

The Edgeworth Conjecture with Small Coalitions and Approximate Equilibria in Large Economies

Federico Echenique¹ and Siddarth Barman²

February 12th, 2020. (First version: February, 2019.)

¹Division of the Humanities and Social Sciences, California Institute of Technology

²Indian Institute of Science.

Abstract

We revisit the connection between bargaining and equilibrium in exchange economies, and study its algorithmic implications. We consider bargaining outcomes to be allocations that cannot be blocked (i.e., profitably re-traded) by coalitions of small size, and show that these allocations must be approximate Walrasian equilibria. Our results imply that deciding whether an allocation is approximately Walrasian can be done in polynomial time, even in economies for which finding an equilibrium is known to be computationally hard.

Contents

1	Introduction	2
1.1	Informal statement of our results	5
2	Notation and Preliminaries	7
3	A Quantitative Core Convergence Theorem	11
3.1	Bounded Hull	13
3.2	Lemma for κ -Core Allocations	16
3.3	Proof of Theorem 3	18
4	Testing Algorithm for Walrasian Allocations	19
4.1	Testing Algorithm for Economies with Piecewise-Linear Concave Utilities	21
A	Supplementary Results for Section 3	29
B	Proof of Lemma 1	29
C	Proof of Lemma 6	30
D	Proof of Theorem 7	31
E	Illustrative Example of Strongly Concave Utilities	34
F	Other Notions of Approximate Equilibria	35

1 Introduction

We present a *quantitative core convergence theorem*, with economic and algorithmic implications. Economists have, since Francis Edgeworth in 1881, been interested in the convergence of bargaining in large finite economies to market equilibrium. The celebrated *Edgeworth conjecture* states that bargaining indeterminacy and monopoly power disappear in large economies. Edgeworth focused on the contract curve, what we think of now as the core. The core, however, requires agents to join coalitions of arbitrary size. In our quantitative core convergence theorem we show that coalitions of a fixed size suffice. The fixed size depends polynomially on the approximation error, the number of goods in the economy, and consumers' heterogeneity.

The economic analysis of markets is based on the notion of Walrasian (competitive) market equilibrium and price-taking behavior. Walrasian equilibrium requires that goods are assigned prices, each agent maximizes her utility subject to what she can afford, and the market clears. Walrasian equilibria provide a conceptually crisp framework for economic analysis: with prices as guides, agents direct themselves, in a decentralized manner, to a Pareto optimal allocation. That agents are price takers means that they treat prices as fixed; they will not renegotiate the terms of trade with other agents.

The Edgeworth conjecture, or core convergence, is the economist's basic justification for the price-taking assumption. Specifically, it postulates that, in a large economy, bargaining collapses to Walrasian equilibrium, i.e., in a large economy, no group of agents would choose to upset an equilibrium by renegotiating among themselves.

In a small economy, agents bargain over the terms of trade. They may not trade at fixed prices. Each individual agent in a small economy is to some degree unique, and can therefore command a certain degree of monopoly power. Imagine the market for professional soccer players, or high-profile academics, in which each individual agent is special. Then terms of trade are determined by haggling and bargaining, not as fixed prices in competitive markets. By contrast, in a large economy with limited heterogeneity, no agent is special. Imagine the market for a commodity, like corn or steel, or for a standard piece of technology, like a personal desktop. In such economies, bargaining outcomes (formally, outcomes in the core) should collapse to the competitive outcomes that correspond to a Walrasian equilibrium. This claim is the Edgeworth conjecture.

Bargaining is a complex process, and difficult—arguably impossible—to model in detail. So economists use the notion of the core to impose discipline on bargaining outcomes. The core of an economy is the set of allocations (redistribution of goods

among agents) that no subset, or coalition, of agents can improve upon via re-trading among themselves. In other words, the core comprises the outcomes that are not “blocked” by any coalition of agents. Any allocation that is not in the core would be renegotiated, or re-contracted, by some set of agents. Note that the definition of the core involves all possible coalitions of agents, of all possible sizes.

The core convergence theorem states that in a large (in an asymptotic sense) economy, any outcome in the core, arrived at through seemingly indeterminate reallocations and bargaining, will, in fact, be a Walrasian equilibrium. In other words, the core convergence theorem captures the meaning of the Edgeworth conjecture, and supports the use of Walrasian equilibria and price-taking behavior in large economies.

