{32}$$

The first family contains all possible breakpoints of the upper envelope $\widetilde{F}_i^\tau$. The second contains all prices at which a local-grid inequality in (31) can change status. If $g = a_{i,\ell}$, the two sides have the same slope, so their ordering is constant in p and no event price is needed. The list may contain duplicates and equalities that do not change a rounded demand correspondence; these harmless extra entries simplify its generation.

Lemma 6 **(Finite-price reduction).** *Fix $\tau \in (0,1]$.*

(a) *The list $\widetilde{\mathcal{L}}^\tau$ has $O(nJ^2 + nJ|\mathcal{G}|)$ entries and can be generated in the same time by direct enumeration.*

(b) *If an allocation* $\boldsymbol{x}$ *satisfies* $x_i \in \widehat{D}_i^\tau(p)$ *for every* i *at some price* $p \geq 0$*, then there exists* $p' \in \widetilde{\mathcal{L}}^\tau$ *such that* $x_i \in \widehat{D}_i^\tau(p')$ *for every* i*.*

Proof. *Candidate-list size and generation.* Fix a buyer i. The first family in Definition 6 has $O(J^2)$ pairs (j,ℓ). Because $a_{i,j} \neq a_{i,\ell}$ for $j \neq \ell$, each equality has at most one solution,

$$(a_{i,\ell} - a_{i,j})p = \tau(v_{i,\ell} - v_{i,j}),$$

which is retained when it is nonnegative.

For the second family, fix a plateau j and a grid point $g \in \mathcal{G}_{i,j}$. There are $O(J)$ choices of ℓ. Whenever $g \neq a_{i,\ell}$, the equality has the unique solution

$$(a_{i,\ell} - g)p = \tau v_{i,\ell} - v_{i,j},$$

again retained only when nonnegative. Appendix A.1 established that $\sum_j |\mathcal{G}_{i,j}| = O(|\mathcal{G}| + J)$, so the second family contributes

$$O\left(J \sum_j |\mathcal{G}_{i,j}|\right) = O(J|\mathcal{G}| + J^2)$$

prices for buyer i. Summing the two bounds over all buyers gives $O(nJ^2 + nJ|\mathcal{G}|)$ entries. Each entry is obtained by solving one linear equality, so direct enumeration has the same running time.

Preservation of feasible allocations. Fix $p \geq 0$ and $\widehat{\boldsymbol{x}}$ as in the statement. If $p \in \widetilde{\mathcal{L}}^\tau$, take $p' = p$. Otherwise, let p^- be the largest distinct value in $\widetilde{\mathcal{L}}^\tau$ that is strictly below p. This value exists because zero belongs to the list and $p > 0$. By construction, the interval $(p^-, p]$ contains no candidate price.

Fix a buyer i. By (31), there are a plateau j and a grid point $g \in \mathcal{G}_{i,j}$ such that $\widehat{x}_i \leq g$ and

$$v_{i,j} - gp \geq \tau v_{i,\ell} - a_{i,\ell}p \qquad \text{for every } \ell. \tag{33}$$

For each ℓ, define the affine function

$$h_\ell(q) \coloneqq (v_{i,j} - gq) - (\tau v_{i,\ell} - a_{i,\ell}q).$$

If $g = a_{i,\ell}$, then h_ℓ is constant, and Equation (33) gives $h_\ell(p^-) \geq 0$. Suppose instead that $g \neq a_{i,\ell}$. If $h_\ell(p^-) < 0$, then $h_\ell(p) \geq 0$ and continuity give a root $q \in (p^-, p]$. This root belongs to the second family in Definition 6, contradicting the choice of p^-. Hence $h_\ell(p^-) \geq 0$ in both cases. Because this holds for every ℓ,

$$v_{i,j} - gp^- \geq \max_\ell \{\tau v_{i,\ell} - a_{i,\ell}p^-\}.$$

The same grid point g, together with $\widehat{x}_i \leq g$, therefore witnesses $\widehat{x}_i \in \widehat{D}_i^\tau(p^-)$ through Equation (31). Applying the argument to every buyer proves that the same allocation is feasible at $p' = p^-$. □

Lemma 6 therefore reduces the continuous price space to a finite list. We next solve the allocation problem induced by a fixed candidate price.

MCSS Approximation Algorithm. At fixed (τ, p), set $S_i \coloneqq \widehat{D}_i^{\tau}(p)$ and $T \coloneqq m$. Definition 5 then asks for the largest total allocation not exceeding supply. Proposition 5 gives the inner algorithm. We restate it and provide its full proof.

PROPOSITION 5 **(MCSS Approximation Algorithm)**. *Suppose that R is a positive integer and each $S_i \subseteq \{0, 1, \ldots, m\}$ is a union of at most R disjoint integer intervals. For any target $T \in \{0, 1, \ldots, m\}$ and $\delta_{\mathrm{mcss}} \in (0, 1)$, there is an algorithm that returns* INFEASIBLE *if and only if the feasible set in Definition 5 is empty. Otherwise, it returns a feasible allocation $\boldsymbol{x}$ satisfying*

$$M^* - \delta_{\mathrm{mcss}} m \le \sum_i x_i \le M^* \le T.$$

Its running time is $O\big((n^2 R/\delta_{\mathrm{mcss}}) \log(nR/\delta_{\mathrm{mcss}})\big)$. If $\delta_{\mathrm{mcss}} < 1/m$, the returned sum equals M^*.

Proof of Proposition 5. Write each $S_i = \bigcup_{r=1}^{R_i} [a_{i,r}, b_{i,r}] \bigcap \mathbb{Z}$ as a union of $R_i \le R$ disjoint integer intervals (with $R_i = 0$ when $S_i = \emptyset$). We give an additive approximation by a min-per-bucket dynamic program over feasible partial sums.

