Walrasian Mechanisms for Non-convex Economies and the Bound-Form First Welfare Theorem

Paul Milgrom* and Mitchell Watt†

June 29, 2022

Abstract

We introduce two extensions of the Walrasian mechanism for quasilinear economies to allow agents to report non-concave values and non-convex costs. The extended mechanisms, which always deliver feasible, near-efficient allocations with no budget deficit, are computationally undemanding and nearly incentive-compatible. We also introduce an extension of the First Welfare Theorem allowing us to upper bound the welfare losses from these mechanisms.

Keywords: Approximate efficiency, Approximate incentive-compatibility, Market design, Non-convexity, Prices, Rationing

JEL Codes: C620, D400, D440, D450, D470, D500, D510, D610.

*Stanford University and Auctionomics. Email: milgrom@stanford.edu

†Stanford University. Email: mwatt@stanford.edu.

Thank you to Mohammad Akbarpour, Ravi Jagadeesan, Fuhito Kojima, Shoshana Vasserman and seminar participants at Stanford University, the University of Zurich, the NBER Market Design Working Group and the 32nd Stony Brook International Conference on Game Theory for helpful comments and suggestions relating to this project. An extended abstract of this paper appeared in *Proceedings of the 23rd ACM Conference on Economics and Computation (EC'22), July 11-15, 2022, Boulder, CO, USA* as "Linear Pricing Mechanisms for Markets without Convexity". Milgrom gratefully acknowledges support from the National Science Foundation (grant SES-1947514). Watt gratefully acknowledges support of the Koret Fellowship and the Ric Weiland Graduate Fellowship.

1 Introduction

We introduce extensions of the Walrasian mechanism for quasilinear economies for application in economic sectors with multiple products and significant fixed costs or other non-convexities. Two examples of such sectors are electrical power, in which products are typically distinguished by location and time-of-day with production subject to fixed costs for generation equipment and for ramping production up and down, and regulated fisheries, in which products are typically distinguished by location and species with production subject to fixed costs for boats and for travel to and from fishing locations.

A mechanism extends the Walrasian mechanism if it always returns a competitive equilibrium when one exists and otherwise outputs an allocation and payments that achieve or approximate the desirable properties of the Walrasian mechanism (like efficiency, individual rationality, budget balance, and large market incentive-compatibility). We focus on extensions that use anonymous linear prices to determine payments. Such pricing mechanisms are simple and often perceived to be fairer than personalized payment mechanisms. Our work shows how prices can help to coordinate individual decisions even in complex, non-convex markets where market-clearing prices may not exist.

There are four main challenges to achieving practical extensions of the Walrasian mechanism. The first is *computational complexity*: non-convex welfare optimization can be hard. We seek mechanisms for which the computations are not much harder than when value reports are restricted to be concave and cost reports convex. The second challenge is to ensure that there are *no budget deficits*. This can be problematic for non-convex problems, because it is possible that for all linear prices, the resulting demands and supplies are always either infeasible or have production that strictly exceeds consumption. The third challenge is to ensure *individual rationality*. Even when efficient allocations can be identified, there may be no linear prices such that all agents prefer the recommended allocation to the zero trade. The fourth and final challenge is to encourage *truthful reporting (despite rationing)*. In the traditional Walrasian mechanism for convex economies, a participant can benefit from a false report only to the extent that the report alters the prices. When agents are assigned non-preferred bundles given the price, the incentive challenge is harder, because even an agent who cannot affect prices may find it profitable to misreport in order to avoid being rationed. The extended Walrasian mechanisms described below overcome these challenges.

Our first extensions of the Walrasian mechanism are the *markup mechanisms*. In a markup mechanism, buyers and sellers report values and costs and the mechanism returns a triple (α, p, ω) consisting of a non-negative scalar *markup* α and price and allocation vectors. The allocation ω assigns each seller one of its preferred bundles at prices p and each buyer one of its preferred bundles at prices $(1 + \alpha)p$, and agents make payments accordingly. Markup mechanisms are further restricted to satisfy *physical feasibility*, so that total production of goods weakly exceeds total consumption, and *budget feasibility*, so that total payments by buyers weakly exceeds the total payments to sellers. If a competitive equilibrium exists, then $\alpha = 0$ and the pair (p, ω) is a competitive equilibrium. More generally, the *minimum markup mechanism* selects the markup α to be the smallest non-negative number such that for some price-allocation pair, the physical and budget feasibility conditions are satisfied. Because minimum markups can be hard to compute, we focus on a *simple markup mechanism*, showing that our simple markups are easy to compute, inversely proportional to the number of market participants N, and at worst proportional to both the number of products and the "size" of the largest non-convexity in the buyers' value functions or the sellers' cost functions.

Elementary economic analysis confirms that every markup mechanism is individually rational and that, with truthful reporting, its allocation allows no further gains from trade either among buyers or among sellers. Using our Bound Form First Welfare Theorem, we show that there is an $O(1/N)$ bound on the percentage welfare loss for both the simple and minimum markup mechanisms. The incentive properties of these two markup mechanisms with non-convex reports are similar to those of original Walrasian mechanism restricted to convex reports. For all of these mechanisms, if excess demand is strongly monotone (Watt, 2022), then the largest amount that any single participant can gain by false reporting is at most of order $O(1/N)$.[1]

Next, we introduce a family of *rationing mechanisms*, designed for sectors in which one side of the market—say, the buyers' side—has convex preferences. These mechanisms return a price and feasible allocation (p, ω) such that total payments to sellers equal total receipts from buyers, but some agents may be allocated non-preferred bundles given the prices p. To maximize welfare, one might seek to identify a *minimum-loss rationing mechanism* which would minimize the total *rationing losses* suffered by all agents, where an agent's rationing loss at (p, ω) is the difference between its payoff from a preferred bundle at price p and its payoff given its allocation in ω given

[1] In addition, when values (costs) are drawn i.i.d. from some underlying distribution for which the *expected* demand (supply) is strongly monotone, then the maximum gain from false reporting about that value (cost) is similarly bounded with high probability. This condition on expected demands and supplies can be helpful to prove conclusions about models with with random values or costs (see Watt (2022) and Section 4 below).

price p. However, we show that identifying such a mechanism requires solving the efficient allocation problem for the economy, which may be computationally intractable. Instead, we introduce a *buyers-only rationing mechanism* which is computed by solving just two convex optimization problems and assigns all sellers their most preferred bundles at price vector p. If buyers have strongly monotone demands, then the buyers-only rationing mechanisms suffers a total welfare loss of only $O(1/N)$ and a percentage welfare loss of $O(1/N^2)$. For this mechanism, too, we prove that the gains from misreporting are also on the order of $O(1/N)$.

Our analyses of welfare losses and incentives in the extended mechanisms apply our new *Bound-Form First Welfare Theorem* for quasilinear economies, which applies even when no competitive equilibrium exists. The theorem asserts that for any price-allocation pair (p, ω), the welfare loss of ω is no greater than the budget deficit from allocation ω at prices p plus the sum over all agents of the agent's loss from receiving its ω-bundle rather than its most preferred bundles at prices p (the rationing losses). For the special case where (p, ω) is a competitive equilibrium, this restates the conclusion of the standard First Welfare Theorem: the welfare loss for ω is zero.

The rest of this paper is organized as follows. In the remainder of Section 1, we consider an example that highlights the challenges of market design with non-convexities and illustrates our proposed mechanisms, before reviewing related literature. In Section 2, we introduce the model and discuss known results for convex economies. In Section 3, we introduce and prove the Bound-Form First Welfare Theorem. In Section 4, we introduce markup mechanisms. Theorem 2 shows that the *relative inefficiency* of the minimum markup mechanism and an easily-computable approximation thereof, is of order $1/N$, while Theorem 3 establishes that these mechanisms have good incentive properties similar to those of the Walrasian mechanism. In Section 5, we consider the special case when one side of the market – consumers but not producers – has convex preferences and introduce our family of rationing mechanisms. In Section 6, we conclude.

1.1 An example

To gain some intuition for the challenges and solutions, we temporarily set aside the computational challenge to focus on the other three using an example with just one good and one producer. The producer's technology is not convex: it can make only nonnegative integer numbers of units at a marginal cost $c = \$0.50$.

There are N buyers indexed $n = 1, ..., N$ and their preferences are not convex. If the price is low enough, buyer n would like to purchase *exactly* q_n units of the good: buyers derive no value from receiving fewer than q_n units and no additional value for additional units beyond q_n. We refer to q_n as buyer n's *order size*. Buyers have quasilinear preferences, where buyer n's per-unit value for the good is v_n, so its utility $U_n(x_n, t)$ from paying \$$t$ for x_n units is

$$U_n(x_n,t) = \begin{cases} -t & \text{if } x_n < q_n \\ q_n v_n - t & \text{if } x_n \geq q_n. \end{cases}$$

The supply and demand curves for a specific instance of this market, using the buyer values and order sizes in Table 1, are depicted in Figure 1. Note that the supply and demand correspondences fail to be lower hemicontinuous due to the non-convexities on both sides of the market.

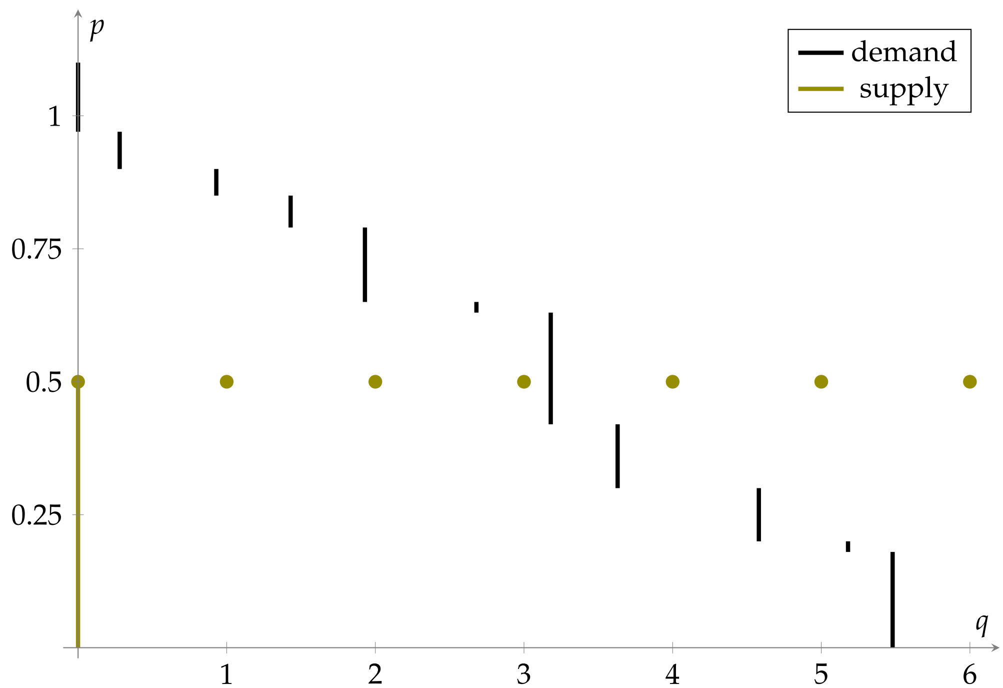

Figure 1: Example market

Buyer n	1	2	3	4	5	6	7	8	9	10
Unit value v_n	0.97	0.9	0.85	0.79	0.65	0.63	0.42	0.3	0.2	0.18
Order size q_n	0.28	0.65	0.5	0.5	0.75	0.5	0.45	0.95	0.6	0.3

Table 1: Buyer values

How should an extended Walrasian mechanism allocate this good? If a market-clearing price exists, then any extension must assign each agent its preferred bundle at the clearing price to implement the competitive equilibrium. In our example, if the sum of q_n over all agents with

$v_n \geq 0.5$ is an integer, then \$0.50 is a competitive equilibrium price. However, if this sum is non-integer, as in Figure 1, then there is no competitive equilibrium.

If there is no competitive equilibrium, a single price mechanism that nearly achieves market clearing may seem attractive. A classic result of Starr (1969) implies that, in a non-convex market with one good, there exists a price vector and allocation such that at most one agent is assigned an infeasible bundle and that bundle is in the convex hull of the agent's preferred set. The resulting price-allocation pair is called a *pseudoequilibrium.* In the example of Figure 1, the pseudoequilibrium allocates the first six buyers their full orders and the seller must produce the resulting 3.18 units at a price \$0.50, which is infeasible given the seller's production technology. In this paper, we limit attention to mechanisms that can select only feasible bundles for every agent.[2]

An alternative to this pseudoequilibrium is to round the allocation so that all agents are allocated their preferred feasible bundles. This approach, pioneered by Starr (1969), results in an *approximate equilibrium*. It is not an exact equilibrium because there is an imbalance in supply and demand. In the market of Figure 1, at a price of \$0.50, the seller has multiple preferred bundles. It could maximize its profit by producing 4 units of the good and fulfilling the orders of the first six buyers, resulting in an oversupply of 0.82 units of the good. Alternatively, the seller could produce 3 units, resulting in an undersupply of 0.18 units. If the approximate equilibrium allocation exhibits undersupply, it is physically infeasible, because buyers are being promised more than the seller has supplied. If it exhibits oversupply, then it is budget-infeasible, because it must pay the seller more than it receives from buyers. In this paper, we consider only mechanisms for which the outcomes are both physically feasible and budget feasible.