Edgeworth’s conjecture was first formalized by Debreu and Scarf (Debreu and Scarf (1963) as well as Aumann (1964)).¹ In our paper we shall follow Debreu and Scarf; Aumann’s model assumes a limit economy with a continuum of agents, hence there is no scope in his model to address the questions we are interested in.² Debreu and Scarf consider a *large economy with limited heterogeneity*. Their idea, which goes back to Edgeworth himself, is to postulate a finite number of agent types, say h , and then imagine a sequence of replica economies. The n -th replica of the economy has n identical copies of each type of agent $t \in [h]$. Under standard assumptions, they show that, in the limit as $n \rightarrow \infty$, the core of the n th replica economy approaches the set of Walrasian equilibrium allocations. Note that the core convergence theorem not only requires n to be asymptotically large, but, to address the core, one has to account for coalitions that are arbitrarily large, both in size and in number.

In contrast to the core, we focus on the κ -core: the set of allocations that cannot be blocked by any coalition of size at most κ . In a large economy, it is unlikely that coalitions of arbitrary size can function effectively. Even if a large coalition can effectively communicate among its members, they would face significant hurdles in aggregating their preferences to achieve an outcome that is collectively better for the members of the coalition. Think of Arrow’s theorem (see e.g. Mas-Colell, Whinston, Green, et al. (1995)): preference aggregation is known to be very problematic, and the size and number of all coalitions in a large economy makes the problem even worse. The κ -core demands much less of the blocking coalitions, and seems to us a more plausible solution concept than the core. In addition, the number of coalitions to check for is much smaller in the κ -core than in the core ($\binom{n}{\kappa}$ with fixed κ , rather than 2^n).

¹See Hildenbrand (1974) and Anderson (1992) for reviews of the literature on core convergence.

²A “fractional” version in atomless economies is still possible, see Schmeidler (1972); Grodal (1972); Vind (1972).

Our main result is that under benign assumptions on agents' preferences, an allocation that is in the κ -core must be an approximate equilibrium (an ε -Walrasian equilibrium). κ depends polynomially on the approximation guarantee and on the agent heterogeneity in the economy. Our result is non-asymptotic; it holds for finite economies, as long as the number of agents, n , exceeds κ . We adopt the model and the assumptions of Debreu and Scarf. In fact, to make our paper self-contained, their result is restated in Section 3.

We are not the first to prove a quantitative version of Edgeworth's conjecture. The closest paper to ours is Mas-Colell Mas-Colell (1979), who also exhibits a bound on the size of blocking coalitions needed for core convergence.³ The main difference between Mas-Colell's result and ours is that he considers an average error as his approximation guarantee. He ensures that the average budget gap, where the average is taken over all agents in the economy, is small. If we translate his results into our stronger objective of a *per-agent* approximation guarantee, his result requires κ to be a function of n .⁴ Our result will instead say that κ is $\mathcal{O}(1/\varepsilon^2)$ (as long as $n \geq \kappa$). We should emphasize that Mas-Colell's theorem requires substantially weaker assumptions than ours and does not rely on the Edgeworth-Debreu-Scarf replica setting.

Our main theorem has important *algorithmic* consequences for the problem of deciding whether an allocation can be supported as an approximate Walrasian equilibrium. Imagine a policy maker who would like to install a given allocation using decentralized markets. The standard advice to such a policy maker comes from the second welfare theorem, which says that any Pareto optimal allocation can be obtained (given certain assumptions on the economy) as a Walrasian equilibrium, using a system of taxes and subsidies to adjust all agents' incomes. We study the same question as in the second welfare theorem, but when the policy maker is unable to freely use taxes or subsidies.

Specifically, given an allocation in an exchange economy, we consider the problem of deciding whether there exists prices that make the allocation an approximate Walrasian equilibrium. Our main algorithmic result is to provide a polynomial-time algorithm that resolves the question. Our finding stands in stark contrast

³Mas-Colell's result builds on techniques due to Anderson Anderson (1978). Anderson's theorem, however, does not bound the size of the blocking coalitions, and hence does not speak to the κ -core. His result connects the size (and other parameters) of the economy with the approximation guarantee.

⁴Per-agent approximation guarantees are the focus of modern results in algorithmic game theory. In the terminology of Anderson's survey Anderson (1992) it reflects uniform convergence (what he terms U3), and it is rare in the older literature on core convergence, even for results that rely on the power of all coalitions.

to the existing hardness results for Walrasian equilibria, see, for example, Chen, Dai, Du, and Teng (2009); Codenotti, Saberi, Varadarajan, and Ye (2006); Vazirani and Yannakakis (2010); Garg, Mehta, Vazirani, and Yazdanbod (2017). There are settings for which it is known that finding an equilibrium is hard, and for which our algorithm efficiently decides whether a given allocation can be supported as an approximate Walrasian equilibrium.