Feasibility test. With $\mathcal{F} \coloneqq \{\boldsymbol{x} \in \prod_i S_i : \sum_i x_i \le T\}$, and since $x_i = \min S_i$ attains $\sum_i \min S_i$ when every $S_i \neq \emptyset$, we have $\mathcal{F} = \emptyset$ iff some $S_i = \emptyset$ or $\sum_i \min S_i > T$, both checkable in $O(\sum_i R_i) = O(nR)$ time; the algorithm returns INFEASIBLE exactly in this case. Assume henceforth that $\mathcal{F} \neq \emptyset$.

The dynamic program. Set $\eta \coloneqq \delta_{\mathrm{mcss}} m/n$ and assign each integer $q \in \{0, 1, \ldots, T\}$ the bucket index $B(q) \coloneqq \lfloor q/\eta \rfloor$; the number of buckets meeting $\{0, \ldots, T\}$ is at most $\lfloor T/\eta \rfloor + 1 \le n/\delta_{\mathrm{mcss}} + 1$. Let $P_i \coloneqq \{\sum_{j \le i} x_j : x_j \in S_j \text{ for } j \le i,\ \sum_{j \le i} x_j \le T\}$ be the feasible partial sums after the first i buyers. We maintain a compressed state $\Phi_i \subseteq P_i$, defined inductively: $\Phi_0 \coloneqq \{0\}$, and given Φ_{i-1}, the transition forms $C_i \coloneqq \{\phi + s : \phi \in \Phi_{i-1},\ s \in S_i,\ \phi + s \le T\}$ and lets Φ_i hold, for each non-empty bucket of C_i, the *minimum* element of C_i in that bucket (with a backpointer). Compression may leave a non-empty bucket of P_i without a representative in Φ_i; the property the analysis relies on is only the downward-approximation invariant established below. After processing all n buyers, output the largest element of Φ_n and follow the backpointers to recover an $\boldsymbol{x} \in \prod_i S_i$ achieving it; then $\boldsymbol{x} \in \mathcal{F}$, which gives the membership $x_i \in S_i$ for every i.

Approximation analysis. We prove the *downward-approximation invariant*: for every i and every $s \in P_i$, there is $\phi \in \Phi_i$ with $\phi \le s$ and $s - \phi \le i\eta$. This holds for $i = 0$. Inductively, write $s = s' + x_i$ with $s' \in P_{i-1}$ and $x_i \in S_i$, and pick $\phi' \in \Phi_{i-1}$ with $\phi' \le s'$ and $s' - \phi' \le (i-1)\eta$. Then $c \coloneqq \phi' + x_i \le s \le T$, so $c \in C_i$; the bucket of c retains its minimum $\phi \le c \le s$ with $c - \phi < \eta$, whence $s - \phi = (s' - \phi') + (c - \phi) \le i\eta$. Applying the invariant to an optimal $\boldsymbol{x}^*$ with $\sum_i x_i^* = M^*$ gives $\phi \in \Phi_n$ with $\phi \ge M^* - n\eta = M^* - \delta_{\mathrm{mcss}} m$. The output is the largest element of Φ_n; it is an achievable sum at most T, hence lies in $\mathcal{F}$ and is at most M^*, so it lies in $[M^* - \delta_{\mathrm{mcss}} m,\ M^*]$.

Runtime analysis. Each Φ_{i-1} has $O(n/\delta_{\text{mcss}})$ elements, and each S_i is a union of $R_i \le R$ intervals. For $\phi \in \Phi_{i-1}$ and interval $[a_{i,r}, b_{i,r}]$, if $\phi + a_{i,r} \le T$, the feasible shifted candidates form the contiguous integer interval $[\phi + a_{i,r}, \min\{T, \phi + b_{i,r}\}]$; if $\phi + a_{i,r} > T$, the shifted interval is empty and is discarded. Thus C_i is a union of $O(|\Phi_{i-1}|R_i) = O(nR/\delta_{\text{mcss}})$ nonempty intervals. Sorting these by left endpoint and sweeping against the (non-integer) bucket boundaries $\{b\eta\}$ yields, for every non-empty bucket, the smallest integer it contains that is covered by some candidate interval, i.e. the bucket minimum of C_i; this produces Φ_i with backpointers in $O\big(\frac{nR}{\delta_{\text{mcss}}} \log \frac{nR}{\delta_{\text{mcss}}}\big)$ time per buyer, hence $O\big(\frac{n^2 R}{\delta_{\text{mcss}}} \log \frac{nR}{\delta_{\text{mcss}}}\big)$ over all n buyers. Reconstructing an explicit $\boldsymbol{x}$ that achieves the output sum follows the backpointers and adds no asymptotic overhead.

Exact case. If $\delta_{\text{mcss}} < 1/m$ then the additive slack $\delta_{\text{mcss}} m < 1$; since the output and M^* are integers with the output in $[M^* - \delta_{\text{mcss}} m,\, M^*]$, the output equals M^* and the algorithm is exact. Choosing $\delta_{\text{mcss}} = 1/(2m)$ makes the algorithm exact with running time $\text{poly}(n, m)$, since then $1/\delta_{\text{mcss}} = 2m$ and $R \le m + 1$. □

In the bicriteria construction, every set $S_i = \widehat{D}_i^{\tau}(p)$ is nonempty. Indeed, choose

$$j^\star \in \arg\max_j \{v_{i,j} - a_{i,j} p\}.$$

For every ℓ,

$$v_{i,j^\star} - a_{i,j^\star} p \ge v_{i,\ell} - a_{i,\ell} p \ge \tau v_{i,\ell} - a_{i,\ell} p,$$

where the second inequality uses $\tau \le 1$ and $v_{i,\ell} \ge 0$. Thus $j^\star \in J_i^{\tau}(p)$ and $a_{i,j^\star} \in \widehat{D}_i^{\tau}(p)$. Moreover, each S_i is a union of at most $J+1$ integer intervals. Proposition 5 therefore applies with $R \le J+1$ and $T = m$; in this application, infeasibility is equivalent to $\sum_i \min S_i > m$.