Given an approximate equilibrium, feasibility could in principle be restored by rationing, that is, assigning some agents larger or smaller bundles than their most preferred ones. In the market of Figure 1, the regulator could require the seller to produce 3 units and fulfill only the orders associated with five of the six buyers with values above \$0.50. But which agents should the mechanism ration? In the example of Figure 1, the most efficient choice is to ration buyer 1, who also has the *highest* per-unit value of the good. If the adopted mechanism always selects the most efficient rationing, buyer 1 could benefit by reporting a value of 1.1 or an order size of 0.31 so that its reported contribution to surplus is larger than that of buyer 6. Such a report would not affect the price, but would nevertheless be profitable, because it allows buyer 1 to avoid rationing.

[2] If we add free disposal for the producer, it becomes feasible for it to supply 3.18 units (by producing 4 and discarding 0.82) but the producer would not be willing to do so at a price of \$0.50, and at higher prices, the seller would want to produce many more units.

In the standard Walrasian model, each buyer is assigned a most preferred bundle at the given prices, so a buyer can gain from a false report only if that report alters the clearing prices. Roberts and Postlewaite (1976) show, if the equilibrium price is always unique,[3] then the ability of any single agent to influence that price tends to zero as the market grows large, and hence each agent finds truthful reporting to be nearly optimal. In contrast, as our example illustrates, a potentially rationed agent may have an incentive to misreport that remains bounded away from zero even as the market grows large.

This paper offers extensions of the Walrasian mechanism with the following desirable properties:

- computations are not much harder than for the convex case, requiring only a small number of convex optimizations,
- solutions are physically- and budget-feasible, individually rational (IR) or near-IR, and nearly-truthful (with vanishing gains to false reports in large markets), and
- the selected outcome is a competitive equilibrium when one exists.

Our first extensions, the markup mechanisms developed in Section 4, introduce one additional scalar parameter to the Walrasian model: a nonnegative percentage markup to be added to sellers' prices to determine the buyers' prices. In the market of Figure 1, a markup mechanism could use the approximate equilibrium allocation in which the seller produces 4 units of the good, receiving \$0.50 per unit, and the orders of the first six buyers are filled, but at the price \$0.63. The revenue earned from sale to the buyers fully covers the seller's cost of production, and all agents are assigned their most preferred bundles at their respective prices. In Section 4, we show that it is always possible to find an allocation and a markup with the property that the welfare loss compared to the maximum total welfare loss is bounded by a constant independent of the number of agents N. In general, the markup may be chosen to be $O(1/N)$, vanishing as the market grows large. The only computations required to approximate the minimum markup mechanism are convex optimizations and a one-dimensional binary search. Moreover, because no agent is rationed given its price vector, the incentives in this mechanism closely resemble those of the Walrasian mechanism.

[3] Or more generally, if the Walrasian equilibrium price correspondence mapping economies to Walrasian prices is continuous at the limit economy. See Section 4.3.

If we modify this example slightly by convexifying the preferences of the buyer side, then our rationing mechanisms can lead to a smaller welfare loss than a markup mechanism. For suppose our buyers were willing to accept *no more* than q_n units at the value v_n (rather than *exactly* q_n units). Then, at the price of \$0.50, the seller could supply 3 units, with buyers 1 through 5 assigned the full q_n units at that price and buyer 6 assigned the remaining 0.32 units. Because buyer 6 would prefer to purchase 0.65 units at that price, buyer 6 is *rationed*. In Section 5, we show that with strongly monotone buyer demands, our buyers-only rationing mechanism changes any single buyer's bundle by at most $O(1/N)$. Then, using the envelope theorem and our Bound-Form First Welfare Theorem, we find that this rationing mechanism has $O(1/N)$ total welfare loss, which corresponds to a fractional welfare loss of $O(1/N^2)$. The maximum gain to any agent from misreporting then diminishes at the same $O(1/N)$ rate as the original Walrasian mechanism with the additional assumption of convex costs.

1.2 Related literature

The problem of non-convexity for the existence of competitive equilibrium was discussed in a series of papers by Farrell (1959), Rothenberg (1960), Koopmans (1961) and Bator (1961). Much of the subsequent classical literature on non-convexity in general equilibrium theory focused on concepts of *approximate* equilibria which replace aggregate feasibility requirements with approximate feasibility, measured in terms of distance in the commodity space between the aggregate supply and demand, while maintaining the requirement that individual agents act optimally given the prices. Starr (1969) showed the existence of such an approximate equilibrium in non-convex production economies, in which the maximum imbalance is proportional to the number of goods and a measure of non-convexity. Heller (1972) proved a similar result with an alternative measure of non-convexity. More recently, Nguyen and Vohra (2022) proved a bound for markets with indivisible goods that depends only on a measure of preference complementarity of agents. We build on some of these results (summarizing the key results we employ in Appendix A), but depart from this literature by requiring that any feasible mechanism must always specify a feasible outcome. Influenced by computer scientists' approaches to approximations in mechanism design, we will be interested in approximate efficiency and truthfulness, rather than approximate feasibility.[4]

A substantial literature has focused on identifying various conditions on preferences in markets with indivisibilities, under which competitive equilibria exist despite non-convexities. Con-

[4] Scarf (1967) also features an approximate efficiency objective.

tributors include Bikhchandani and Mamer (1997), Gul and Stacchetti (1999), Danilov, Koshevoy, and Murota (2001), Sun and Yang (2006), Milgrom and Strulovici (2009), Hatfield, Kominers, Nichifor, Ostrovsky, and Westkamp (2013), Baldwin and Klemperer (2019) and Nguyen and Vohra (2022). Milgrom (2009) emphasizes the reporting language that agents use when goods are substitutes. None of these papers treat markets with fixed costs such as those described above, in which competitive equilibria do not generally exist. Our analysis seeks to develop practical mechanisms for those settings.

An alternative approach to establish equilibrium existence in markets with non-convexities is to study the large market limit with a continuum of agents. Aumann (1966) showed that the convexity assumption is not necessary for equilibrium existence in an economy with a continuum of traders and divisible goods, while Azevedo, Weyl, and White (2013) demonstrated a similar result for an economy with indivisibilities. In this paper, we study mechanisms that can operate in finite economies and investigate rates of convergence as suggestive of the potential performance of the mechanisms. Yet another approach is to allow nonlinear or personalized pricing rules, as explored by Wilson (1993), Chavas and Briec (2012), Azizan, Su, Dvijotham, and Wierman (2020) and others, but mechanisms that use anonymous linear prices may sometimes be preferred for reasons outside our models, such as familiarity and perceived fairness.

A two-price solution to equilibrium nonexistence has also been proposed in a contemporaneous contribution of Feldman, Shabtai, and Wolfenfeld (2021). The key difference between our approaches is the structure and role of the two prices: Feldman et al. (2021) consider (one-sided) exchange economies in which buyers who are allocated a good face one price for the good and buyers who are not allocated a good face a different price for the same good. The role of the two prices in their mechanism is to prevent buyers from wanting to change their bundle of goods from the one allocated by the market designer. We restrict attention to mechanisms that use the same price vector for all buyers (and similarly all sellers) regardless of whether they are allocated a good, which makes achieving incentive-compatibility more difficult.

Our study is motivated by several important applies of linear pricing mechanisms with non-convex production. In particular, we have taken inspiration from the novel market design for fisheries rights in New South Wales, Australia, introduced by Bichler, Fux, and Goeree (2018, 2019), in which the need to implement sustainable catches led to exit of fishing boats, with an associated loss of fixed costs. Other sectors with non-convexities that have used linear prices electricity generation, with their large start-up and ramping costs, radio spectrum, where geographical complementarities can cause exposure problems.[5]

Notation In this paper, we will view consumption and production bundles as vectors in Euclidean space $\mathbb{R}^L$, and let $x \cdot y$ denote the standard inner product of x and y and $\|x\| = \sqrt{x \cdot x}$ denote the Euclidean norm of x. We use $\geq$ to denote the partial order on $\mathbb{R}^L$ so that $x \geq y$ if and only if $x_\ell \geq y_\ell$ for $\ell = 1, ..., L$. The set $\mathbb{R}^L_+$ is $\{x \in \mathbb{R}^L : x \geq 0\}$. The radius of a set $S \subseteq \mathbb{R}^L$ is defined by $\text{rad}(S) = \inf_{x \in \mathbb{R}^L} \sup_{y \in S} \|x - y\|$ (intuitively, the greatest distance from the "center" of the set). The notation $|\cdot|$ represents either the absolute value (if its argument is a number) or the cardinality (if its argument is a set). The Hausdorff distance between sets $S, S' \subseteq \mathbb{R}^L$ is denoted d_H and is defined by letting $\text{dist}(x, S) = \inf_{y \in S} \|x - y\|$ and $d_H(S, S') = \max\{\sup_{x \in S'} \text{dist}(x, S), \sup_{x \in S} \text{dist}(x, S')\}$.

The convex hull of a set $S \subseteq \mathbb{R}^L$ is denoted by $\text{co}(S) = \{x \in \mathbb{R}^L : x = \lambda y + (1 - \lambda)z \text{ for } y, z \in S \text{ and } \lambda \in [0, 1]\}$. The upper concave envelope of $f : S \to \mathbb{R}$ is the function $\text{cav}(f) : \text{co}(S) \to \mathbb{R}$ which is the (pointwise) smallest concave function g on $\text{co}(S)$ that satisfies $g(x) \geq f(x)$ for all $x \in S$. The lower convex envelope of $f : S \to \mathbb{R}$ is the function $\text{vex}(f) : \text{co}(S) \to \mathbb{R}$ which is the largest convex function g on $\text{co}(S)$ that satisfies $g(x) \leq f(x)$ for all $x \in S$.

We make regular use of the asymptotic notation due to Knuth (1976). For functions f and g both mapping $\mathbb{R}$ to $\mathbb{R}$, we write $f(x) = O(g(x))$ if $\limsup_{x \to \infty} \frac{|f(x)|}{g(x)} < \infty$; $f(x) = \Omega(g(x))$ if $\liminf_{x \to \infty} \frac{|f(x)|}{g(x)} > 0$; and $f(x) = \Theta(g(x))$ if $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$. We write $f(x) = O_P(g(x))$ if there is some $M > 0$ such that $\lim_{x \to \infty} \Pr\left(|f(x)/g(x)| > M\right) = 0$.

2 Model and preliminaries

2.1 Model

We employ a Walrasian model with a set of *buyers* N and a set of *firms* or *sellers* F, both finite. Together $A = N \cup F$ is the set of *agents*. There are L varieties of consumable *goods* and a numeraire, money.

Each buyer $n \in N$ chooses a consumption bundle in X, a compact subset of $\mathbb{R}^L_+$ containing 0, called the *consumption possibility set*. Buyer n has quasilinear preferences[6] over bundles in X with

[5]See Liberopoulos and Andrianesis (2016) for a summary of pricing mechanisms used in electricity markets with non-convexities, most of which include "uplift" (or side-payments) in addition to linear pricing, and Ausubel and Milgrom (2002) for a discussion of complementarities in spectrum auctions.

[6]The quasilinearity assumption allows our analysis to abstract from income effects, as is usual in mechanism design analyses. For more discussion of the role of income effects see Morimoto and Serizawa (2015).

a continuous *valuation function* $u_n : X \to \mathbb{R}$, so that the buyer's *utility* associated with receiving allocation x_n and paying t is $U_n(x_n, t) = u_n(x_n) - t$. We suppose that the valuation functions are bounded, nondecreasing with respect to the partial order $\geq$ on $\mathbb{R}^L_+$ and normalized so that $u_n(0) = 0$. We let $\mathscr{U}$ be the space of possible valuation functions for the buyers, which we assume is admissible in the sense of Aumann (1963).[7]

Each seller $f \in F$ chooses a production bundle in the *production possibility set* Y, a compact subset of $\mathbb{R}^L_+$ containing 0. Seller f has a *cost function*[8] $c_f : Y \to \mathbb{R}_+$ which allows us to write f's *profit* from producing $y_f \in Y$ and receiving payment t as $\pi_f(y_f, t) = t - c_f(y_f)$. The cost functions are nondecreasing with respect to the partial order $\geq$ on $\mathbb{R}^L_+$ and normalized so that $c_f(0) = 0$. Let $\mathscr{C}$ be the space of sellers' cost functions, which we also assume to be admissible.

An *economy* $\mathscr{E}$ consists of buyers with their valuation functions and sellers with their cost functions, so that we may write $\mathscr{E} = \langle N, (u_n)_{n\in N}, F, (c_f)_{f\in F}\rangle$. When it is clear, we use the shorthand $\mathscr{E} = \langle N, F\rangle$. At times, it is also convenient to associate $\mathscr{E}$ with the normalized counting measures μ on $\mathscr{U}$ and ν on $\mathscr{C}$ defined by

$$\mu(u_n) = \frac{\text{\# of buyers in } \mathscr{E} \text{ with valuation function } u_n}{|N|},$$
$$\chi(c_f) = \frac{\text{\# of sellers in } \mathscr{E} \text{ with cost function } c_f}{|F|},$$

and to let $\phi = \frac{|F|}{|N|}$, so that $\langle N, \mu, \phi, \chi\rangle$ is an alternative specification of economy $\mathscr{E}$.