We are, to the best of our knowledge, the first to study the question of “testing” a putative equilibrium allocation. That is, without use of taxes and subsidies, and thus outside the scope of the second welfare theorem. Our paper resolves this question with a satisfactory algorithm.

1.1 Informal statement of our results

Our paper makes an economic and an algorithmic contribution. The economic contribution is in the form of a non-asymptotic core convergence theorem, for a relaxed notion of the core: one that only requires robustness to blocks by small coalitions. The algorithmic contribution is to the problem of testing whether an allocation is a (approximate) Walrasian equilibrium allocation. The main lemma in our core convergence result allows us to develop an efficient algorithm that decides whether an allocation can be supported as an approximate Walrasian equilibrium (and find the supporting/equilibrium prices, if they exist).

The κ -core is the set of allocations that are not blocked by coalitions of size at most κ . We show that core convergence is obtained as long as κ is polynomially large (in the approximation parameter and agent heterogeneity) and the size of the economy is at least κ . Thus, core convergence is obtained in finite economies.

Specifically, we prove that in any replica economy with well-behaved utilities, allocations in the κ -core are, in fact, ε -Walrasian, as long as κ is $\mathcal{O}\left(\frac{h^2\ell}{\varepsilon^2}\right)$; here, h denotes the heterogeneity of the economy (the number of different types of agents), ℓ denotes the number of goods, and $\varepsilon > 0$ is the approximation parameter; see Theorem 3. Note that our result does not require the number of agents n , or the size of blocking coalitions, to be arbitrarily large. The result is applicable whenever n is $\Omega(\kappa)$.

Our work develops new techniques for addressing core convergence. Instead of relying on limiting, or measure-theoretic, arguments as in the classical literature (see Hildenbrand (1974)), our proof builds upon geometric insights and dimension-free results. In that sense, we show how research in algorithmic game theory can be used to address a foundational issue in economics, such as the price-taking assumption. Specifically, we employ the approximate version of Carathéodory’s theorem

(see Barman (2015) and references therein). The classical theorem of Debreu and Scarf requires a “rounding to rational numbers” argument—the size of the economy needs to be arbitrarily large so that such a rounding is possible. Our proof avoids such a rounding by invoking the approximate Carathéodory theorem. In this manner, we ascertain that the largest coalitions needed to ensure an approximate Walrasian equilibrium are only polynomially large and independent of the size of the economy.

Our key technical insight is provided through an efficiently implementable characterization of approximate Walrasian equilibrium allocations (Lemma 4 in Section 3.1). Our characterization yields, not only the quantitative version of core convergence, but also an efficient algorithm that tests whether a given allocation can be supported as an approximate equilibrium.

Our algorithm takes as input an allocation, and determines whether there are prices that make the allocation into an approximate Walrasian equilibrium. This result is potentially useful to understand the policy objectives that can be obtained in a decentralized fashion, as a Walrasian equilibrium. It also contributes to the algorithmic literature on equilibrium computation, as it identifies a notable dichotomy between testing and computing Walrasian equilibria.

The problem of computing Walrasian equilibria has been studied extensively in algorithmic game theory. Most results in this direction are negative: finding a Walrasian equilibrium is computationally hard for general settings; see, e.g., Chen, Dai, Du, and Teng (2009); Codenotti, Saberi, Varadarajan, and Ye (2006); Vazirani and Yannakakis (2010); Garg, Mehta, Vazirani, and Yazdanbod (2017). By contrast, we show that the testing counterpart is computationally tractable. In particular, we develop a polynomial-time algorithm that, for a given allocation, finds prices (if they exist) that make the allocation an approximate Walrasian equilibrium (Theorem 7). Note that, in particular settings, it might be possible to efficiently test if a given allocation is Pareto optimal and, hence (using the characterization provided by the welfare theorems), determine whether the allocation can be supported as an *exact* equilibrium. However, in and of itself, this observation does not extend to approximate equilibria. Furthermore, it does not directly lead to an efficient method for finding the supporting equilibrium prices.

The above-mentioned results require the utilities in the economy to be strongly concave. Hence, as is, they do not address piecewise-linear concave (PLC) utilities, which, while concave, are not strongly concave. The PLC case is particularly important in light of the hardness results in ?Chen, Dai, Du, and Teng (2009); Garg, Mehta, Vazirani, and Yazdanbod (2017). However, we show that the developed ideas can be adapted to obtain an efficient testing algorithm even for PLC economies (Theorem 10).