For each fixed τ, Lemma 6 reduces price search to finitely many instances, and Proposition 5 approximately maximizes the allocated quantity in every instance. The next subsection combines these components with the rounding guarantees and the outer search over relaxation levels.

A.3. Proof of the Bicriteria FPTAS

Algorithm 2 is fully specified in Section 5.3. This subsection does not repeat that construction. It combines the rounding guarantees established in Appendix A.1 with the finite-price reduction, the MCSS Approximation Algorithm, and the geometric search over relaxation levels to prove Proposition 4.

PROPOSITION 4 **(Bicriteria FPTAS)**. *There exists an algorithm* BICRITERIA*(ε,δ) (Algorithm 2) such that, for every $\varepsilon, \delta \in (0,1)$ and arbitrary monotone valuations $\{f_i\}_{i\in[n]}$ with $f_i(0) = 0$, it returns $\widehat{\tau} \in (0,1]$, $\widehat{p} \ge 0$, and $\widehat{\boldsymbol{x}} \in \{0,\dots,m\}^n$ satisfying*

(i) $\widehat{\tau} \ge (1-\varepsilon)\tau^*$;

(ii) $\widehat{x}_i \in D_i^{\widehat{\tau}}(\widehat{p})$ *for every* i*;*

(iii) $\sum_i \widehat{x}_i \in [(1-\delta)m, m]$.

The running time is polynomial in n, $\log m$, $1/\varepsilon$, *and* $1/\delta$.

Proof. Use the calibration in (12):

$$\varepsilon_1 = \frac{\varepsilon}{6}, \qquad \varepsilon_2 = \frac{\delta}{2}, \qquad \varepsilon_{\mathrm{bs}} = \frac{\varepsilon}{3}, \qquad \delta_{\mathrm{mcss}} = \frac{\delta}{2}, \qquad \alpha = 1+\varepsilon_1.$$

Let $(\tau^{\mathrm{acc}}, \widehat{p}, \widehat{\boldsymbol{x}})$ denote the tuple recorded by Algorithm 2, so the returned relaxation is $\widehat{\tau} = \tau^{\mathrm{acc}}/\alpha$.

Acceptance at and below τ^*/α. Lemma 2 gives $\tau^* \geq 1/m$, and the maximum in (TR) is attained. Fix an optimal solution $(\tau^*, p^*, \boldsymbol{x}^*)$ and write $\widetilde{p}^* := p^*/\alpha$. We show that every queried relaxation level $\lambda \leq \tau^*/\alpha$ is accepted.

Lemma 5(b) gives an allocation $\widehat{\boldsymbol{x}}^\circ$ satisfying

$$\widehat{x}_i^\circ \in \widehat{D}_i^{\tau^*/\alpha}(\widetilde{p}^*) \quad \text{for every } i, \qquad (1-\varepsilon_2)m \leq \sum_i \widehat{x}_i^\circ \leq m.$$

For $\lambda \leq \tau^*/\alpha$, Lemma 8 implies

$$\widehat{D}_i^{\tau^*/\alpha}(\widetilde{p}^*) \subseteq \widehat{D}_i^{\lambda}(\widetilde{p}^*),$$

so the same allocation is feasible at $(\lambda, \widetilde{p}^*)$. Lemma 6(b) identifies a price $p' \in \widetilde{\mathcal{L}}^\lambda$ at which this allocation remains feasible. At this candidate price, the MCSS optimum therefore satisfies

$$M^* \geq \sum_i \widehat{x}_i^\circ \geq (1-\varepsilon_2)m.$$

Proposition 5 returns an allocation $\boldsymbol{x}'$ with

$$\begin{aligned}\sum_i x_i' &\geq M^* - \delta_{\mathrm{mcss}} m \\ &\geq (1-\varepsilon_2-\delta_{\mathrm{mcss}})m = (1-\delta)m.\end{aligned}$$

The clearing test therefore accepts every queried $\lambda \leq \tau^*/\alpha$.

Acceptance of the smallest grid value. The smallest grid value satisfies

$$\begin{aligned}\lambda_{T_\lambda} &\leq \frac{1}{(1+\varepsilon_{\mathrm{bs}})m} \\ &\leq \frac{1}{\alpha m} \leq \frac{\tau^*}{\alpha}.\end{aligned}$$

The second inequality uses $\varepsilon_{\mathrm{bs}} \geq \varepsilon_1$, and the third uses $\tau^* \geq 1/m$. Hence the smallest grid value is accepted whenever it is queried. In particular, the algorithm always records an accepted tuple.

Approximation of τ^*. If $\lambda_0 = 1$ is accepted, then

$$\widehat{\tau} = \frac{1}{\alpha} \geq 1-\varepsilon_1 \geq (1-\varepsilon)\tau^*,$$

using $\tau^* \le 1$. Otherwise, the search terminates with adjacent indices $t_+ = t_- + 1$, where λ_{t_-} is rejected and $\tau^{\mathrm{acc}} = \lambda_{t_+}$ is accepted. Because every queried level at most τ^*/α is accepted, rejection implies

$$(1+\varepsilon_{\mathrm{bs}})\tau^{\mathrm{acc}} = \lambda_{t_-} > \frac{\tau^*}{\alpha}.$$

Consequently,

$$\begin{aligned}\widehat{\tau} = \frac{\tau^{\mathrm{acc}}}{\alpha} &> \frac{\tau^*}{\alpha^2(1+\varepsilon_{\mathrm{bs}})} \\ &\ge (1-\varepsilon_1)^2(1-\varepsilon_{\mathrm{bs}})\tau^* \\ &\ge (1-2\varepsilon_1-\varepsilon_{\mathrm{bs}})\tau^* \\ &= \left(1-\frac{2\varepsilon}{3}\right)\tau^* \ge (1-\varepsilon)\tau^*.\end{aligned}$$

Here we used $1/(1+a) \ge 1-a$ for $a \ge 0$ and discarded only nonnegative cross terms. This proves Proposition 4(i).