Throughout, we will suppose that agent types—that is, the valuation functions u_n of buyers and cost functions c_f of sellers—are private information, but $\mathscr{U}, \mathscr{C}$, $|N|$ and $|F|$ are common knowledge. In some results, we specialize to an *independent private valuations (IPV)* model, in which buyer types are drawn i.i.d. from a common knowledge distribution μ on $\mathscr{U}$ and seller types are drawn i.i.d. from common knowledge χ on $\mathscr{C}$.

Allocations and efficiency An *allocation* $\omega = ((x_n)_{n\in N}, (y_f)_{f\in F})$ is an assignment of consumption bundles $x_n \in X$ to each buyer $n \in N$ and production bundles $y_f \in Y$ to each seller $f \in F$. An allocation is *feasible* if $\sum_{n\in N} x_n \leq \sum_{f\in F} y_f$. We denote by Ω the set of all feasible allocations.

[7] That is, it is possible to define a measure on $\mathscr{U}$, equipped with an appropriate σ-algebra. For example the set of bounded, continuous functions on a compact subset of $\mathbb{R}^L$ is admissible, as is the set of bounded functions with discontinuities of the first kind, or more generally, any subset of a Baire class (Aumann, 1963).

[8] Note that sellers in this economy could equivalently be thought of as buyers with valuations $-c_f(y_f)$ and payments $-t$. However, we will be interested in mechanisms that may charge buyers and sellers different prices, and so it is convenient to distinguish the two groups in our notation.

We define the *surplus* $\mathcal{S}(\omega)$ associated with allocation $\omega \in \Omega$ by

$$\mathcal{S}(\omega) = \sum_{n \in N} u_n(x_n) - \sum_{f \in F} c_f(y_f).$$

The *efficient allocation* problem is to solve

$$\max_{\omega \in \Omega} \mathcal{S}(\omega), \tag{P}$$

with a solution denoted by $\omega^* \in \arg\max_{\omega \in \Omega}$ and the resulting surplus $\mathcal{S}^* = \mathcal{S}(\omega^*)$.

For any allocation $\omega \in \Omega$, we will refer to $\mathcal{S}(\omega) - \mathcal{S}^*$ as the *deadweight loss* of ω and the ratio $\frac{\mathcal{S}(\omega) - \mathcal{S}^*}{\mathcal{S}^*}$ as the *percentage loss* at ω.[9]

Pricing rules To extend the Walrasian mechanism to a wider class of *anonymous linear prices*, we allow two different price vectors $p^b, p^s \in \mathbb{R}^L_+$ for buyers and sellers such that if any buyer n purchases a bundle x, it makes a payment of $t = p^b \cdot x$ and if seller f supplies y, it receives a payment of $t = p^s \cdot y$.

Denote buyer n's utility by $U_n(x, p^b \cdot x) = u_n(x) - p^b \cdot x$. Its *demand correspondence* $D_n : \mathbb{R}^L_+ \rightrightarrows X$ maps each price vector p^b to the set of utility-maximizing bundles $D_n(p^b)$. Its *indirect utility* function is $\hat{u}_n(p^b) = \max_{x \in X} u_n(x) - p^b \cdot x$. Similarly, denote seller f's profit by $\pi_f(y, p^s \cdot y)$ and its *supply correspondence* by $S_f : \mathbb{R}^L_+ \rightrightarrows Y$, which maps each price vector p^s to the set of profit-maximizing bundles $S_f(p^s)$. Its *indirect profit* function is $\hat{\pi}_f(p^s) = \max_{y \in Y} p^s \cdot y - c_f(y)$.

2.2 Convex quasilinear economies

Convexity is defined with respect to the set of payoff-improving allocations for an agent in the economy. The *$\bar{u}$-upper contour set* of buyer $n \in N$ is defined by

$$UC_n^{\bar{u}} = \{(x, t) \in X \times \mathbb{R} : U_n(x, t) \geq \bar{u}\},$$

while the *$\bar{\pi}$-upper contour set* of seller $f \in F$ is given by

$$UC_f^{\bar{\pi}} = \{(y, t) \in Y \times \mathbb{R} : \pi_f(y, t) \geq \bar{\pi}\}.$$

[9]Later we will make assumptions to rule out cases where $\mathcal{S}^* = 0$ so that this ratio is well-defined.

We say that buyer n has *convex preferences* if X is convex and $UC_n^{\bar{u}}$ is convex for all $\bar{u} \in \mathbb{R}$, which is equivalent to the quasiconcavity of U_n and the concavity of the valuation function v_n. Seller f has *convex technology* if Y is convex and $UC_f^{\bar{\pi}}$ is convex for all $\bar{\pi} \in \mathbb{R}$, which is equivalent to the quasiconcavity of π_f and the convexity of cost function c_f.

Under the assumption of quasilinearity and the convexity of agents' preferences and technologies, we have the following statement of the fundamental welfare theorems of Arrow (1951) and Debreu (1951). [10]

Proposition 1. *Suppose in (quasilinear) economy $\mathscr{E}$ that all buyers $n \in N$ have convex preferences and all sellers $f \in F$ have convex technologies. Then a feasible allocation $\omega \in \Omega$ is efficient if and only if there exists $p \in \mathbb{R}_+^L, p \neq 0$ such that for all $n \in N$, $x_n \in D_n(p)$; for all $f \in F$, $y_f \in S_f(p)$; and $\sum_{n \in N} p \cdot x_n = \sum_{f \in F} p \cdot y_f$. The pair (p, ω) is a* competitive *or* Walrasian equilibrium*.*

2.3 Measures of non-convexity and approximate equilibria

The non-convexity of a set S can be measured in several ways by comparing S and $\text{co}(S)$. We will work with the following measures of non-convexity of a set:

- The *inner radius* of S is $r(S) = \sup_{x \in \text{co}(S)} \inf_{T \subseteq S : x \in \text{co}(T)} \text{rad}(T)$.
- The *inner distance* of S is $\rho(S) = \sup_{x \in \text{co}(S)} \inf_{y \in S} \|x - y\|$.

For a convex set S, $r(S) = 0 = \rho(S)$. The two functions, which are illustrated in Figure 2, measure the size of the set of points in $\text{co}(S)$ but missing from S.

The non-convexity of the preferences of buyer $n \in N$ may be measured by the largest inner radius or inner distance of their upper contour sets, that is $r_n = \sup_{\bar{u} \in \mathbb{R}} r(UC_n^{\bar{u}})$ or $\rho_n = \sup_{\bar{u} \in \mathbb{R}} \rho(UC_n^{\bar{u}})$. Similarly, the non-convexity of the technology of seller $f \in F$ may be measured by $r_f = \sup_{\bar{\pi} \in \mathbb{R}} r(UC_f^{\bar{\pi}})$ or $\rho_f = \sup_{\bar{\pi} \in \mathbb{R}} \rho(UC_f^{\bar{\pi}})$. Let $r_{\mathscr{E}}$ and $\rho_{\mathscr{E}}$ denote the largest of such measures among all the buyers and sellers in economy $\mathscr{E}$.

[10] The statement of Proposition 1 is stronger than the classic statements of the welfare theorems in the 'only if' direction, which is possible due to the quasilinear form of the utility and profit functions. Without quasilinearity or an alternative assumption, it may only be possible to find prices so that agents are expenditure-minimizing for a given level of utility or profit, that is, a price quasiequilibrium with transfers. With quasilinearity and convexity, the efficient allocation program is a convex program. Since there exists a feasible allocation, Slater's Theorem (see, for example, Boyd and Vandenberghe (2004)) implies strong duality. A solution p^* to the dual program, $\inf_{p \in \mathbb{R}_+^L} \sum_{n \in N} \hat{u}_n(p) + \sum_{f \in F} \hat{\pi}_f(p)$, and an efficient allocation ω^* comprise a saddle point for the Lagrangian $\mathscr{L}(\omega, p) = \sum_{n \in N} u_n(x_n) - \sum_{f \in F} c_f(y_f) - p \cdot \left(\sum_{n \in N} x_n - \sum_{f \in F} y_f\right)$, so that for any $\omega \in \Omega$, $\mathscr{L}(\omega, p^*) \leq \mathscr{L}(\omega^*, p^*)$. Because the Lagrangian is separable across agents, the saddle point condition implies $x_n^* \in D_n(p^*)$ and $y_f^* \in S_f(p^*)$.

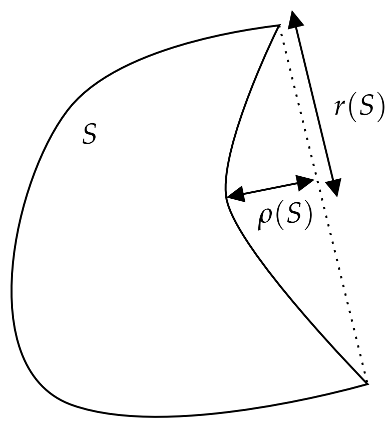

Figure 2: Measures of non-convexity of a set

When agents' upper contour sets are not convex, the second welfare theorem may not hold and there may be no competitive equilibrium. Proposition 2, the Shapley-Folkman Lemma, assists in identifying allocations which are nearly competitive equilibria.

Proposition 2 (Shapley-Folkman Lemma[11])**.** *Let $S_i \subseteq \mathbb{R}^L$ for $i = 1, ..., M$, $S = \bigoplus_{i=1}^{M} S_i$ and $L' = \min(L, M)$. Then for any $x \in \mathrm{co}(S)$, $x = \sum_{i=1}^{N} x_i$ where $x_i \in \mathrm{co}(S_i)$ and $|i : x_i \in \mathrm{co}(S_i) \setminus S_i| \leq L'$. Moreover, there exists $y, y' \in S$ such that $\|x - y\| \leq r_{\mathscr{E}} \sqrt{L'}$ and $\|x - y'\| \leq \rho_{\mathscr{E}} L'$.*

Proposition 2 has been used to establish results about *approximate equilibria*, which are constructed as follows. First, consider a convexified version of the non-convex economy in which the upper contour sets of all agents are replaced by their convex hulls. This is equivalent to replacing the buyers' valuation functions by their upper concave envelopes and sellers' cost functions by their lower convex envelopes. The *convexified economy* $\hat{\mathscr{E}}$ is then

$$\hat{\mathscr{E}} = \langle N, (\mathrm{cav}(u_n))_{n \in N}, F, (\mathrm{vex}(c_f))_{f \in F} \rangle.$$

The convexified economy has a competitive equilibrium which is efficient (according to the concavified valuation functions and convexified cost functions) by Proposition 1. Since the convexified economy's efficient allocation problem is a relaxation[12] of the same problem for the original economy, the efficient surplus of the convexified economy is an upper bound on the efficient sur-

[11]It is perhaps most accurate to refer only to the result in the second sentence of Proposition 2 as the Shapley-Folkman Lemma, althought it was first reported by Starr (1969) as a result of private communication with Lloyd Shapley and Jon Folkman. Starr (1969) then proved the first half of sentence three of Proposition 2, while Heller (1972) proved the second half. For simplicity, we will refer to the whole of Proposition 2 as the Shapley-Folkman Lemma.

[12]That is, the constraint set is weakly larger than the original constraint set and the objective function is pointwise weakly larger than the original objective.

plus of the original economy. We call the resulting price-allocation pair (p, ω) a *pseudoequilibrium* of the actual economy $\mathscr{E}$. Proposition 2 implies that ω can be chosen so at most L' agents in $\mathscr{E}$ are not utility- or profit-maximizing at ω given prices p and that there is a nearby allocation ω' such that *all* agents are maximizing given prices p,[13] but markets may not exactly clear at ω'. The price-allocation pair (p, ω') is called an *approximate equilibrium*.

Pseudoequilibria and approximate equilibria describe allocations rather than mechanisms. These allocations can be infeasible or may impose large losses on some agents. We utilize these ideas to devise mechanisms that are computable, select feasible allocations, and have the other desirable properties that we listed earlier.

3 Rationing, inefficiency and incentives

When competitive equilibrium does not exist, no feasible allocation is supported by a single anonymous price vector that is the same for buyers and sellers. Feasibility can be restored only by varying the price vector for one side of the market or by rationing some agents, which means requiring them to accept a bundle that is not most preferred at the specified prices. Given an allocation and prices, we measure the extent of rationing for buyers and sellers by their *rationing losses*, which are defined as the excess of the payoff an agent would obtain from its most preferred bundle given the prices compared to the payoff it receives in the prescribed allocation.

Definition 3.1. The *rationing loss $\mathscr{R}_n(p, x)$ of buyer n* at price p and allocation x is

$$\mathscr{R}_n(p, x) = \hat{u}_n(p) - U_n(x, p \cdot x).$$

The *rationing loss $\mathscr{R}_f(p, y)$ of seller f* is

$$\mathscr{R}_f(p, y) = \hat{\pi}_f(p) - \pi_f(y, p \cdot y).$$

[13] To see this, note that if a buyer is assigned a bundle x_n in ω that is not utility-maximizing at p, then x_n must be the convex combination of bundles (x_n') in X which are *exposed points* in u_n (i.e. where $\text{cav}(u_n) = u_n$), and that the agents in the convexified economy must be indifferent between x_n and these bundles. That is, the concavified portions of buyers' utility functions consist of (patches of) hyperplanes, and if an agent is assigned a bundle on such a patch, then the price vector must be normal to that hyperplane. This implies that the original buyer must be maximizing at bundles in (x_n'), which are on the relative boundaries of the patch of hyperplane.