2 Notation and Preliminaries

Notational Conventions. For vectors $x, y \in \mathbb{R}^\ell$, write $x \leq y$ iff $x_j \leq y_j$ for all $j \in [\ell]$. We will use Δ to refer to the standard simplex in \mathbb{R}^ℓ ,

$$\Delta := \left\{ x \in \mathbb{R}^\ell \mid \sum_{j=1}^{\ell} x_j = 1 \text{ and } x_j \geq 0 \text{ for all } j \right\}.$$

The convex hull of a set of vectors $A \subset \mathbb{R}^\ell$ will be denoted by $\text{cvh}(A)$. In addition, the all-ones and all-zeros vector will be denoted by $\mathbf{1}$ and $\mathbf{0}$, respectively.

Exchange economies. An exchange economy \mathcal{E} comprises of a set of consumers, $[h] := \{1, 2, \dots, h\}$, and a set of goods, $[\ell] := \{1, 2, \dots, \ell\}$. Each consumer is endowed with different quantities of the goods as an endowment; specifically, for every consumer $i \in [h]$, the vector $\omega_i \in \mathbb{R}_+^\ell$ denotes (componentwise) the amount of each good endowed to i . The preference of consumer $i \in [h]$, over bundles of goods, is specified by a utility function, $u_i : \mathbb{R}_+^\ell \mapsto \mathbb{R}$. In particular, every consumer is described by a pair (u_i, ω_i) . An exchange economy \mathcal{E} is a tuple $((u_i, \omega_i))_{i=1}^h$.

We shall adopt some standard assumptions on consumers' utilities: u_i s are continuous and monotone increasing. We further assume that the utilities are continuously differentiable⁵ and α -strongly concave, with $\alpha > 0$. Strong concavity (convexity) is a well-studied property in convex optimization (see, for example, Boyd and Vandenberghe (2004)). It provides a parametric strengthening of concavity. Formally, a differentiable function, $u : \mathbb{R}^\ell \mapsto \mathbb{R}$, is said to be α -strongly concave within a set $\mathcal{R} \subset \mathbb{R}^\ell$ iff the following inequality holds for all $x, y \in \mathcal{R}$:

$$u(y) \leq u(x) + \nabla u(x)^T (y - x) - \frac{\alpha}{2} \|y - x\|^2.$$

Here, $\nabla u(x)$ is the gradient of the function u at point x and $\|\cdot\|$ denotes the Euclidean norm. Note that if $\alpha = 0$, then u is simply a concave function. Furthermore, the case of $\alpha > 0$ corresponds to strict concavity.

The set \mathcal{R} specifies the subdomain over which strong concavity holds. We assume that \mathcal{R} is appropriately large and, in particular, that it contains the Euclidean ball of radius $r \in \mathbb{R}_+$ and center ω_i (the endowment vector), for each i . We assume that the radius r satisfies $\alpha r^2 \geq \frac{2\varepsilon\lambda\ell}{h} + 2$, with $\varepsilon > 0$ being the approximation parameter and λ denoting the Lipschitz constant of the utilities. At a high level, the condition

⁵This additional smoothness assumption is primarily for ease of exposition. One can extend the arguments developed in this work to continuous functions by replacing gradients with subgradients.

asserts that agent i 's utility function has sufficient curvature close to i 's endowment ω_i . The strong concavity assumption is satisfied by standard examples of utility functions used in economics; in particular, utilities of the form $u(y) := \left(\sum_{j=1}^{\ell} y_j^{\rho}\right)^{\frac{1}{\rho}}$, with $0 < \rho < 1$ —i.e., utilities within the CES (constant elasticity of substitution) family of utilities—are strongly concave; see Appendix E for an illustrative example. Throughout, for ease of presentation, we will simply say that the utilities are α -strongly concave, with the set \mathcal{R} being implicit.

Core. For an exchange economy with h consumers and ℓ goods, $\mathcal{E} = ((u_i, \omega_i))_{i=1}^h$, we define the following central notions:

- An *allocation* in \mathcal{E} is a vector, $\bar{x} = (\bar{x}_i)_{i=1}^h \in \mathbb{R}_+^{h\ell}$, such that $\sum_{i=1}^h \bar{x}_i = \sum_{i=1}^h \omega_i$. In other words, an allocation corresponds to a redistribution of the endowments among the consumers.
- A nonempty subset $S \subseteq [h]$ is a *coalition*. Let S be a coalition, then a vector $(y_i)_{i \in S}$ is an S -*allocation* if $\sum_{i \in S} y_i = \sum_{i \in S} \omega_i$.
- A coalition S *blocks* the allocation $\bar{x} = (\bar{x}_i)_{i=1}^h$ in \mathcal{E} if there exists an S -allocation $(y_i)_{i \in S}$ such that $u_i(y_i) > u_i(\bar{x}_i)$ for all $i \in S$. That is, a blocking coalition is a group of consumers that are better off trading among themselves, than they would have been under \bar{x} .
- The *core* of \mathcal{E} is the set of all allocations that are not blocked by any coalition.
- The κ -*core* of \mathcal{E} , for $\kappa \in \mathbb{Z}_+$, is the set of allocations that are not blocked by any coalition of cardinality at most κ . Hence, in an economy with n consumers, the core is the n -core and the set of individually rational allocations are the 1-core.