Demand feasibility and clearing. The recorded tuple comes from an accepting clearing test. Therefore,

$$\widehat{x}_i \in \widehat{D}_i^{\tau^{\mathrm{acc}}}(\widehat{p}) \qquad \text{for every } i,$$

and the MCSS target and acceptance rule give

$$(1-\delta)m \le \sum_i \widehat{x}_i \le m.$$

Lemma 5(a) yields

$$\widehat{x}_i \in D_i^{\tau^{\mathrm{acc}}/\alpha}(\widehat{p}) = D_i^{\widehat{\tau}}(\widehat{p}) \qquad \text{for every } i.$$

This proves parts (ii) and (iii).

Runtime. The preprocessing makes

$$O\big(n(1+J\log m)\big)$$

value queries and stores all plateau and grid data. No later step uses the value oracle. Let $G := |\mathcal{G}|$. The relaxation grid contains

$$T_\lambda + 1 = O(1+\varepsilon_{\mathrm{bs}}^{-1}\log m)$$

levels, and the search queries

$$Q_\lambda = O(1+\log T_\lambda) = O\big(\log\big(2+\varepsilon^{-1}\log m\big)\big)$$

of them.

For each queried level, Lemma 6(a) generates $O(nJ^2 + nJG)$ candidate prices in the same asymptotic time. At a fixed candidate price, forming all rounded RDCs requires $O(nJ\log(G+1))$ operations, using predecessor queries in the local grids. Proposition 5, with $R = O(J)$, then requires

$$O\left(\frac{n^2 J}{\delta_{\mathrm{mcss}}}\log\frac{nJ}{\delta_{\mathrm{mcss}}}\right)$$

operations. Thus the total number of arithmetic and comparison operations is bounded by

$$O\Bigg(T_\lambda + n(1+J\log m) + G + n(G+J) \\ + Q_\lambda(nJ^2+nJG)\left(nJ\log(G+1) + \frac{n^2 J}{\delta}\log\frac{nJ}{\delta}\right)\Bigg).$$

Finally,

$$J = O(1+\varepsilon^{-1}\log m), \qquad G = O\big(\delta^{-1}(1+\log m)\big), \qquad T_\lambda = O(1+\varepsilon^{-1}\log m).$$

Substitution proves that both the oracle-query count and the operation count are polynomial in n, $\log m$, $1/\varepsilon$, and $1/\delta$. □

Appendix B: Proof of Theorem 2

This appendix proves Theorem 2. Corollary 1 handles the small-supply regime, while Lemma 3 provides the completion step for the large-supply regime. We establish these two ingredients before combining them to prove the theorem.

B.1. Exact-Clearing Specialization and Completion

Small-supply exact specialization.

Corollary 1 **(Exact-clearing specialization)**. *Fix $\varepsilon \in (0,1)$. Then* Bicriteria*$(\varepsilon, 1/(2m))$ returns a $\widehat{\tau}$-WE satisfying $(1-\varepsilon)\tau^* \le \widehat{\tau} \le \tau^*$, in time* $\mathrm{poly}(n, m, 1/\varepsilon)$.

Proof. Apply Proposition 4 with clearing slack $\delta = 1/(2m)$. It returns $(\widehat{\tau}, \widehat{p}, \widehat{\boldsymbol{x}})$ satisfying

$$\widehat{\tau} \ge (1-\varepsilon)\tau^*, \qquad \widehat{x}_i \in D_i^{\widehat{\tau}}(\widehat{p}) \quad \text{for every } i, \qquad m - \frac{1}{2} \le \sum_i \widehat{x}_i \le m.$$

The total allocation is integral, so $\sum_i \widehat{x}_i = m$. Hence $(\widehat{p}, \widehat{\boldsymbol{x}})$ is a $\widehat{\tau}$-WE, and the triple $(\widehat{\tau}, \widehat{p}, \widehat{\boldsymbol{x}})$ is feasible for (TR). It follows that $\widehat{\tau} \le \tau^*$.

The running-time guarantee in Proposition 4 is polynomial in n, $\log m$, $1/\varepsilon$, and $1/\delta$. Substituting $1/\delta = 2m$ gives a running time polynomial in n, m, and $1/\varepsilon$. □

Multiplicative completion.

LEMMA 3 **(Multiplicative Completion)**. *Let each f_i be monotone with $f_i(0)=0$. Fix $\tau \in (0,1]$, $p \geq 0$, and an integral allocation $\boldsymbol{x}$ with $x_i \in D_i^{\tau}(p)$ for every i. Write $s := \sum_i x_i$, $q := |\{i : x_i > 0\}|$, and $c := s/(m+q-1)$, and assume $1 \leq s < m$. In $O(n)$ operations one can construct a market-clearing allocation $\boldsymbol{y}$ such that $y_i \in D_i^{c\tau}(cp)$ for every i. If $s \geq (1-\delta)m$ for some $\delta \in [0,1)$, then $c \geq (1-\delta)/(1+(q-1)/m)$.*

Proof. Set $I := \{i \in [n] : x_i > 0\}$. Since $1 \leq s < m$, the set I is nonempty, $q = |I|$, and $c = s/(m+q-1) \in (0,1)$.