The *rationing losses* of allocation $\omega = ((x_n)_{n\in N}, (y_f)_{f\in F})$ at price p is defined by

$$\mathscr{R}(p,\omega) = \sum_{n\in N} \mathscr{R}_n(p, x_n) + \sum_{f\in F} \mathscr{R}_f(p, y_f).$$

If competitive equilibrium does not exist, any price-allocation pair must entail rationing or wasted supply (and thus budget deficit) or both. In our first main result, we show that the extent of such rationing and budget losses fully characterize the efficiency of the allocation.

Theorem 1 (Bound-Form First Welfare Theorem)**.** *Let $p \in \mathbb{R}^L_+$ be a price vector and $\omega = ((x_n)_{n\in N}, (y_f)_{f\in F})$ be any allocation. Then, the deadweight loss of allocation ω satisfies*

$$\underbrace{\mathcal{S}^* - \mathcal{S}(\omega)}_{\textit{deadweight loss}} \leq \underbrace{\mathscr{R}(p,\omega)}_{\textit{rationing loss}} + \underbrace{p\cdot\left(\sum_{f\in F} y_f - \sum_{n\in N} x_n\right)}_{\textit{budget deficit}}.$$

Proof. Fix any efficient allocation ω^*. By the definitions of indirect utility and consumer surplus, the following must hold for any prices:

$$\hat{u}_n(p) \geq u_n(x_n^*) - p\cdot x_n^*$$
$$\hat{\pi}_f(p) \geq p\cdot y_f^* - c_f(y_f^*).$$

Summing these inequalities, we obtain

$$\sum_{n\in N} \hat{u}_n(p) + \sum_{f\in F} \hat{\pi}_f(p) + p\cdot\left(\sum_{n\in N} x_n^* - \sum_{f\in F} y_f^*\right) \geq \sum_{n\in N} u_n(x_n^*) - \sum_{f\in F} c_f(y_f^*) = \mathcal{S}^*.$$

Since ω^* is feasible, the third term on the left side is nonpositive, so

$$\sum_{n\in N} \hat{u}_n(p) + \sum_{f\in F} \hat{\pi}_f(p) \geq \mathcal{S}^*$$

Subtracting $\mathcal{S}(\omega)$ and applying the definitions,

$$\mathcal{S}^* - \mathcal{S}(\omega) \leq \sum_{n\in N} \hat{u}_n(p) + \sum_{f\in F} \hat{\pi}_f(p) - \mathcal{S}(\omega) = \mathscr{R}(p,\omega) + p\cdot\left(\sum_{f\in F} y_f - \sum_{n\in N} x_n\right),$$

which is what we sought to prove. $\square$

The Bound-Form First Welfare Theorem extends the First Welfare Theorem for quasilinear economies by applying to any price-allocation pair (p, ω) rather than just to competitive equilibria. If (p, ω) is such an equilibrium, then both the budget deficit and the rationing losses are zero, so the theorem asserts that the welfare loss is zero, or equivalently that any Walrasian equilibrium is efficient.

One interpretation of the First Welfare Theorem is that prices act as a "certificate of optimality": if supporting prices exist for some allocation, then that allocation is efficient. As Scarf (1994) lamented, in the absence of convexity, there is, in general, no such optimality test. However, one interpretation of Theorem 1 is as an *approximate* optimality test: if we can have a price-allocation pair for which the rationed surplus plus the net budget deficit is small, then the given allocation is approximately efficient. Theorem 1 is also begins to link incentives to efficiency: since for any fixed price, agents would prefer not to be rationed, Theorem 1 suggests that a pricing mechanism with little rationing and in which individual agents have little influence over prices will have both good incentive properties and small deadweight losses. This observation is the key to our extensions in Sections 4 and 5.

4 Two-sided non-convexities: the markup mechanism

4.1 Pricing mechanisms and approximate mechanism design

We study direct pricing mechanisms that map profiles of reports of sellers' cost functions $(c_f)_{f \in F}$ and buyers' value functions $(u_n)_{n \in N}$ to a feasible allocation $\omega \in \Omega$ and anonymous prices for buyers and sellers, $\boldsymbol{p} = (p^b, p^s)$.[14] We also require the mechanism to be *budget-feasible* so that for all report profiles, $p^s \cdot \sum_{f \in F} y_f \leq p^b \cdot \sum_{n \in N} x_n$. We do not delve into the important question of how agents communicate their potentially complicated costs and values to the mechanism.[15]

For convenience, we restate some familiar definitions. A pricing mechanism is

(a) *efficient* if the output allocation ω is an efficient allocation given the reported value and cost functions,

[14] We choose not to consider randomized mechanisms, both because these are unnecessary to achieve our objectives and because they raise daunting practical issues, including most importantly very high trust requirements in the mechanism designer and the possible failure of ex post individual rationality for ex ante individually rational lotteries.

[15] The design of reporting languages to report complex preferences for economic mechanisms has been studied by Milgrom (2009), Bichler, Goeree, Mayer, and Shabalin (2014), Bichler, Milgrom, and Schwarz (2022) and others.

(b) *ex post incentive-compatible (EPIC)* if truthful reporting is an *ex post* Nash equilibrium of the reporting game induced by the mechanism,

(c) *interim incentive-compatible (IIC)* if truthful reporting maximizes each agents' expected payoffs under the mechanism, and

(d) *individually-rational (IR)* if, given reported value and cost functions, the allocation and prices determined by the mechanism delivers each agent a payoff (utility or profit) no worse than non-participation (here, 0).

It is typically impossible for a mechanism to exactly satisfy these desirable properties; instead, we seek mechanisms satisfying appropriate approximations to these goals. A pricing mechanism is

(a) *ε-nearly-efficient* if the deadweight loss of ω is bounded by ε given the reports,

(b) *ε-EPIC* if truthful reporting is an *ε-ex post* Nash equilibrium,

(c) *ε-IIC* if, for each agent, the expected payoff associated with any report in the mechanism is no more than ε greater than that of the truthful report, and

(d) *ε-individually-rational* if each agent obtains a payoff of at least $-\varepsilon$ given its report.

A well-known pricing mechanism is the *Walrasian mechanism*, which inputs reports of value and cost functions and outputs a Walrasian equilibrium price and allocation. There may be multiple Walrasian equilibria - in this case, we suppose that the mechanism designer has a predetermined rule for selecting among them (and we make a similar assumption when multiplicity arises in the other mechanisms discussed in this paper). Walrasian mechanisms are efficient, individually rational and have good large-market incentive properties (discussed in Section 4.3).

We are interested in extensions of the Walrasian mechanism that also perform well in large markets, so that these approximations depend on $|N|$. Let $\mathscr{E}_t = \langle N_t, \mu_t, \phi_t, \chi_t \rangle$ be a sequence of economies indexed by $t = 1, 2, ...$ We make the following additional assumptions as the economy grows large.

Assumption 1 (Existence of limit economy). *As $t \to \infty$, $|N_t| \to \infty$ and $\phi_t \to \phi \in (0,1)$. Furthermore, μ_t converges weakly to probability measure μ_∞ on $\mathscr{U}$ and χ_t converges weakly to measure χ_∞ on $\mathscr{C}$.*

Assumption 2 (Individual non-convexities are bounded). *There exists $R > 0$ with $R_{\mathscr{E}_t} < R$ for all t.*

Assumption 3 (Growing gains from trade). *As $t \to \infty$, the efficient surplus $\mathcal{S}_t^*$ grows at least as quickly as $|N_t|$ asymptotically, that is, $\mathcal{S}_t^* = \Omega(|N_t|)$. Since utilities are bounded, equivalently, $\mathcal{S}_t^* = \Theta(|N_t|)$.*

Assumption 4 (Prices are bounded). *There exists $M > 0$ such that $d_H\left(\sum_{n \in N_t} D_n(p), \sum_{f \in F_t} S_f(p)\right) \geq R$ for sufficiently large t and $p \in \mathbb{R}_+^L$ such that $\|p\| < \frac{1}{M}$ or $\|p\| > M$.*

Assumption 1 asserts that the only important variation in economies in the sequence is their scale: the proportions of various types converge to a limit. Assumption 2 asserts that there is a uniform bound on the size of any non-convexity across the sequence of markets, so no single firm's or consumer's non-convexity can be problematic in a large economy.[16] Assumption 3 is the condition that the efficient surplus per participant is bounded away from zero. Assumption 4 implies the compactness of the set of possible prices for which the aggregate demand and supply to be close to equal, so that pseudoequilibrium prices do not tend to zero or infinity in large economies.

4.2 Markup mechanisms

We now introduce our first class of new mechanisms, the *markup mechanisms*. These mechanisms maintain the no-rationing property of the Walrasian mechanism. If agents are unrationed, it may be impossible to find allocations in which all supply is demanded by buyers. In order to permit for this non-market clearing, a markup on the prices paid by buyers is used to cover payments for any unallocated supply.

Definition 4.1 (Markup mechanisms). A *markup mechanism* outputs a Walrasian equilibrium if it exists, or else a triple (α, p, ω) consisting of a markup parameter $\alpha \geq 0$, a price $p \in \mathbb{R}_+^L$ and a feasible allocation ω such that:

(a) payments for sellers are determined by price vector p and sellers are not rationed given these prices, so $y_f \in S_f(p)$;

(b) payments for buyers are determined by price vector $(1+\alpha)p$ and buyers are not rationed given these prices, so $x_n \in D_n((1+\alpha)p)$; and

(c) budgets are balanced, so $\sum_{n \in N}(1+\alpha)p \cdot x_n - \sum_{f \in F} p \cdot y_f \geq 0$.

[16] It suffices here to assume that the consumption possibility set X does not grow with t, since $R \leq \operatorname{rad}(X)$.

In order to minimize losses, a markup mechanism would ideally use an α close to zero and leave few goods unallocated. This follows by applying Theorem 1 at the price-allocation pair (p, ω): if few goods are unallocated, the budget deficit at price p is small, while if prices p and $(1+\alpha)p$ are close, the rationing losses for each buyer at price p are small. This latter claim follows by an envelope theorem argument, which formalized in Proposition 3.

Proposition 3. *Let bundle x be demanded by buyer n at price p, so $x \in D_n(p)$. Consider some other price $p' \neq p$. Then the rationing loss of buyer n at allocation x given price p', $\mathscr{R}_n(x, p')$ is $O(\|p - p'\|)$.*

If computational challenges were not a concern, a market designer may seek to identify a markup mechanism with the smallest loss, which we call a *minimum markup mechanism*. Zero markups and zero budget imbalance are exactly achievable whenever Walrasian equilibria exist, so that the minimum markup mechanism reduces in that case to a Walrasian mechanism. However, solving for the minimum markup mechanism in non-convex economies can be a difficult non-convex optimization problem, requiring optimization over both the space of allocations and prices. We now show that an $O(1)$-nearly-efficient markup mechanism can be identified using only convex optimization problems and a one-dimensional binary search. Before providing a technical description of this pricing mechanism and its analysis, we sketch intuitively the steps of our approach.

For a fixed α, we select (p, ω) to be the equilibrium price-and-allocation pair of a related economy with three changes from the actual economy: (1) every buyers' value function is replaced by the smaller function $u_n/(1+\alpha)$, (2) all values and costs are then replaced by their concave or convex hulls, respectively,[17] and (3) we add demand for each good by the auctioneer in the amount of $R := \min\{r_{\mathscr{G}}\sqrt{L}, \rho_{\mathscr{G}} L\}$. Step (1) in this construction implies that when buyers' demands are computed using their actual value functions u_n and their marked-up prices of $(1+\alpha)p$, the buyers' demands are just as computed in the related economy. We apply the Shapley-Folkman Lemma (Proposition 2) to round ω to one of the demanded allocations for each agent, while changing the net demand for each good by at most R units. To maintain feasibility after that change, we balance with an offsetting change from the units allocated to the auctioneer in step (3).

Excluding the auctioneer, the resulting final allocation always has supply greater than or equal to demand—it is feasible—and its excess supply is no more than $2R$ units of each good. Any

[17]This need not be computationally expensive. For example, if the value and cost functions are reported to the mechanism using a mixed integer program, the mechanism may simply convert integer variables to real variables to obtain the convex hulls in the form of linear or quadratic programs.

excess supply can result in a loss of efficiency (units allocated to the auctioneer are wasted), but the quantity allocated to the auctioneer is bounded by a constant, that is, by an amount that is independent of the size of the market.[18]

Since the excess supply of goods is bounded, the budget imbalance at price p is bounded as well. As trade increases with the size of the economy, the markup, α, needed to guarantee budget balance is inversely related to market size: it is $O(1/N)$. Thus, the total welfare loss from the markup mechanism is bounded by a constant plus a term that is inversely proportional to market size.

Definition 4.2 (simple markup mechanism)**.** The *simple markup mechanism* is the markup mechanism with parameters $(\alpha^*, p^*, \omega^*)$ determined as follows. If a Walrasian equilibrium exists according to the reported preferences, set $\alpha^* = 0$ and choose (p^*, ω^*) to be some Walrasian equilibrium. Otherwise, for each $\alpha > 0$, consider the following convex program:

$$\min_{p \in \mathbb{R}^L_+} \max_{x_n \in \mathrm{co}(X), y_f \in \mathrm{co}(Y)} \sum_{n \in N} \frac{\mathrm{cav}(u_n)(x_n)}{1+\alpha} - \sum_{f \in F} \mathrm{vex}(c_f)(y_f) - p \cdot \left(\sum_{n \in N} x_n + R 1_L - \sum_{f \in F} y_f \right),$$

where 1_L is the vector of ones in $\mathbb{R}^L$. Let $(p^\alpha, \tilde{\omega}^\alpha)$ denote any solution to this program.