Equilibrium and approximate equilibrium. In an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$, a *Walrasian equilibrium* is a pair $(p, \bar{x}) \in \mathbb{R}_+^{\ell} \times \mathbb{R}_+^{h\ell}$ in which

1. $p \in \mathbb{R}_+^{\ell}$ is a price vector for the ℓ goods in the economy.
2. $\bar{x} = (\bar{x}_i)_{i \in [h]} \in \mathbb{R}_+^{h\ell}$ is an allocation, i.e., under \bar{x} supply equals the demand, $\sum_{i=1}^h \bar{x}_i = \sum_{i=1}^h \omega_i$.
3. Every consumer $i \in [h]$ maximizes its utility u_i , while consuming its endowment ω_i . That is, $p^T \bar{x}_i = p^T \omega_i$ and, for all bundles $y \in \mathbb{R}_+^{\ell}$ with the property that $u_i(y) > u_i(\bar{x}_i)$, we have $p^T y_i > p^T \omega_i$.

Next, we describe what we mean by an approximate Walrasian equilibrium. We use a notion of approximation in which consumers are optimizing exactly, subject to budget constraints that are satisfied approximately.

Formally, in an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$, ε -Walrasian equilibrium is a pair $(p, \bar{x}) \in \mathbb{R}_+^\ell \times \mathbb{R}_+^{h\ell}$ in which the (normalized) price vector $p \in \Delta$ and the allocation $\bar{x} \in \mathbb{R}_+^{h\ell}$ satisfy the following two conditions, for all consumers $i \in [h]$:

- (i) $|p^T \bar{x}_i - p^T \omega_i| \leq \varepsilon$ and
- (ii) for any bundle $y \in \mathbb{R}_+^\ell$, with the property that $u_i(y) > u_i(\bar{x}_i)$, we have $p^T y > p^T \omega_i - \varepsilon/h$.

Our notion of approximation is different from the standard use in algorithmic game theory, where agents are approximately optimizing, but it conforms to the most used definitions in economic theory; see Starr (1969); Arrow and Hahn (1971); Hildenbrand, Schmeidler, and Zamir (1973); Hildenbrand (1974); Anderson (1978, 1982); Anderson, Khan, and Rashid (1982); Mas-Colell (1978, 1979). Economists do not view utilities as objective, observable, entities. Hence it is natural in economics to measure the approximation error in monetary terms, using consumers' expenditure, instead of measuring approximation in terms of utility loss.

That said, it is straightforward to move between different notions of approximate equilibrium, and our result continues to apply. One can, instead, require that the approximation guarantee be in terms of the utilities. Or that budget exhaustion holds exactly (so the value of agents' consumption equals the value of their endowment), but the overall supply is only approximately equal to the demand. In Appendix F we show that one can obtain these kinds of approximations as well from the results in the current paper.

Also, note that part (ii) of the definition is stronger than the symmetric (to part (i)) version that requires $p^T y > p^T \omega_i - \varepsilon$ for every bundle $y \in \mathbb{R}_+^\ell$, with the property that $u_i(y) > u_i(\bar{x}_i)$.

Walrasian allocation. The term (*approximate*) *Walrasian allocation* refers to an allocation $\bar{x} \in \mathbb{R}_+^{h\ell}$ in \mathcal{E} for which there exists a price vector $p \in \Delta$ such that (p, \bar{x}) is a (approximate) Walrasian equilibrium.

Utility normalization. Our work addresses additive approximations, hence, we will follow the standard assumption (used in the context of absolute-error bounds) that the utilities are normalized so that $u_i(x_i) \in [0, 1]$ for all consumers $i \in [h]$ and all allocations $(x_i)_i \in \mathbb{R}_+^{h\ell}$. In fact, we will normalize such that $u_i(x_i) \in [0, 1 - \bar{\eta})$,

where $\bar{\eta} \in (0, 1)$ for all consumers $i \in [h]$ and all allocations $(x_i)_i \in \mathbb{R}_+^{h\ell}$. These normalizations are possible as the set of all allocations is compact and, hence, each utility function will achieve a maximum over the space of allocations.