We begin by showing that joint scaling preserves relaxed demand when a buyer's quantity is increased by at most the factor $1/c$. Fix a buyer i and a quantity y_i satisfying

$$x_i \leq y_i \leq \min\{m, x_i/c\}.$$

Let $u_i := f_i(x_i) - px_i$. The demand inequality for x_i against the benchmark $k=0$, together with $f_i(0)=0$, gives $u_i \geq 0$. Moreover, for every $k \in \{0,\dots,m\}$,

$$u_i \geq \tau f_i(k) - pk.$$

If the right-hand side is nonnegative, multiplying it by $c \leq 1$ weakens the inequality. If it is negative, the same conclusion follows from $u_i \geq 0$. Hence

$$u_i \geq c\tau f_i(k) - cpk \qquad \text{for every } k.$$

Because $cy_i \leq x_i$, monotonicity of f_i and nonnegativity of p give

$$\begin{aligned} f_i(y_i) - cpy_i &\geq f_i(x_i) - cpy_i \\ &= u_i + p(x_i - cy_i) \\ &\geq u_i \\ &\geq c\tau f_i(k) - cpk. \end{aligned}$$

Thus every such y_i belongs to $D_i^{c\tau}(cp)$.

It remains to show that the positive allocations have enough integer capacity to absorb the $m-s$ unallocated units. For each $i \in I$, define

$$U_i := \min\left\{m, \left\lfloor \frac{x_i}{c} \right\rfloor\right\}.$$

The strict inequality $\lfloor z \rfloor > z-1$ yields

$$\sum_{i \in I} \left\lfloor \frac{x_i}{c} \right\rfloor > \frac{s}{c} - q = m-1.$$

The sum on the left is integral and therefore at least m. If every uncapped quantity is below m, then the cap does not change the sum. If some uncapped quantity is at least m, its capped value alone equals m. Consequently,

$$\sum_{i\in I} U_i \geq m.$$

Since $c<1$ and $x_i \leq m$, we also have $U_i \geq x_i$ for every $i \in I$. Thus the total residual capacity $\sum_{i\in I}(U_i - x_i)$ is at least $m-s$.

Initialize $y_i = x_i$ for $i \in I$ and set $r = m-s$. Scan the buyers in I, increase each y_i by $\min\{r, U_i - y_i\}$, and reduce r by the same amount. The preceding capacity bound ensures that the scan ends with $r=0$. Set $y_i = 0$ for $i \notin I$. Then $\sum_i y_i = m$, and

$$x_i \leq y_i \leq \min\{m, x_i/c\} \qquad \text{for every } i.$$

The demand-preservation argument above gives $y_i \in D_i^{c\tau}(cp)$ for every buyer, including buyers with $x_i = y_i = 0$. The construction uses one scan and therefore takes $O(n)$ operations.

Finally, if $s \geq (1-\delta)m$, then

$$c = \frac{s/m}{1+(q-1)/m} \geq \frac{1-\delta}{1+(q-1)/m},$$

which proves the stated bound. □

B.2. Combining the Two Supply Regimes

THEOREM 2 **(Exact-clearing FPTAS)**. *For every $\varepsilon \in (0,1)$ and arbitrary monotone valuations $\{f_i\}_{i\in[n]}$ with $f_i(0)=0$, there exists an FPTAS (Algorithm 1) that finds a $(1-\varepsilon)\tau^*$-WE in time polynomial in n, $\log m$, and $1/\varepsilon$.*

Proof. *Feasibility for* (TR). In the small-supply branch, Corollary 1 gives demand feasibility and exact clearing. In the large-supply branch, Proposition 4 gives a demand-feasible allocation with

$$s \geq (1-\varepsilon/4)m > 3m/4 \geq 3/2.$$

Thus integrality gives $s \geq 2$. If $s = m$, Algorithm 1 returns this allocation. If $s < m$, Lemma 3 applies and produces a market-clearing allocation that is demand-feasible at $(c\tau, cp)$. In both branches, the returned relaxation lies in $(0,1]$, and the returned triple is feasible for (TR). Consequently, $\widehat{\tau} \leq \tau^*$.

Approximation ratio. The small-supply branch gives $\widehat{\tau} \geq (1-\varepsilon)\tau^*$ directly from Corollary 1. In the large-supply branch, Proposition 4 gives

$$\tau \geq (1-\varepsilon/4)\tau^*, \qquad s \geq (1-\varepsilon/4)m.$$

If $s = m$, set $c = 1$. If $s < m$, let $q = |\{i : x_i > 0\}|$. Since $q \le n$ and $m > 4(n-1)/\varepsilon$,

$$\frac{q-1}{m} \le \frac{n-1}{m} < \frac{\varepsilon}{4}.$$

Lemma 3 and $1/(1+a) \ge 1-a$ for $a \ge 0$ yield

$$c = \frac{s/m}{1+(q-1)/m} \ge \frac{1-\varepsilon/4}{1+\varepsilon/4} \ge (1-\varepsilon/4)^2.$$

The same lower bound holds when $s = m$ and $c = 1$. Hence

$$\begin{aligned} \widehat{\tau} = c\tau &\ge (1-\varepsilon/4)^3 \tau^* \\ &\ge (1-3\varepsilon/4)\tau^* \\ &\ge (1-\varepsilon)\tau^*, \end{aligned}$$

where $(1-a)^3 \ge 1-3a$ for $a \in [0,1]$ gives the second inequality.