From $\tilde{\omega}^\alpha$, we obtain, via Proposition 2, an allocation ω^α with $\|\omega^\alpha - \tilde{\omega}^\alpha\| \leq R$ such that $x_n^\alpha \in \max_{x \in X} \frac{1}{1+\alpha} u_n(x) - p \cdot x$ for each $n \in N$ and $y_f^\alpha \in S_f(p)$. By construction, this ω^α will be feasible in $\mathcal{E}$. Let

$$\alpha^* = \min \left\{ \alpha \,\middle|\, \sum_{n \in N} (1+\alpha) p^\alpha \cdot x_n^\alpha - \sum_{f \in F} p^\alpha \cdot y_f^\alpha \geq 0 \right\}, \tag{A}$$

and define $p^* = p^{\alpha^*}$ and $\omega^* = \omega^{\alpha^*}$.

In Theorem 2, we show that the simple markup mechanism is well-defined (that is, the minimum in (A) exists) and the resulting mechanism is $O(1/|N_t|)$-near-efficient. Before doing so, we note some other important properties of the markup mechanism that follow directly from the construction. First, the equilibrium is physically feasible and weakly budget-balanced. Second, the equilibrium allocation and payments are individually rational for each agent. For sellers, this follows because their profits are identical to those in the pseudoequilibrium used in the construction.

[18]The choice of R units of each good as a set-aside for the auctioneer in step (3) is a theoretical guarantee. It might be possible to allocate fewer units to the auctioneer in step (3) and arrive at a more efficient feasible allocation using the same approach. We note that an alternative approach could be to start by checking for feasible allocations with zero units set aside (these would correspond to competitive equilibria) and then increase the set-aside gradually until a budget-feasible markup mechanism is identified, but we leave such details for future research.

For buyers, the pseudoequilibrium price p^α and consumption allocation in $\tilde{\omega}^\alpha$ satisfy

$$\frac{1}{1+\alpha}u_n(\tilde{x}_n^\alpha) - p^\alpha \cdot \tilde{x}_n^\alpha = \frac{1}{1+\alpha}u_n(x_n^\alpha) - p^\alpha \cdot x_n^\alpha \geq 0$$

so that $u_n(x_n^\alpha) - (1+\alpha)p^\alpha \cdot x_n^\alpha \geq 0$ as well.

Theorem 2. *Let $\mathscr{E}_t$ be a sequence of economies satisfying Assumptions 1–4. Then*

(a) the simple markup mechanism is well-defined (that, is the minimum in (A) *is attained),*

(b) the simple markup mechanism's markup $\alpha^* \leq O(1/|N_t|)$, and

(c) the deadweight loss of the simple markup mechanism's allocation is $O(1)$, so that the relative inefficiency is $O\left(1/|N_t|\right)$.

Although the rates of convergence in Theorem 2 are stated in terms of $|N_t|$, by Assumption 1, the same asymptotic rate of convergence holds with respect to $|F_t|$ or $|A_t|$.

4.3 Incentives

In the Walrasian and markup mechanisms, both buyers and sellers receive their optimal bundles given their prices, so an agent can profit from false reports only to the extent that it can influence its prices. Moreover, because the prices in the simple markup mechanism are Walrasian equilibrium prices of a related convex economy, the limited ability of agents to manipulate Walrasian prices in large economies implies a similar difficulty for the markup mechanism, which is the underpinning of our incentive analysis.

We now briefly discuss the most relevant literature related to the agents' ability to influence prices in large markets. Roberts and Postlewaite (1976) initiated the formal literature of this subject in a study of a sequence of exchange economies with the number of agents going to infinity. They represented the sequence of economies by measures μ_t on $\mathscr{U}$ with $\lim_{t\to\infty} \mu_t = \mu_\infty$, showing that if the Walrasian price correspondence is continuous at μ_∞, then each agent's influence on the price goes to zero as t increases. Jackson (1992) shows in the same model that an agent's optimal reported demand converges in the L^∞ norm to the demand associated with that agent's true preferences.

Watt (2022) studies the rate of this convergence, showing that a condition on the demand and supply correspondences, called *strong monotonicity*, ensures fast convergence of incentives in the Walrasian mechanism.

Definition 4.3 (Strong monotonicity)**.**

(a) A buyer n is *active* at price p if $D_n(p) \neq \{0\}$.

(b) Buyer n's demand correspondence $D_n : \mathbb{R}^L_+ \rightrightarrows X$ is *strongly monotone* if there exists some $m > 0$ such that for all p, p' at which buyer n is active and for all $d \in D_n(p), d' \in D_n(p')$,

$$(p - p') \cdot (d' - d) \geq m\|p - p'\|^2.$$

(c) A seller f is *active* at price p if $S_f(p) \neq \{0\}$ and there is some $\beta > 1$ such that $S_f(\beta p) \neq S_f(p)$.

(d) Seller f's supply correspondence $S_f : \mathbb{R}^L_+ \rightrightarrows Y$ is *strongly monotone* if for all p, p' at which seller f is active and for all $s \in S(p), s' \in S(p')$,

$$(p - p') \cdot (s - s') \geq m\|p - p'\|^2.$$

Strong monotonicity is a condition on how quickly demand or supply changes in response to price changes: in settings with one good, strong monotonicity is equivalent to a lower bound on the slope of the firm's supply curve and an upper bound on the slope of a buyer's demand curve. If each buyer has strongly monotone demand and each seller has strongly monotone supply, Watt (2022) shows that the resulting sequence of economies is *perturbation-proof*, which implies that the maximum influence of any one agent on Walrasian prices is $O\left(1/|N^a_t|\right)$, where N^a_t is the number of active agents at the Walrasian price. Furthermore, if each buyer is drawn independently at random from some distribution μ over $\mathcal{U}$ and each seller is drawn independently at random from χ over $\mathcal{C}$ for which the *expected* demand and supply correspondences are strongly monotone, the maximum benefit of misreporting under the Walrasian mechanism is $O_P\left(1/|N_t|^{1-\varepsilon}\right)$ for all $\varepsilon > 0$.[19]

While the above results considered ex post incentives, Azevedo and Budish (2019) studied interim incentives and show that the Walrasian mechanism (with a finite set of buyer and seller types) is *strategy-proof in the large*, which implies that the benefit to any agent of misreporting against any full-support, independent and identically-distributed distribution of agent types tends to zero at $O(1/|N_t|^{\frac{1}{2}-\varepsilon})$ for all $\varepsilon > 0$.

Theorem 3, proved in Appendix B.3, adapts these results to the markup mechanism.

[19] Here the expectation is with respect to the measure over economies induced by draws of agents from μ and χ. We clarify the meaning of the "expected demand and supply correspondences" in Appendix B.3.

Theorem 3. *Let $(\mathscr{E}_t)_{t\in\mathbb{N}} = (\langle N_t, \mu_t, \phi_t, \chi_t\rangle)_{t\in\mathbb{N}}$ be a sequence of economies satisfying Assumptions 1–4 and suppose that a markup mechanism with markups α_t is applied to $\mathscr{E}_t$.*[20]

(a) Suppose $(\mathscr{E}_t)_{t\in\mathbb{N}}$ are IPV economies drawn according to full-support probability distributions μ and χ defined on finite *type spaces $\mathscr{U}$ and $\mathscr{C}$. Then the markup mechanism is strategy-proof-in-the-large and $O(1/|N_t|^{\frac{1}{2}-\varepsilon})$-IIC for any $\varepsilon > 0$.*

Now suppose that the markup satisfies $\alpha_t \leq O(1/|N_t|)$ as in the minimal and simple markup mechanisms.

(b) Suppose that each buyer in each $\mathscr{E}_t$ has strongly monotone demand and that each seller has strongly monotone supply. Then the markup mechanism is $O(1/|N_t^a|)$-EPIC, where N_t^a is the number of active buyers and sellers at the mechanism's respective prices.

(c) Suppose that $(\mathscr{E}_t)_{t\in\mathbb{N}}$ are IPV economies drawn from distributions μ and χ for which the expected demand and supply correspondences are strongly monotone. Then, the maximum ex post benefit of misreporting for any agent is $O_P\left(1/|N_t|^{1-\varepsilon}\right)$, for any $\varepsilon > 0$, and the mechanism is $O(1/|N_t|^{1-\varepsilon})$-IIC.

4.4 Computational properties

While equilibrium computation is hard in general,[21] the Walrasian equilibrium problem in concave quasilinear economies reduces to a convex optimization problem. A wide class of such optimization problems are efficiently solvable, including problems with self-concordant or strongly convex objectives. For example, Walrasian prices in economies with strongly monotone supply and demand may be efficiently computed via tâtonnement (Watt, 2022).

In contrast, efficient allocation problems in many non-convex economies are known to be computationally complex. For example, the problem of identifying an optimal allocation in the fisheries market of Bichler et al. (2018) involved solving a large integer programming problem, which is known to be NP-complete. No efficient algorithms are known for such problems, although heuristic approaches may make specific instances of such problems tractable.

Conditional on α, our simple markup mechanisms—like the Walrasian mechanism—solve only convex optimization problems. Identifying the optimal markup α^* in the simple markup mechanism may be more challenging, although a binary search algorithm for α might be em-

[20] The markups may be determined endogenously, as in the minimal and simple markup mechanisms.
[21] See, for example, Chen, Dai, Du, and Teng (2009) and Daskalakis, Goldberg, and Papadimitriou (2009).

ployed, in practice, to identify some markup that ensures weak budget balance. Since the loss of the simple markup mechanism is $O(\alpha)$, small markups are associated with small losses.

4.5 Alternative mechanisms

Our markup mechanisms are related to a linear pricing mechanism which was proposed for allocating commercial fisheries licenses in New South Wales: the *maximum surplus anonymous pricing* mechanism described by Bichler et al. (2018).[22] That mechanism solves the usual surplus optimization problem

$$\max_{\omega \in \Omega} \sum_{n \in N} u_n(x_n) - \sum_{f \in F} c_f(y_f),$$

but subject to the constraint that there exist prices p^b and p^s satisfying

(a) Individual rationality: for all $n \in N, f \in F$, $u_n(x_n) - p^b \cdot x_n \geq 0$ and $p^s \cdot y_f - c_f(y_f) \geq 0$.

(b) Budget balance: $\sum_{n \in N} p^b \cdot x_n \geq \sum_{f \in F} p^s \cdot y_f$.

A corollary of Theorem 2 is that this mechanism has a deadweight loss that is bounded by a constant independently of the market size $|N|$. We have focused our analysis on the class of markup mechanisms for several reasons. Firstly, the maximum surplus anonymous pricing mechanism requires the solution of non-convex optimization problems, while the markup mechanism can be approximated by solving convex optimization problems (plus a binary search). Secondly, optimizers of the alternative mechanisms may involve rationing at the prevailing prices, which as we have already seen can give agents an additional incentive for false reporting that is avoided by the markup mechanism.

5 One-sided non-convexities: the rationing mechanism

In this section, we show that it is sometimes possible to preserve large-market incentives and significantly reduce the welfare loss by adopting a certain rationing mechanism. These mechanisms identify a price-allocation pair in which the rationing losses are small and the budget is exactly balanced so that the expression on the right-hand-side of the Bound-Form First Welfare Theorem is small.

[22] A variant of this mechanism was later implemented in New South Wales.

Definition 5.1 (Rationing mechanisms)**.** A *rationing mechanism* outputs a Walrasian equilibrium if one exists, or else a price and feasible allocation (p, ω) such that $p \cdot \sum_{n \in N} x_n = p \cdot \sum_{f \in F} y_f$ with payments determined for buyers and sellers by price vector p. Allocations need not lie in the demand sets of buyers or the supply sets of sellers given price p. A *minimum-loss rationing mechanism* is a rationing mechanism with the least total rationing losses.

Rationing mechanisms could be applied in settings with two-sided non-convexities, but we focus our analysis in this section on economies in which agents on one side of the market (for concreteness, the buyers) have convex preferences.[23]

Assumption 5 (Buyers' Side Convexity)**.** *Each valuation $u_n \in \mathcal{U}$ is concave and the consumption possibility set X is a Cartesian product of intervals $\prod_{\ell=1}^{L} [0, \bar{x}^\ell]$.*[24]

There are two reasons to conjecture that one-sided convexity might allow for more efficient pricing mechanisms. First, the markup mechanism handles non-convexity by wasting some supply, but convexity on the buyers' side ensures that whatever quantities producers may supply, it is possible to identify allocations in which buyers make valuable use of the total supply. Second, in other specific contexts, the total loss of welfare from some mechanisms has been shown to converge to zero as the economy grows. For example, in the double auction model of Rustichini, Satterthwaite, and Williams (1994), the expected *total* welfare loss is just $O\left(1/|N|\right)$, which is better than the $O(1)$ total loss of the markup mechanism. Also simulation studies of the maximum surplus anonymous pricing mechanism of Bichler et al. (2018) find suggestive evidence of similarly small deadweight losses.

The minimum-loss rationing mechanism may be challenging to compute. In fact, the following proposition shows that its computation is at least as hard as the associated efficient allocation problem.

Proposition 4. *The allocation ω outputted by the minimum-loss rationing mechanism is efficient.*

We now identify a *buyers-only rationing mechanism* that uses only two convex optimization problems and has total deadweight losses converging quickly to zero in economies satisfying strong monotonicity assumptions.