Note that the utilities are normalized only for feasible allocations in the underlying exchange economy. A bundle which is, say, component-wise greater than the total endowment in the exchange economy—or, allocations in a replica economy (defined below)—can have utility arbitrarily greater than one.

Also, we will use λ to denote the Lipschitz constant of the utility functions, u_i s.

Replica Economies. Let $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ be an exchange economy over the commodity space \mathbb{R}_+^ℓ . The n -th replica of \mathcal{E} , for $n \geq 1$, is the exchange economy $\mathcal{E}^n = ((u_{i,t}, \omega_{i,t}))_{i \in [n], t \in [h]}$, with nh consumers. In \mathcal{E}^n the consumers are indexed by (i, t) , with index $i \in [n]$ and type $t \in [h]$, and they satisfy:

$$u_{i,t} = u_t \text{ and } \omega_{i,t} = \omega_t.$$

Replica economies constitute the basic model of a large economy with limited heterogeneity. See (Mas-Colell, Whinston, Green, et al., 1995, Chapter 18) for a textbook treatment.⁶

As mentioned previously, the utility normalization is considered only with respect to feasible allocations in the underlying economy \mathcal{E} .⁷ Indeed, the consumers' utilities scale up for feasible bundles in \mathcal{E}^n , since the amount of each good in the economy increases as a function of n .

An allocation $\bar{x} = (\bar{x}_{i,t})$ of \mathcal{E}^n is said to have the *equal treatment property* iff all the consumers with the same type are allocated identical bundles, i.e., iff $\bar{x}_{i,t} = \bar{x}_{i',t}$ for every $t \in [h]$ and $i, i' \in [n]$.

The following lemma is established in the work of Debreu and Scarf (1963) and it asserts that allocations contained in the h -core necessarily satisfy the equal treatment property. Debreu and Scarf state the result as a consequence of the core, or, in the terminology of our paper, the $(n \times h)$ -core. In actuality, the stronger statement below holds true. We include a proof in Appendix B for completeness, but the proof is really the same as that of Debreu and Scarf.

⁶The idea of replicas was suggested by Edgeworth himself in his discussion of the competitive hypothesis in large economies. He used X and Y to refer to two different agents, and wrote: "Let us now introduce a second X and a second Y ; so that the field of competition consists of two X s and two Y s. And for the sake of illustration (not of argument) let us suppose that the new X has the same requirements, the same nature as the old X ; and similarly that the new Y is equal-natured with the old." (Edgeworth, 1881).

⁷This, in particular, ensures that this normalization is compatible with the strict monotonicity of the utilities.

Lemma 1 (Equal treatment property). *Let the utilities of each consumer be strictly monotonic, continuous, and strictly concave in an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ and, hence, in the corresponding replica economy \mathcal{E}^n . Then, every h -core allocation of \mathcal{E}^n satisfies the equal treatment property.*

By Lemma 1, we can express the h -core allocations of \mathcal{E}^n as vectors in $\mathbb{R}_+^{h\ell}$; with the convention that consumers that have the same type receive the same bundle. This succinct representation also applies to our main result, which considers κ -cores, with $\kappa \geq h$. Hence, for brevity, we represent κ -core allocations⁸ of the replica economy \mathcal{E}^n as vectors in $\mathbb{R}_+^{h\ell}$.

The definition of a Walrasian equilibrium (both exact and approximate) naturally extends to replica economies \mathcal{E}^n , wherein the price vector $p \in \mathbb{R}_+^\ell$ and the allocation $\bar{x} \in \mathbb{R}_+^{nh\ell}$.

Furthermore, note that if (p, \bar{x}) is a (approximate) Walrasian equilibrium in \mathcal{E} , then p and an n -times replicated version of \bar{x} constitute a (approximate) Walrasian equilibrium in \mathcal{E}^n . Analogously, if \bar{x} is a Walrasian allocation in \mathcal{E} , then it⁹ is a Walrasian allocation in \mathcal{E}^n as well, for all $n \geq 1$.

Throughout, we will address allocations that satisfy the equal treatment property. Hence, we will simply state that $(p, \bar{x}) \in \mathbb{R}_+^\ell \times \mathbb{R}_+^{h\ell}$ is a (approximate) Walrasian equilibrium—or that $\bar{x} \in \mathbb{R}_+^{h\ell}$ is a (approximate) Walrasian allocation—without distinguishing between \mathcal{E} and \mathcal{E}^n in this context.