Welfare. Exact clearing and demand feasibility make the output a $\widehat{\tau}$-WE. Proposition 2 therefore gives

$$\sum_{i=1}^{n} f_i(\widehat{x}_i) \ge \widehat{\tau} \cdot \mathrm{OPT}(\mathrm{SWM}).$$

Value-query complexity. In the large-supply branch, Proposition 4 is called with both parameters of order ε, and completion makes no value queries. In the small-supply branch, the proposition's displayed bound $O\big(n(1+J\log m)\big)$ is independent of δ. Thus setting $\delta = 1/(2m)$ still gives a value-query count polynomial in n, $\log m$, and $1/\varepsilon$.

Unit-cost operation complexity. In the large-supply branch, Proposition 4 is called with $1/\delta = 4/\varepsilon$, and the completion step adds $O(n)$ operations. In the small-supply branch, Corollary 1 gives $\mathrm{poly}(n, m, 1/\varepsilon)$ operations. Substituting $m \le 4(n-1)/\varepsilon$ makes this count polynomial in n and $1/\varepsilon$, and hence in the parameters stated in the theorem.

□

Appendix C: Remaining Proofs

This section collects the remaining proofs for the structural results in Sections 2 and 4. We establish the properties of the concave envelope, show that the tightest relaxation is well defined and NP-hard to compute, and derive the relative distance from submodularity and the tightest relaxation for Markets 1–3.

C.1. Concave-Envelope Properties

Proof of Lemma 1. Extend the mixture representation in (2) from integer quantities to every $x \in [0, m]$ by defining

$$g(x) \coloneqq \max \left\{ \sum_{q=0}^{m} \lambda_q f(q) : \sum_{q=0}^{m} \lambda_q = 1, \quad \sum_{q=0}^{m} q\lambda_q = x, \quad \lambda_q \geq 0 \right\}. \tag{34}$$

The feasible set is nonempty because any $x \in [0, m]$ is a convex combination of two adjacent integers. It is also a closed subset of the probability simplex and is therefore compact, so the maximum is attained. For every integer k, Equation (34) coincides with (2); hence $g(k) = \overline{f}(k)$. The point mass at k is feasible, so $g(k) \geq f(k)$, and nonnegativity of f gives $g(x) \geq 0$ for every $x \in [0, m]$.

We next show that g is concave. Fix $x, y \in [0, m]$, choose maximizers $\boldsymbol{\lambda}^x$ and $\boldsymbol{\lambda}^y$ in Equation (34), and let $\theta \in [0, 1]$. The vector $\theta\boldsymbol{\lambda}^x + (1-\theta)\boldsymbol{\lambda}^y$ is feasible at mean $\theta x + (1-\theta)y$ and has value $\theta g(x) + (1-\theta)g(y)$. Therefore,

$$g\big(\theta x + (1-\theta)y\big) \geq \theta g(x) + (1-\theta)g(y).$$

Concavity of g on $[0, m]$ implies that the increments $g(k) - g(k-1)$ are weakly decreasing in k. Thus the restriction $\overline{f} = g|_{\{0,\ldots,m\}}$ is submodular.

The same representation shows that $\overline{f}$ is the pointwise smallest submodular valuation that weakly dominates f. Let $h : \{0, \ldots, m\} \to \mathbb{R}$ be any submodular valuation satisfying $h(q) \geq f(q)$ for every q, and extend h linearly between consecutive integers. Submodularity of h makes this extension concave. For any $x \in [0, m]$ and any feasible $\boldsymbol{\lambda}$ in Equation (34), Jensen's inequality gives

$$h(x) \geq \sum_{q=0}^{m} \lambda_q h(q) \geq \sum_{q=0}^{m} \lambda_q f(q).$$

Maximizing over $\boldsymbol{\lambda}$ yields $h(x) \geq g(x)$ and, in particular, $h(k) \geq \overline{f}(k)$ at every integer k.

It remains to establish normalization and monotonicity. At $x = 0$, the mean constraint admits only the point mass at zero, so $g(0) = f(0) = 0$. Now fix $0 \leq x < y \leq m$ and set $\theta = (m-y)/(m-x)$. Mix a maximizer for $g(x)$ with the point mass at m, using weights θ and $1-\theta$. The resulting distribution has mean y. Moreover, monotonicity of f implies $g(x) \leq f(m)$ because $g(x)$ is a convex combination of the values $f(q)$. Hence

$$g(y) \geq \theta g(x) + (1-\theta) f(m) \geq g(x).$$

Therefore, $\overline{f}$ is normalized, nonnegative, monotone, and submodular.

Suppose that f is submodular. Its piecewise-linear extension is concave, so Jensen's inequality gives, for every feasible $\boldsymbol{\lambda}$ with integer mean k,

$$f(k) \geq \sum_{q=0}^{m} \lambda_q f(q).$$

Taking the maximum over $\boldsymbol{\lambda}$ shows $f(k) \geq \overline{f}(k)$. The reverse inequality follows from the feasible point mass at k, and hence $\overline{f} = f$. Conversely, if $\overline{f} = f$, then submodularity of $\overline{f}$ implies that f is submodular. □

C.2. Existence, Attainment, and NP-Hardness

Proof of Lemma 2. Let $f_{\max} \coloneqq \max_{i\in[n]} f_i(m)$ and fix a buyer $r \in \arg\max_i f_i(m)$. Set

$$\tau_0 \coloneqq \frac{1}{m}, \qquad p_0 \coloneqq \frac{f_{\max}}{m},$$

and allocate all m units to buyer r. Every buyer's realized utility is zero. For every buyer i and every positive quantity k, monotonicity gives

$$\tau_0 f_i(k) - k p_0 \leq \frac{f_{\max}}{m} - k\frac{f_{\max}}{m} \leq 0,$$

and the benchmark at quantity zero is also zero. Hence the allocation and price p_0 form a τ_0-WE, proving $\tau^* \geq 1/m$.