[23]One reason for this focus is that in our formulation of markets with two-sided non-convexities, exactly market-clearing feasible allocations may not even exist which makes obtaining budget balance with non-zero prices impossible. If we allowed for free disposal, this concern disappears, but we have not been able to improve on the $O(1)-$near-efficiency result of the simple markup mechanism even under this assumption.

[24]The assumption on the shape of the consumption possibility set X plays a minor technical role in the proof of Theorem 4.

Definition 5.2 (Buyers-only rationing mechanism)**.** The *buyers-only rationing mechanism* inputs reports of values and costs from sellers and outputs the price-allocation pair $(p, \mathbf{x}', \mathbf{y})$ determined as follows.

Identify $(p, \tilde{\omega})$, a solution to

$$\min_{p \in \mathbb{R}^L_+} \max_{\omega \in X^{|N|} \times Y^{|F|}} \sum_{n \in N} u_n(x_n) - \sum_{f \in F} \mathrm{vex}(c_f)(y_f) - p \cdot \left(\sum_{n \in N} x_n - \sum_{f \in F} y_f - R_{\mathscr{E}} 1_L \right). \tag{ARM1}$$

Using Proposition 2, identify $\omega = (\mathbf{x}, \mathbf{y})$ satisfying $\|\tilde{\omega} - \omega\| \leq R_{\mathscr{E}}$ with $y_f \in S_f(p)$ for each $f \in F$ and let $s = \sum_{f \in F} y_f$ be the resulting supply vector.

Now identify $\mathbf{x}' = (x'_n)_{n \in N}$ solving

$$\max_{(x_n)_{n \in N} \in X^{|N|}} \sum_{n \in N} u_n(x_n) \text{ subject to } \sum_{n \in N} x_n \leq s \tag{ARM2}$$

such that $p \cdot \sum_{n \in N} x'_n = p \cdot s$.

As in our simple markup mechanism, we use the pseudoequilibrium of a perturbed economy to identify a price vector, but the rationing mechanism solves for the pseudoequilibrium with an additional R units of *supply* rather than demand. After rounding the supply using Proposition 2 to obtain a feasible allocation ω at which all agents optimize with respect to p, there may be up to $2R$ units of excess demand. While the simple markup mechanism aimed to avoid excess demand, the excess demand and convexity assumption ensure that we are able to identify a reallocation to buyers with no wasted supply at the given prices. We do so by identifying the optimal allocation of the fixed supply vector to the buyers. This new allocation for buyers is associated with a shadow price vector p' at which these bundles are exactly demanded by buyers.

In order to ensure that this reallocation is associated with a small change in shadow prices (from p' to p) and thus correspondingly small rationing losses, we impose an additional strong monotonicity assumption on demand. This analysis uses the findings of Watt (2022) showing that, with strong monotonicity, the effect of a small change in total supply on the price vector is $O\left(1/|N_t|\right)$. Moreover, Watt shows that in IPV economies such that *expected* supply and demand are strongly monotone, this conclusion holds with high probability. The following result is proved in Appendix B.5.

Theorem 4. *Let $(\mathscr{E}_t)_{t \in \mathbb{N}}$ be a sequence of economies satisfying Assumptions 1–5.*

(a) The buyers-only rationing mechanism is well-defined: that is, a budget-balanced solution to (ARM2) *exists.*

(b) Suppose each buyer has strongly monotone demand. Then the buyers-only rationing mechanism is $O(1/N_t^a)$*-near-efficient,* $O(1/N_t^a)$*-EPIC, exactly IR for sellers and* $O(1/N_t^a)$*-IR for buyers, where* N_t^a *is the number of active buyers at the mechanism's prices.*

(c) Suppose that $(\mathcal{E}_t)_{t\in\mathbb{N}}$ *are IPV economies drawn from distributions* μ *and* χ *for which the expected demand and supply correspondences are strongly monotone. The buyers-only rationing mechanism on* $(\mathcal{E}_t)$ *is* $O_P(1/|N_t|^{1-\varepsilon})$*-near-efficient,* $O(1/|N_t|^{1-\varepsilon})$*-IIC, exactly IR for sellers and* $O_P(1/|N_t|^{1-\varepsilon})$*-IR for buyers.*

If in addition the buyers' preferences are separable across goods, the conclusions in (b) and (c) hold with exact *IR for buyers.*

6 Discussion

This paper adopts a market-design perspective to combine and extend two older traditions: one that evaluated reporting incentives in the Walrasian mechanism in convex economies and another that suggested extensions of Walrasian equilibrium theory for situations without convexity. Unlike the latter tradition, the mechanism design approach disallows physically infeasible allocations and often requires at least weak budget balance. Our markup extension has no close antecedent in either tradition.

There is little doubt that non-convexities in production or consumption are significant in many sectors, which makes the restriction of the traditional Walrasian mechanism to concave value reports and convex cost reports unsuitable for those applications. Similarly, there is little doubt that anonymous linear prices have been attractive to policymakers in practice, perhaps because they are seen as familiar, fair, and robust, so theories that require personalized or non-linear prices are sometimes deemed unsatisfactory. Despite the importance of these applications, theory has so far offered only limited guidance about the best ways to design suitable market rules for them.

We have shown that one can build upon traditional general equilibrium analysis to design mechanisms with good participation and reporting incentives that are also computable, physically feasible, weakly budget-balanced, and nearly efficient. In mechanism design analyses, it is increasingly common for researchers to describe large-market performance in terms of rates of convergence, and our analysis does that. We find that with additional restrictions on buyer preferences, the rate of convergence to efficiency that can be achieved by a rationing mechanism is much faster than for our markup mechanisms. The markup mechanism has the advantage that it is simple, always individually rational, and never requires rationing. Both classes of mechanisms have large market reporting incentives similar to the traditional Walrasian mechanism.

Our formal analysis rests in important part on our new Bound Form First Welfare Theorem, which extends the traditional statement about efficiency of exact competitive equilibrium outcomes to tight bounds on inefficiency of our approximate competitive equilibria. The theorem produces a tight bound on the inefficiency of the allocation in any price-allocation pair (p, ω), which allows the theorem to be widely applied to assess and compare different extensions of the traditional Walrasian mechanism. The applications of this theorem to evaluate and compare our two new market designs demonstrate its usefulness and how it can be applied.

References

Arrow, K. J. (1951). An extension of the basic theorems of classical welfare economics. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability.*

Aumann, R. J. (1963). On choosing a function at random. In F. Wright (Ed.), *Ergodic theory.* Academic Press.

Aumann, R. J. (1965). Integrals of set-valued functions. *Journal of Mathematical Analysis and Applications*, *12*(1), 1–12.

Aumann, R. J. (1966). Existence of competitive equilibria in markets with a continuum of traders. *Econometrica*, *34*(1), 1–17.

Ausubel, L. M., & Milgrom, P. R. (2002). Ascending auctions with package bidding. *Advances in Theoretical Economics*, *1*(1).

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

Azevedo, E. M., Weyl, E. G., & White, A. (2013). Walrasian equilibrium in large, quasilinear markets. *Theoretical Economics*, *8*(2), 281–290.

Azizan, N., Su, Y., Dvijotham, K., & Wierman, A. (2020). Optimal pricing in markets with non-convex costs. *Operations Research*, *68*(2), 480–496.

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

Bator, F. M. (1961). On convexity, efficiency, and markets. *Journal of Political Economy*, *69*(5), 480–483.

Bichler, M., Fux, V., & Goeree, J. (2018). A matter of equality: Linear pricing in combinatorial exchanges. *Information Systems Research*, *29*(4), 1024–1043.

Bichler, M., Fux, V., & Goeree, J. K. (2019). Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences*, *116*(3), 786–791.

Bichler, M., Goeree, J., Mayer, S., & Shabalin, P. (2014). Spectrum auction design: Simple auctions for complex sales. *Telecommunications Policy*, *38*(7), 613–622.

Bichler, M., Milgrom, P., & Schwarz, G. (2022). Taming the communication and computation complexity of combinatorial auctions: The FUEL bid language. *Management Science*.

Bikhchandani, S., & Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, *74*(2), 385–413.

Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Chavas, J.-P., & Briec, W. (2012). On economic efficiency under non-convexity. *Economic Theory*, *50*(3), 671–701.

Chen, X., Dai, D., Du, Y., & Teng, S.-H. (2009). Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *2009 50th annual IEEE symposium on foundations of computer science* (pp. 273–282).

Danilov, V., Koshevoy, G., & Murota, K. (2001). Discrete convexity and equilibria in economies with indivisible goods and money. *Mathematical Social Sciences*, *41*(3), 251–273.

Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, *39*(1), 195–259.

Debreu, G. (1951). The coefficient of resource utilization. *Econometrica*, *19*, 273–292.

Farrell, M. J. (1959). The convexity assumption in the theory of competitive markets. *Journal of Political Economy*, *67*(4), 377–391.

Feldman, M., Shabtai, G., & Wolfenfeld, A. (2021). Two-price equilibrium. *arXiv preprint arXiv:2112.08215*.

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

Hatfield, J. W., Kominers, S. D., Nichifor, A., Ostrovsky, M., & Westkamp, A. (2013). Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, *121*(5), 966–1005.

Heller, W. P. (1972). Transactions with set-up costs. *Journal of Economic Theory*, *4*(3), 465 - 478.

Jackson, M. O. (1992). Incentive compatibility and competitive allocations. *Economics Letters*, *40*(3), 299–302.

Knuth, D. E. (1976). Big omicron and big omega and big theta. *ACM Sigact News*, *8*(2), 18–24.

Koopmans, T. C. (1961). Convexity assumptions, allocative efficiency, and competitive equilibrium. *Journal of Political Economy*, *69*(5), 478–479.

Liberopoulos, G., & Andrianesis, P. (2016). Critical review of pricing schemes in markets with non-convex costs. *Operations Research*, *64*(1), 17–31.

Milgrom, P. (2009). Assignment messages and exchanges. *American Economic Journal: Microeconomics*, *1*(2), 95–113.

Milgrom, P., & Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, *70*(2), 583–601.

Milgrom, P., & Strulovici, B. (2009). Substitute goods, auctions, and equilibrium. *Journal of Economic Theory*, *144*(1), 212–247.

Morimoto, S., & Serizawa, S. (2015). Strategy-proofness and efficiency with non-quasi-linear preferences: A characterization of minimum price Walrasian rule. *Theoretical Economics*, *10*(2), 445–487.

Nguyen, T., & Vohra, R. (2022). (Near) substitute preferences and equilibria with indivisibilities. *Working paper*.

Roberts, D. J., & Postlewaite, A. (1976). The incentives for price-taking behavior in large exchange economies. *Econometrica*, *44*(1), 115–127.

Rockafellar, R. T., & Wets, R. J.-B. (2009). *Variational analysis* (Vol. 317). Springer Science & Business Media.

Rothenberg, J. (1960). Non-convexity, aggregation, and Pareto optimality. *Journal of Political Economy*, *68*(5), 435–468.

Rustichini, A., Satterthwaite, M. A., & Williams, S. R. (1994). Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 1041–1063.

Scarf, H. (1967). The approximation of fixed points of a continuous mapping. *SIAM Journal on Applied Mathematics*, *15*(5), 1328–1343.

Scarf, H. (1994). The allocation of resources in the presence of indivisibilities. *Journal of Economic Perspectives*, *8*(4), 111–128.

Starr, R. M. (1969). Quasi-equilibria in markets with non-convex preferences. *Econometrica*, *37*(1), 25–38.

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

Watt, M. (2022). Strong monotonicity and perturbation-proofness of exchange economies. *Working paper*.

Wilson, R. B. (1993). *Nonlinear pricing*. Oxford University Press.

A Non-convexity and approximate equilibria

We begin with a slightly stronger statement of the Shapley-Folkman Lemma that is used in general equilibrium theorem with non-convexities.

Proposition 5. *Let $S_i \subseteq \mathbb{R}^L$ for $i = 1, ..., N$, $S = \bigoplus_{i=1}^N S_i$ and $L' = \min(L, N)$. Then for any $x \in \mathrm{co}(S)$:*

(a) (Shapley-Folkman Lemma) $x = \sum_{i=1}^N x_i$ where $x_i \in \mathrm{co}(S_i)$ and $|i : x_i \in \mathrm{co}(S_i) \setminus S_i| \leq L'$.

(b) (Starr, 1969) If S_i is ordered so that $r(S_i)$ is nonincreasing in i, then there is $y \in S$ such that $|x - y| \leq \sqrt{\sum_{i=1}^{L'} r(S_i)^2}$.

(c) (Heller, 1972) If S_i is ordered so that $\rho(S_i)$ is nonincreasing in i, then there is $y \in S$ such that $|x - y| \leq \sqrt{\sum_{i=1}^{L'} \rho(S_i)^2}$.

These results have been used in the general equilibrium context to obtain **approximate equilibria**, which are price-allocation pairs (p, ω) such that $x_n \in D_n(p)$ for all n, $y_f \in S_f(p)$ for all f but $\left|\sum_{n\in N} x_n - \sum_{f\in F} y_f\right| \leq s$ for some $s > 0$. In particular, the allocation associated with an approximate equilibrium may have excess demand and therefore be infeasible. The approximate equilibrium is obtained by identifying the competitive equilibrium associated with a convexified version of the economy (in which each agent's upper contour set is replaced by its convex hull) and applying the results of Proposition 5 to the resulting allocation. The approximate equilibrium analogues of Proposition 5 are contained in Proposition 6 below.