3 A Quantitative Core Convergence Theorem

“... the reason why the complex play of competition tends to a simple uniform result – what is arbitrary and indeterminate in contract between individuals becoming extinct in the jostle of competition – is to be sought in a principle which pervades all mathematics, the principle of limit, or law of great numbers as it might perhaps be called.”

(Francis Ysidro Edgeworth Edgeworth (1884))

The following classic result of Debreu and Scarf Debreu and Scarf (1963) establishes that, as n tends to infinity, each allocation in the core of \mathcal{E}^n is a Walrasian allocation.

⁸Note that in \mathcal{E}^n an arbitrary allocation (say, one that does not satisfy equal treatment) is, in fact, a vector in $\mathbb{R}_+^{nh\ell}$.

⁹Specifically, a replicated version of \bar{x}

Theorem 2 (Debreu-Scarf Core Convergence Theorem). *Assume that in an exchange economy \mathcal{E} —with h consumers and ℓ number of goods—the consumers’ utilities are strictly monotonic, continuous, and strictly quasiconcave. Furthermore, suppose that an allocation $\bar{x} \in \mathbb{R}_+^{h\ell}$ belongs to the core of \mathcal{E}^n for all $n \geq 1$. Then, \bar{x} is a Walrasian allocation, i.e., there exists a nonzero price vector $p \in \mathbb{R}_+^\ell$ such that (p, \bar{x}) is a Walrasian equilibrium.*

By contrast, our result is nonasymptotic—it shows that as long as n is quadratic in h , ℓ , and $1/\varepsilon$, a core allocation is an ε -Walrasian allocation. In fact, our result holds for κ -cores, where κ is $\mathcal{O}\left(\frac{h^2\ell}{\varepsilon^2}\right)$.

Theorem 3 (Main Result). *Assume that in an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ —with h consumers and ℓ number of goods—the utilities, u_i s, are strictly monotonic, continuously differentiable, and α -strongly concave. Furthermore, suppose that an allocation $\bar{x} \in \mathbb{R}_+^{h\ell}$ belongs to the κ -core of \mathcal{E}^n , for any*

$$n \geq \kappa \geq \frac{16}{\alpha} \left(\frac{\lambda \ell h}{\varepsilon} + \frac{h^2}{\varepsilon^2} \right).$$

Then, \bar{x} is an ε -Walrasian allocation (i.e., there exists a price vector $p \in \Delta$ such that (p, \bar{x}) is an ε -Walrasian equilibrium). Here, $\varepsilon > 0$ is the approximation parameter and λ is the Lipschitz constant of the utilities.

We will establish this theorem by showing that if an allocation \bar{x} is in the κ -core, then there exists a price vector $p \in \Delta$ such that (p, \bar{x}) is a ε -Walrasian equilibrium. It is relevant to note that, instead of relying on limiting arguments, we use nonasymptotic results (with relevant approximation guarantees), such as the approximate version of Carathéodory’s theorem (see Barman (2015) and references therein). Developing such a proof ensures that the underlying parameters (e.g., the core size κ) are only polynomially large (and not arbitrarily large). This approach not only provides us with a quantitative version of the core convergence theorem, but also an efficiently implementable characterization (Section 3.1). In Section 4 we will use the developed characterization to design an efficient algorithm that tests whether a given allocation is approximately Walrasian, or not.

We will begin by establishing (in Section 3.1) a useful geometric property that is satisfied by all allocations (Walrasian or otherwise). At a high level, this property shows that one can focus on a bounded subset of a specific convex hull, which in itself is unbounded. This bounding exercise essentially enables us to bypass asymptotic arguments, and prove the quantitative version of the core convergence theorem in Section 3.3.

The fact that we can work with a bounded set (and not an unbounded one) is also essential from an algorithmic standpoint. Specifically, it lets us apply the Ellipsoid method and develop (in Section 4) an efficient, testing algorithm for Walrasian allocations.

3.1 Bounded Hull

The result developed in this section applies to arbitrary allocations; allocations that might or might not be in the core. Consider an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ with h consumers and ℓ goods—in which the utilities, u_i s, are strictly monotonic, continuously differentiable, and α -strongly concave. As before, λ denotes the Lipschitz constant of the utility functions. Let $\delta := \varepsilon/h$.

Let $\bar{y} = (\bar{y}_i)_{i \in [h]}$ be an allocation in \mathcal{E} . With parameter $\eta \in (0, 1)$, for each consumer $i \in [h]$ we define $V_i^\eta := \{y \in \mathbb{R}_+^\ell \mid u_i(y) \geq u_i(\bar{y}_i) + \eta\}$ to denote the upper contour set with respect to the allocated bundle \bar{y}_i and margin η . Also, write Q_i^η to denote the set of trades from the endowment ω_i that render consumer i better off than the allocated bundle, $Q_i^\eta := \{z \in \mathbb{R}^\ell \mid u_i(z + \omega_i) \geq u_i(\bar{y}_i) + \eta\}$. By definition, $z \in Q_i^\eta$ iff $(z + \omega_i) \in V_i^\eta$.