It remains to show that the maximum is attained. We may restrict τ to $[1/m, 1]$. At any feasible triple, market clearing gives a buyer i with $x_i \geq 1$. Comparing x_i with quantity zero in the relaxed demand inequality yields

$$f_i(x_i) - x_i p \geq 0, \qquad \text{so} \qquad p \leq \frac{f_i(x_i)}{x_i} \leq f_{\max}.$$

For each fixed market-clearing allocation $\boldsymbol{x}$, the relaxed demand constraints are finitely many weak linear inequalities in (τ, p). They therefore define a closed subset of the compact rectangle $[1/m, 1] \times [0, f_{\max}]$. There are finitely many market-clearing allocations, so the feasible set is a nonempty finite union of compact sets. The continuous objective τ attains its maximum on this set. □

Proof of Proposition 3. It suffices to show that deciding whether $\tau^* = 1$ is NP-hard, because exact computation of τ^* solves this decision problem. The reduction constructs a market whose exact-demand allocations at its only possible clearing price correspond to subsets of a Subset Sum instance.

Subset Sum is NP-complete (Kellerer et al. 2004, Karp 2009). Given positive integers $a_1, \ldots, a_n$ and a target T, the problem asks whether there exists a subset $S \subseteq [n]$ such that $\sum_{i\in S} a_i = T$. We restrict attention to instances satisfying $a_i \leq T$ for every i, $\sum_i a_i > T$, and $T \geq 2$. These restrictions follow by standard preprocessing and, for the last one, scaling all integers by two.

Set the supply to $m = T$ and give buyer i the single-minded valuation

$$f_i(k) = \begin{cases} 0, & k < a_i, \\ a_i, & k \geq a_i. \end{cases}$$

For $p < 1$, every demanded quantity is at least a_i, so any demand-feasible allocation uses more than m units. For $p > 1$, every buyer demands zero. At $p = 1$, buyer i demands exactly 0 and a_i. A Walrasian equilibrium therefore exists if and only if some subset of the a_i sums to T. The construction is polynomial, and Lemma 2 implies that a Walrasian equilibrium exists if and only if $\tau^* = 1$. Thus deciding whether $\tau^* = 1$, and consequently computing τ^*, is NP-hard. □

C.3. Analysis of Markets 1–3

The calculations below establish the relative distance from submodularity and the tightest relaxation reported in Section 4. For Market 1, we also derive the exact value of τ^* underlying the strict comparison in Table 1.

Market 1. Let $v_0 = 0$ and $v_q := f_2(q) = q + (q-1)\xi$ for $q \geq 1$. We first show that, for $m \geq 3$,

$$\tau^* = \tau_1 := \frac{m}{m-1}\frac{v_{m-1}}{v_m} = \frac{m}{m-1}\frac{m-1+(m-2)\xi}{m+(m-1)\xi}. \tag{35}$$

For fixed (τ, p), buyer 1's scaled benchmark is maximized at quantity 0 or 1. Buyer 2's marginal values are nondecreasing, so her scaled benchmark is maximized at quantity 0 or m. Consequently, for a clearing allocation $(k, m-k)$, their relaxed demand constraints reduce to

$$\begin{aligned}
k = 0: &\quad p \geq 2\tau, \\
k \geq 1: &\quad p \leq \frac{2}{k}, \qquad \tau \leq 1 - \frac{(k-1)p}{2}, \\
m-k = 0: &\quad p \geq \frac{\tau v_m}{m}, \\
m-k \geq 1: &\quad p \leq \frac{v_{m-k}}{m-k}, \qquad \tau \leq \frac{v_{m-k} + kp}{v_m}.
\end{aligned} \tag{36}$$

At allocation $(1, m-1)$, set $p_1 = v_{m-1}/(m-1)$. Equation (36) holds with $\tau = \tau_1$. In particular, $p_1 < 2$ and

$$1 - \tau_1 = \frac{\xi}{(m-1)v_m} > 0.$$

Thus $(p_1, (1, m-1))$ is a τ_1-WE and $\tau^* \geq \tau_1$.

To prove the reverse inequality, consider the remaining clearing allocations. If $k = 0$, Equation (36) gives $\tau \leq v_m/(2m) < \tau_1$. The strict inequality is equivalent to

$$2m^2 v_{m-1} - (m-1)v_m^2 = m^2(m-1) - 2m\xi - (m-1)^3\xi^2 > 0.$$

The expression on the right is concave in ξ and is positive at both $\xi = 0$ and $\xi = 1$, where its values are $m^2(m-1)$ and $2m^2 - 5m + 1$, respectively.

If $2 \leq k \leq m-1$, combining the two bounds on τ in Equation (36) yields

$$\tau \leq \frac{v_{m-k} + 2k/(k-1)}{v_m + 2k/(k-1)}.$$

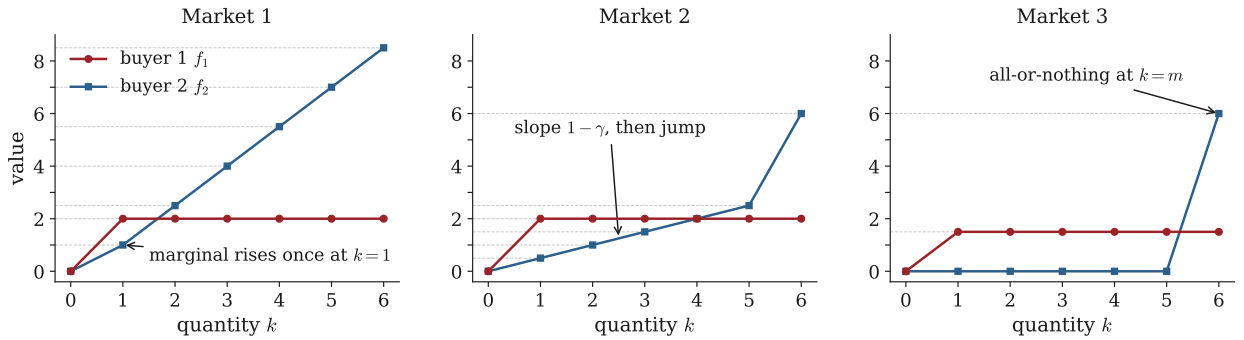


Figure 3 Valuation functions of the three markets, drawn with $m=6$. Markets 1 and 3 use $\xi=\frac{1}{2}$, and Market 2 uses $\gamma=\frac{1}{2}$. Buyer 1 is unit-demand throughout, while Buyer 2's valuation differs across the three markets.