Proposition 6. *For economy $\mathscr{E} = (N, F)$:*

(a) (Starr, 1969) There is $\omega \in \mathrm{co}(\Omega)$ and $p \in \mathbb{R}^L_+, p > 0$ such that $|n : x_n \in \mathrm{co}(X)| + |f : y_f \in \mathrm{co}(Y)| \leq L$ and for all other agents, $x_n \in D_n(p)$ and $y_f \in S_f(p)$.

(b) Let $r_n = \sup_{\bar{u}\in\mathbb{R}} r(UC_n^{\bar{u}})$ and $r_f = \sup_{\bar{\pi}\in\mathbb{R}} r(UC_f^{\bar{\pi}})$. Let $\delta \geq 0$ satisfy $r_n \leq \delta$ for all $n \in N$ and $r_f \leq \delta$ for all $f \in F$. Then there exists $p \in \mathbb{R}^L_+, p > 0$, $x_n \in X$ and $y_f \in Y$ such that $x_n \in D_n(p)$ for all $n \in N$, $y_f \in S_f(p)$ for all $f \in F$ and $\left|\sum_{n\in N} x_n - \sum_{f\in F} y_f\right| \leq \delta\sqrt{L}$.

(c) Let $\rho_n = \sup_{\bar{u}\in\mathbb{R}} \rho(UC_n^{\bar{u}})$ and $\rho_f = \sup_{\bar{\pi}\in\mathbb{R}} \rho(UC_f^{\bar{\pi}})$. Let $\delta' \geq 0$ satisfy $\rho_n \leq \delta'$ for all $n \in N$ and $\rho_f \leq \delta'$ for all $f \in F$. Then there exists $p \in \mathbb{R}^L_+, p > 0$, $x_n \in X$ and $y_f \in Y$ such that $x_n \in D_n(p)$ for all $n \in N$, $y_f \in S_f(p)$ for all $f \in F$ and $\left|\sum_{n\in N} x_n - \sum_{f\in F} y_f\right| \leq \delta' L$.

The statement of Proposition 6(a) is standard, but the statements of Proposition 6(b) and (c) are stronger than the classical statement due to Starr (1969) and Heller (1972). Again, the quasi-linearity of agent preferences allows us to conclude that agents are utility- and profit-maximizing, rather than just expenditure-minimizing.

Finally, we introduce a more general class of quasilinear preferences to which many of our results also apply and which offer a meaningful interpretation in terms of perceived complementarity and substitutability of goods. Nguyen and Vohra (2022) introduced the concept of the **generalized Δ-single improvement property**, which is a generalization of the well-known single improvement property.

Definition A.1. The preferences of buyer $n \in N$ satisfies the generalized Δ-single improvement property (or satisfy $\Delta-$substitutes) for some $\Delta > 0$ if for any price vector $p > 0$, any two bundles $x, y \in D_n(p)$ and any price change δp such that $\delta p \cdot x > \delta p \cdot y$, there exist $a \leq (x-y)^+$ and $b \leq (y-x)^+$ such that:

(a) $|a| + |b| \leq \Delta$

(b) $\delta p \cdot a > \delta p \cdot b$, and

(c) $x - a + b \in D_n(p)$.

Here $(x-y)^+$ denotes the vector whose ℓ^{th} component is $\max(x_\ell - y_\ell, 0)$.

The Δ in this definition captures a measure of the substitutability and complementarity between goods. Preferences with the single improvement property of Gul and Stacchetti (1999) are contained in the class with $\Delta = 2$.

By our assumption on the compactness of X and Y, all preferences and technologies satisfy the general Δ-improvement property for some Δ (as noted by Nguyen and Vohra (2022)). But the following stronger relationship between the inner radii of preferences and the Δ-single improvement property also holds.

Proposition 7. *Let $r_n = \sup_{\bar{u} \in \mathbb{R}} r(UC_n^{\bar{u}})$. Then the preferences of buyer $n \in N$ satisfy the generalized $\Delta-$single improvement property for all $\Delta > 2\sqrt{2} r_n$.*

Proof. Let the preferences of buyer n satisfy $r_n = \sup_{\bar{u} \in \mathbb{R}} r(UC_n^{\bar{u}})$. Let $x, y \in X$ and $p \in \mathbb{R}_+^L$ be given such that $x, y \in D_n(p)$. Suppose $|(x-y)^+| + |(y-x)^+| \geq 2r_n$ (else the preferences immediately satisfy the Δ improvement property for $\Delta = 2r_n$).

For any $\epsilon > 0$, let $z \in \mathbb{R}^L_+$ be the unique convex combination of x and y such that $|x - z| = r_n + \epsilon$ and write $z = \lambda x + (1-\lambda)y$. By construction $(z, p \cdot z) \in \text{co}(UC_n^{u_n(x)-p\cdot x})$. Then by the bound on the non-convexity of the preferences, there is a set $T \subseteq UC_n^{u(x)-p\cdot x}$ with $\text{rad}(T) \leq r_n$ such that $(z, p \cdot z) = \sum_{(x',t')\in T} \alpha_{(x',t')}(x', t')$ where $\sum_{(x',t')\in T} \alpha_{(x',t')} = 1$.

We now argue that for all $(x', t') \in T$, $x' \in D_n(p)$ and $t' = p \cdot x'$. To see this, note that $x \in D_n(p)$ implies $u_n(x') - p \cdot x' \leq u_n(x) - p \cdot x$. Summing, we have

$$\begin{aligned} u_n(x) - p \cdot x &\geq \sum_{(x',t')\in T} \alpha_{(x',t')}[u_n(x') - p \cdot x'] \\ &= \sum_{(x',t')\in T} \alpha_{(x',t')} u_n(x') - p \cdot z \\ &= \sum_{(x',t')\in T} \alpha_{(x',t')}[u_n(x') - t'] \end{aligned}$$

On the other hand, since $(x', t') \in UC_n^{u(x)-p\cdot x}$ we have $u_n(x') - t' \geq u(x) - p \cdot x$. The only way these can simultaneously hold is if $u_n(x') - t' = u(x) - p \cdot x$ for all $(x', t') \in T$.

However, we then have $\sum_{(x',t')\in T} \alpha_{(x',t')}[u_n(x') - p \cdot x'] = u(x) - p \cdot x$. This implies that at least one of $u_n(x') - p \cdot x' \geq u_n(x) - p \cdot x$. But then $x \in D_n(p)$ implies that $u_n(x') - p \cdot x' = u_n(x) - p \cdot x$ for all x', so $x' \in D_n(p)$.

By construction, $|x - x'| \leq 2r_n + \epsilon$. But then $||x - x'||_1 \leq 2\sqrt{2}r_n + \epsilon$ as well. □

Clearly, the generalized Δ-single property can be readily extended to sellers, by replacing the expressions for utility with those for profits, and an analogue of Proposition 7 also holds.

Nguyen and Vohra (2022) demonstrate the following approximate equilibrium result in a setting with indivisibilities (so that $X \subseteq \mathbb{Z}^L_+$ and $Y \subseteq \mathbb{Z}^L_+$).

Proposition 8. *Suppose all buyers' preferences and sellers' technologies satisfy the generalized Δ-improvement property and that $X \subseteq \mathbb{Z}^L_+$ and $Y \subseteq \mathbb{Z}^L_+$. Then there exists $p \in \mathbb{R}^L_+$, $p > 0$, $x_n \in X$ and $y_f \in Y$ such that $x_n \in D_n(p)$ for all $n \in N$, $y_f \in S_n(p)$ and for each $\ell \in L$, $\left|\sum_{n\in N} x_{nl} - \sum_{f\in F} y_{fl}\right| \leq \Delta - 1$.*

Note that the concept of approximate equilibrium in this result is somewhat stronger than the previous results since the maximum imbalance in supply and demand is bounded good-by-good, rather than in terms of Euclidean distance in commodity space. However, depending on the relative size of Δ, the inner radii of non-convexity and the breadths of non-convexity of preferences, any of the approximate equilibrium bounds in Proposition 6(b), 6(c) or 8 may be strongest for our purposes.

B Proofs omitted from the main text

B.1 Proof of Proposition 3

Proof. We offer two proofs of this claim: the first directly from the definitions and the second highlights the relationship with the envelope theorem.

For the first proof, let $x' \in D_n(p')$. We have that

$$\mathscr{R}_n(p', x) = u_n(x') - p' \cdot x' - (u_n(x) - p' \cdot x).$$

But since $x \in D_n(p)$, we have $u(x') - p \cdot x' \leq u(x) - p \cdot x$, so that

$$\mathscr{R}_n(p', x) \leq p \cdot x' - p \cdot x + p' \cdot x - p' \cdot x = (p - p') \cdot (x' - x),$$

which is $O(\|p - p'\|)$ since $x', x \in X$, a compact set.

For the second proof, write

$$\begin{aligned}
\mathscr{R}_n(p', x) &= \hat{u}_n(p') - (u_n(x) - p' \cdot x) \\
&= \hat{u}_n(p') - (u_n(x) - p \cdot x) - p \cdot x + p' \cdot x \\
&= \hat{u}_n(p') - \hat{u}_n(p) + (p' - p) \cdot x.
\end{aligned}$$

Now let $p(t) = (1 - t)p + tp'$ for $t \in [0, 1]$ and apply the Milgrom and Segal (2002) envelope theorem to the parametrized utility maximization problem

$$\hat{u}_n(p(t)) = \max_{x \in X} u_n(x) - p(t) \cdot x,$$

to give

$$\hat{u}_n(p') = \hat{u}_n(p) - \int_0^1 (p' - p) \cdot d(t) dt.$$

for selections $d(t) \in D_n(p(t))$. Substituting into the expression for $\mathscr{R}_n(p', x)$ above, we obtain

$$\mathscr{R}_n(p', x) = -\int_0^1 (p' - p) \cdot d(t) dt + (p' - p) \cdot x = \int_0^1 (p' - p) \cdot (x - d(t)) dt$$

which is bounded above by $(p' - p) \cdot (x - x')$, the same expression as before, since $(p' - p) \cdot (x - d(t))$ is increasing in t by the law of demand. $\square$

B.2 Proof of Theorem 2

Part (a) Fix some $\mathscr{E}$ and consider any sequence $\alpha_i \to \alpha$ and selections R_i of revenues $\sum_{n\in N}(1+\alpha_i)p^{\alpha_i}\cdot x_n^{\alpha_i} - \sum_{f\in F} p^{\alpha_i}\cdot y_f^{\alpha_i}$ associated with some markup mechanisms $(\alpha_i, p^{\alpha_i}, \omega^{\alpha_i})$ constructed as in Definition 4.2. We will show that $\lim_i R_i$ is obtainable as the revenue of some markup mechanism $(\alpha, p^\alpha, \omega^\alpha)$ so that the infimum in equation (A) is attained (and thus the minimum exists).

By the saddle point condition associated with the objective in Definition 4.2, we have that there are some $\tilde{\omega}^{\alpha_i}$ maximizing over $\mathrm{co}(\Omega)$ the objective $\sum_{n\in N}\frac{\mathrm{cav}(u_n)(x_n)}{1+\alpha_i} - \sum_{f\in F}\mathrm{vex}(c_f)(y_f)$. As $\alpha_i \to \alpha$, this objective hypo-converges (since it is continuous and bounded) to $\sum_{n\in N}\frac{\mathrm{cav}(u_n)(x_n)}{1+\alpha} - \sum_{f\in F}\mathrm{vex}(c_f)(y_f)$, so that $\tilde{\omega}^{\alpha_i}\to\tilde{\omega}^\alpha$ for some $\tilde{\omega}^\alpha$ that maximizes this latter objective. By optimality, each p^{α_i} lies in the superdifferential ∂^* of the concavified valuation functions of each buyer and the subdifferential of the convexified cost functions of each seller at $\tilde{\omega}^{\alpha_i}$. For these concave / convex functions, the super- and subdifferential correspondences are upper hemicontinuous, so that the sequence of p^{α_i} must converge to some p^α in the super- and subdifferentials at $\tilde{\omega}^\alpha$. Finally, since the demand and supply correspondences are upper hemicontinuous, the convergence of prices implies that ω^{α_i} must approach some ω^α such that $x_n^\alpha \in D_n((1+\alpha)p^\alpha)$ and $y_f^\alpha \in S_f(p^\alpha)$. Thus the limit of R_i is attained as the revenue of some markup mechanism $(\alpha, p^\alpha, \omega^\alpha)$ as required.

Part (b) For notational simplicity, we drop the index for t in the prices, premiums and allocations.

The construction in Definition 4.2 ensures, via Proposition 2, that $\sum_{f\in F} y_f^\alpha - \sum_{n\in N} x_n^\alpha \le (2R)1_L$. Thus, it suffices to show that for sufficiently large $|N_t|$, there is an α such that

$$\alpha\sum_{n\in N} p^\alpha\cdot x_n^\alpha \ge (2R)p^\alpha\cdot 1_L.$$

Moreover, if this $\alpha = O\left(\frac{1}{|N_t|}\right)$, then since $\alpha^* < \alpha$, (b) will follow. To arrive at this result, we will show that for fixed $\alpha > 0$ close enough to zero, $\sum_{n\in N} p^\alpha\cdot x_n^\alpha$ is $\Omega(|N_t|)$, while $p^\alpha\cdot 1_L$ is $O(1)$.

Let $\mathcal{S}^\alpha$ be the value of the saddle point problem

$$\min_{p\in\mathbb{R}_+^L}\max_{\omega\in\Omega}\sum_{n\in N}\frac{\mathrm{cav}(u_n)(x_n)}{1+\alpha} - \sum_{f\in F}\mathrm{vex}(c_f)(y_f) - p\cdot\left(\sum_{n\in N} x_n + R1_L - \sum_{f\in F} y_F\right).$$

First, note that $\mathcal{S}^\alpha$ is $\Theta(|N_t|)$ for sufficiently small α by Assumption 3. To see this, denote by $f(|N_t|) = \sum_{n\in N_t} u_n(x_n^*)$ and $g(|N_t|) = \sum_{f\in F_t} c_f(y_f^*)$. Then by Assumption 3 and the definition of $\Theta(\cdot)$, $\liminf_{N\to\infty}\frac{f(N)}{N} = u > 0$, say, and $\limsup_{N\to\infty}\frac{g(N)}{N} = c > 0$, with $u - c > 0$. Then

$\mathcal{S}^\alpha \geq \liminf_{N\to\infty} \frac{f(N)}{(1+\alpha)N} - g(N) = \frac{u}{1+\alpha} - c$, which is positive for sufficiently small α.

Now we show that this implies $\sum_{n\in N} p^\alpha \cdot x_n^\alpha$ is $\Omega(|N_t|)$ for small, fixed α. To see this, note that since $\sum_{f\in F_t} c_f(y_f^\alpha)$ is $\Omega(|N_t|)$, individual rationality of the sellers (in the perturbed economy) implies that $\sum_{f\in F_t} p^\alpha \cdot y_f^\alpha$ is $\Omega(|N_t|)$. But then by complementary slackness $\sum_{n\in N_t} p^\alpha \cdot x_n^\alpha = \sum_{f\in F_t} p^\alpha \cdot y_f^\alpha - Rp^\alpha \cdot 1_L$, and since Assumption 4 implies $\|p\| \leq M$, we must have that $\sum_{n\in N} p^\alpha \cdot x_n^\alpha$ is $\Omega(|N_t|)$.

Since for α near zero, $\sum_{n\in N} p^\alpha \cdot x_n^\alpha$ is $\Omega(|N_t|)$ and $(2R)p^\alpha \cdot 1_L$ is $O(1)$ (where R is $O(1)$ by assumption 2), for sufficiently large $|N_t|$, there is some α (and thus some least α by (a)) such that

$$\alpha \sum_{n\in N} p^\alpha \cdot x_n^\alpha \geq (2R)p^\alpha \cdot 1_L,$$

and furthermore, this α is $O\left(\frac{1}{|N_t|}\right)$. Since $\alpha^* < \alpha$, we have that α^* is $O\left(\frac{1}{|N_t|}\right)$, as required.

Part (c) We now apply the First Welfare Bound to show that the allocation ω^{α^*} is approximately efficient. In order to satisfy the assumptions on prices in Theorem 1, we imagine ω^{α^*} was implemented with a single price vector p^{α^*} and (therefore) a budget deficit. Theorem 1 tells us that

$$\mathcal{S}(\omega^*) - \mathcal{S}(\omega^{\alpha^*}) \leq \mathcal{R}(p^{\alpha^*}, \omega^{\alpha^*}) + p^{\alpha^*} \cdot \left(\sum_{f\in F_t} y_f^{\alpha^*} - \sum_{n\in N_t} x_n^{\alpha^*} \right). \tag{1}$$

By construction, in ω^{α^*} at prices p^{α^*}, no sellers are rationed, while at prices $(1+\alpha^*)p^{\alpha^*}$, no buyers are rationed. But (1) requires the rationing of buyers at price p^{α^*}. For this, we use Proposition 3: since α^* is $O\left(\frac{1}{|N_t|}\right)$ and $\|p^{\alpha^*}\|$ is bounded (by assumption 4), this implies that $\mathcal{R}_n(p^{\alpha^*}, x_n^{\alpha^*})$ is $O\left(\frac{1}{|N_t|}\right)$, and so $\mathcal{R}(p^{\alpha^*}, \omega^{\alpha^*}) = \sum_{n\in N_t} \mathcal{R}_n(p^{\alpha^*}, x_n^{\alpha^*})$ is $O(1)$.

Finally, we note that the budget deficit (the second term on the right of Equation (1)) is $O(1)$ since the excess supply is bounded by construction and each component of p^{α^*} is $O(1)$. Thus Theorem 2 follows.

B.3 Proof of Theorem 3

Part (a) This follows by simply noting that all markup mechanisms are envy-free, and so Theorem 1 of Azevedo and Budish (2019) implies the result.

Parts (b) and (c) First we formally define the expected supply and demand correspondences-Given distribution ν on $\mathscr{V}$, The expected indirect utility function is defined pointwise for $p \in \mathscr{P}$ by

$$\mathbb{E}_\nu[\hat{u}(p)] = \int_{\mathscr{V}} \hat{u}_n(p) d\nu(u_n),$$

and similarly the expected indirect profit function is

$$\mathbb{E}_\chi[\hat{\pi}(p)] = \int_{\mathscr{C}} \hat{\pi}_f(p) d\chi(\pi_f).$$

The expected demand correspondence is then $\mathbb{E}_\nu[D(p)] = -\partial \mathbb{E}_\nu[u(p)]$ and the expected supply correspondence is $\mathbb{E}_\chi[S(p)] = \partial \mathbb{E}_\chi[\pi(p)]$.[25]

Parts (b) and (c) follow directly from the corresponding theorems for the Walrasian mechanism—namely, Theorems 1, 3 and 4 adapted as in Appendix C of Watt (2022). To see this, we show that the objective for the markup mechanism—both under truthful reporting and after misreporting by a single agent—differs from the objective of the Walrasian mechanism for the convexified economy by a $O(1)-$Lipschitz convex function, which constitutes a perturbation under the definition in Watt (2022). Suppose that under truthful reporting, the mechanism chooses some markup $\alpha \geq 0$ and that under a misreport by an agent, the mechanism chooses markup $\alpha' \geq 0$. The markup mechanism (seller) prices p^α under truthful reporting minimize the dual objective

$$\frac{1}{1+\alpha} \sum_{n \in N} \hat{u}_n(p) + \sum_{f \in F} \hat{\pi}_f(p) - p \cdot R1_L,$$

while under the misreport, the (seller) price vector $p^{\alpha'}$ minimizes

$$\frac{1}{1+\alpha'} \sum_{n \in N} \hat{u}_n(p) + \sum_{f \in F} \hat{\pi}_f(p) - p \cdot R1_L.$$

Note that these objectives do not have cav or vex in them since the indirect utility and profit functions are the same for the original valuations and costs as their concavified/convexified versions.

The Walrasian mechanism for the convexified economy minimizes

$$\sum_{n \in N} \hat{u}_n(p) + \sum_{f \in F} \hat{\pi}_f(p).$$

[25]The expected demand and supply correspondences can also be defined using the set-valued integral of Aumann (1965), but we refer the reader to Watt (2022) for details on this construction.

This objective and the $\alpha-$objective differ by

$$\frac{\alpha}{1+\alpha}\sum_{n\in N}\hat{u}_n(p)-p\cdot R1_L.$$

Since $\alpha\leq O(1/N)$, we have $\frac{\alpha}{1+\alpha}\leq O(1/N)$ as well, while $\sum_{n\in N}\hat{u}_n(p)$ is $O(N)-$Lipschitz since its subdifferential is total demand at p which is $O(N)$ (and the Lipschitz constant is the largest selection from the subdifferential). This implies that the perturbation above is $O(1)$-Lipschitz. A similar analysis applies for the α'-objective.

B.4 Proof of Proposition 4

Proof. The minimum-loss rationing mechanism solves the non-convex problem

$$\min_{p,\omega:p\cdot\sum_n x_n=p\cdot\sum_f y_f}\left\{\sum_{n\in N}\hat{u}_n(p)-u_n(x_n)+p\cdot x_n+\sum_{f\in F}\hat{\pi}_f(p)-p\cdot y_f+c_f(y_f)\right\}.$$

Conditional on the minimizing price p, the allocation ω must minimize the same objective. But with respect to ω, the $\hat{u}_n(p)$ and $\hat{\pi}_f(p)$ terms are constant, while $\sum_{n\in N}p\cdot x_n-\sum_{f\in F}p\cdot y_f=0$ by assumption. Thus

$$\omega\in\operatorname*{arg\,min}_{\omega:p\cdot\sum_n x_n=p\cdot\sum_f y_f}\sum_{n\in N}-u_n(x_n)+\sum_{f\in F}c_f(y_f),$$

which is equivalent to the efficient allocation problem, except for the constraint that $p\cdot\sum_n x_n=p\cdot\sum_f y_f$, which is always binding at the efficient allocation problem by the complementary slackness condition of the related Lagrangian problem. □

B.5 Proof of Theorem 4

Part (a) Budget balance holds if for each good ℓ, either $\sum x_n'^{\ell}=s^\ell$ or $p^\ell=0$. Let us suppose that $p^\ell\neq 0$, then by the complementary slackness condition of (ARM1), we must have that at price p^ℓ, the total demand for good ℓ equals $\sum_f\tilde{y}^\ell+R_{\mathscr{G}}$, which is greater than or equal to s^ℓ by construction. Now suppose that $\sum_n x_n'^{\ell}<s^\ell$. Then we could increase $x_n'^{\ell}$ for each agent up to a selection of their total demand for good ℓ at p until the total consumption equals s^ℓ, since the feasible consumption region is box-shaped (and thus a lattice). This cannot decrease the objective $\sum_n u_n(x_n)$ of (ARM2) since u_n is monotone. Thus we obtain a new x' solving (ARM2) such that $\sum_n x_n'^{\ell}=s^\ell$, as required.

Part (b) From the saddle point of (ARM1) in Definition 5.2, define $\tilde{s} = \sum \tilde{y}_f + R_{\mathscr{G}}$. Since each buyers' demand is strongly monotone, we have via Theorem 1 in Watt (2022) that the price p (which satisfies $\tilde{s} \in \sum_{n \in N} D_n(p)$) and any shadow price p' (which satisfies $s \in \sum_{n \in N} D_n(p')$) must satisfy $\|p - p'\| \leq O(1/|N_t^a|)$, where N_t^a is the number of active buyers (since $s - \tilde{s} \sim O(1)$).

Then by the expression for $\mathscr{R}_n(p, x_n')$ obtained in the proof of Proposition 3, we have for all $d_n \in D_n(p)$ and $d_n' \in D_n(p')$

$$\mathscr{R}_n(p, x_n') \leq (p' - p) \cdot (d_n - d_n').$$

Thus, by summing, we have

$$\mathscr{R}(p, x_n') \leq (p' - p) \cdot R1_L \leq O\left(1/|N_t^a|\right).$$

Employing Theorem 1 with the prices and allocation from the single-price mechanism with rationing (noting that it is budget-balanced and no sellers are rationed) thus implies that the deadweight loss is $O\left(1/|N_t^a|\right)$.

The proof of incentive-compatibility is almost identical to the proof of Theorem 3(c), with the added complication that a buyer may have an incentive to misreport in order to avoid rationing. However, since each buyer's rationing is $O\left(1/|N_t^a|\right)$, the maximum expected benefit of misreporting is $O\left(1/|N_t^a|\right)$ as well. To see this, consider any misreport by a single agent and the resulting price, say $p^\dagger$, produced in the mechanism. Then $\hat{u}_n(p^\dagger)$ is an upper bound on the agent's utility under the misreport and $\hat{u}_n(p) - O(1/|N_t^a|)$ is a lower bound on the agent's utility under truthful reporting. But by strong monotonicity, $\|p^\dagger - p\| \leq O(1/|N_t^a|)$ and so by the Lipschitz property of $\hat{u}_n$ in this environment,[26] we obtain that the difference in utilities is $O(1/|N_t^a|)$ as well.

The approximate individual rationality follows from the observation that $x_n' \in \max_x u_n(x) - p' \cdot x$ so that $u_n(x_n') - p' \cdot x_n \geq 0$. Then $u_n(x_n') - p \cdot x_n' = u_n(x_n') - p' \cdot x_n' + (p' - p) \cdot x_n' \geq (p' - p) \cdot x_n'$, which may be less than zero. However, since $\|p' - p\|$ is $O\left(1/|N_t^a|\right)$ and $\|x\| \leq O\left(1\right)$, at worst we have that the resulting allocation is $O\left(1/|N_t^a|\right)$ $-$individually rational.

Finally, note that if the agent's preferences are separable in the goods, then since p' is the clearing price for a supply vector with strictly fewer of each good, we have that $p' - p$ is nonnegative in this component. This then implies that $u_n(x_n') - p \cdot x_n' \geq (p' - p) \cdot x_n' \geq 0$, as required.

[26] The Lipschitz property of $\hat{u}_n$ follows since the Lipschitz constant of a proper, convex function is the magnitude of the largest selection from a subderivative of that function (see Theorem 9.13 in Rockafellar and Wets (2009)). For $\hat{u}_n$, this is a demand bundle, of bounded magnitude by the assumption that X is compact.

Part (c) The arguments now follows almost exactly as in Part (b), but instead exploits Theorem 2 in Watt (2022), which gives that $\|p - p'\|$ is $O_P(1/|N_t|^{1-\varepsilon})$.