The gap between this upper bound and one is at least

$$\frac{2(1+\xi)}{v_m+4} > \frac{1}{m+2} > \frac{1}{m(m-1)} > 1-\tau_1,$$

where we use $v_m < 2m$ and $m \geq 3$. Hence this upper bound is strictly smaller than τ_1. If $k=m$, buyer 2's deviation to m gives $p \geq \tau v_m/m \geq \tau$, while buyer 1's deviation to one unit gives $\tau \leq 1-(m-1)p/2$. Therefore, $\tau \leq 2/(m+1) \leq 1/2 < \tau_1$, where the last inequality follows from $v_{m-1} \geq m-1$ and $v_m < 2m$. These cases exhaust all clearing allocations and prove Equation (35).

Buyer 1 is submodular and contributes no gap. Buyer 2 is supermodular, and her concave envelope is the chord joining $(0,0)$ and (m,v_m): $\overline{f}_2(k)=kv_m/m$. For $1 \leq k \leq m$,

$$1-\frac{f_2(k)}{\overline{f}_2(k)} = \frac{\xi(m-k)}{k\big(m+(m-1)\xi\big)}.$$

This expression is decreasing in k, so

$$\gamma = \frac{(m-1)\xi}{m+(m-1)\xi}, \qquad 1-\gamma = \frac{m}{m+(m-1)\xi}.$$

Together with Equation (35), this gives

$$\tau^* = (1-\gamma)\left(1+\frac{(m-2)\xi}{m-1}\right) > 1-\gamma.$$

Market 2. Buyer 1 is submodular, so $\overline{f}_1 = f_1$. The valuation $k \mapsto k$ is a submodular majorant of buyer 2's valuation, which gives $\overline{f}_2(k) \leq k$. Conversely, the piecewise-linear extension of $\overline{f}_2$ is concave and satisfies $\overline{f}_2(0)=0$ and $\overline{f}_2(m) \geq m$. Hence $\overline{f}_2(k) \geq k$, so $\overline{f}_2(k)=k$. It follows that

$$\gamma = 1-\beta.$$

The triple $(\tau, p, \boldsymbol{x}) = (\beta, 1, (0, m))$ is feasible for (TR). Buyer 1's scaled benchmark is $1 - k \le 0$ for $k \ge 1$. Buyer 2's scaled benchmark is $(\beta^2 - 1)k \le 0$ for $k < m$ and $(\beta - 1)m \le 0$ at $k = m$. Both buyers have realized utility zero, so $\tau^* \ge \beta$.

Now let $(p, \boldsymbol{x})$ be any τ-WE and write $x_1 = m - x_2$. The three possible ranges of x_2 give the matching upper bound:

- If $x_2 = m$, buyer 2's deviation to zero gives $p \le 1$, while buyer 1's deviation to one unit gives $p \ge \tau/\beta$. Thus $\tau \le \beta$.
- If $1 \le x_2 \le m - 1$, buyer 2's deviation to zero gives $p \le \beta$. Her deviation to m then gives
$$\tau \le \frac{\beta x_2 + p(m - x_2)}{m} \le \beta.$$
- If $x_2 = 0$, buyer 2's deviation to m gives $p \ge \tau$. Buyer 1's deviation to one unit gives $(1-\tau)/\beta \ge p(m-1)$ and therefore
$$\tau \le \frac{1}{1 + \beta(m-1)} \le \beta,$$
where the last inequality follows from $m \ge 1 + (1-\beta)/\beta^2$.

Therefore, $\tau^* = \beta = 1 - \gamma$.

Market 3. Buyer 1 is submodular. As in Market 2, the valuation $k \mapsto k$ is a submodular majorant of buyer 2's valuation, while concavity of the piecewise-linear extension of any submodular majorant places it weakly above the chord joining $(0, 0)$ and (m, m). Thus $\overline{f}_2(k) = k$ and $\gamma = 1$.

Set $A := 1 + \xi$. We next show that, under $m \ge A^2$, $\tau^* = 1/A$. The possible clearing allocations are as follows.

- At $(0, m)$, buyer 1's deviation to one unit gives $p \ge \tau A$, and buyer 2's deviation to zero gives $p \le 1$. Hence $\tau \le 1/A$. Equality is feasible at $(\tau, p) = (1/A, 1)$, so $\tau^* \ge 1/A$.
- At $(m, 0)$, buyer 2's deviation to m gives $p \ge \tau$. Buyer 1's deviation to one unit gives $A(1-\tau) \ge (m-1)p$. Hence
$$\tau \le \frac{A}{A + m - 1} \le \frac{1}{A},$$
where the last inequality follows from $m \ge A^2$.
- At $(k, m-k)$ with $0 < k < m$, buyer 2's deviation to zero forces $p = 0$. Her deviation to m then rules out every $\tau > 0$.

It follows that $\tau^* = 1/(1+\xi) > 1 - \gamma = 0$.