For each consumer i , we also consider \widehat{Q}_i^η , a bounded subset of Q_i^η ; specifically,

$$\widehat{Q}_i^\eta := Q_i^\eta \cap \left\{ z \in \mathbb{R}^\ell : \|z\| \leq \sqrt{\frac{2(\lambda\ell\delta + 1)}{\alpha}} \right\},$$

The set \widehat{Q}_i^η is obtained by intersecting Q_i^η with the Euclidean ball of radius $\sqrt{\frac{2(\lambda\ell\delta + 1)}{\alpha}}$ and center $\mathbf{0}$.

Since the utility u_i is continuous and concave, the set Q_i^η is closed and convex. Therefore, the subset \widehat{Q}_i^η is compact (closed and bounded) as well as convex. The following lemma shows that, by bounding the sets in this manner, we do not lose out on an important containment property.

Lemma 4. *Let \bar{y} be an allocation in an economy \mathcal{E} with strictly monotonic, continuously differentiable, and strongly concave utilities. Suppose that the sets Q_i^η and \widehat{Q}_i^η , for $i \in [h]$, are as defined above, with parameters $\delta > 0$ and $\eta \in [0, \bar{\eta}]$. Then,*

$$(-\delta)\mathbf{1} \in \text{cvh} \left(\bigcup_{i=1}^h Q_i^\eta \right) \quad \text{iff} \quad (-\delta)\mathbf{1} \in \text{cvh} \left(\bigcup_{i=1}^h \widehat{Q}_i^\eta \right).$$

Proof. The reverse direction of the claim is direct, since $\widehat{Q}_i^\eta \subset Q_i^\eta$ for all $i \in [h]$.

For the forward direction, we have vectors $z_i \in Q_i^\eta$ and a convex combination $\lambda_i \geq 0$, for $i \in [h]$, such that $\sum_{i=1}^h \lambda_i = 1$ and

$$\sum_{i=1}^h \lambda_i z_i = (-\delta)\mathbf{1} \quad (1)$$

Let $R := \max_i \{\|z_i\| : i \in [h]\}$. By definition, the z_i s are contained in the (closed) Euclidean ball $B(R)$ of radius R and center $\mathbf{0}$. Note that, for each $i \in [h]$, the intersection $Q_i^\eta \cap B(R)$ is a compact set.

Let Z denote the collection of all tuples

$$(z'_1, z'_2, \dots, z'_h) \in (Q_1^\eta \cap B(R)) \times (Q_2^\eta \cap B(R)) \times \dots \times (Q_h^\eta \cap B(R))$$

for which there exists convex coefficients λ'_i s such that $\sum_i \lambda'_i z'_i \leq (-\delta)\mathbf{1}$, i.e., there exists a convex combination of z'_i s which is component-wise upper bounded by $(-\delta)\mathbf{1}$.¹⁰

From (1), we know that Z is nonempty. Given that the sets $(Q_i^\eta \cap B(R))$ s are compact, one can show that Z is compact as well. Hence, the problem of minimizing $\max \{\|z'_i - (-\delta)\mathbf{1}\| \mid (z'_i)_i \in Z\}$ admits an optimal solution, say $(z_t^*)_t$. Note that, by definition of Z , there exists convex coefficients $(\lambda_t^*)_{t \in H^*}$ that satisfy $\sum_{t \in H^*} \lambda_t^* z_t^* \leq (-\delta)\mathbf{1}$; here, subset $H^* \subseteq [h]$ is selected to ensure that $\lambda_t^* > 0$ for all $t \in H^*$.

Next, we will prove that $\|z_t^*\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$, for all $t \in H^*$. Subsequently, we will show that using z_t^* s we can obtain vectors $\tilde{z}_t \in Q_t^\eta$ that satisfy the same norm bound ($\|\tilde{z}_t\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$) and whose convex combination is equal to $(-\delta)\mathbf{1}$. This norm bound implies that $\tilde{z}_t \in \widehat{Q}_t^\eta$ and, hence, leads to the desired containment: $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \widehat{Q}_i^\eta\right)$.

Assume, by way of contradiction, that $\|z_i^*\| > \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$, for some $i \in H^*$. For such an i , consider bundle $x_i^* := z_i^* + \omega_i$. By definition, $z_i^* \in Q_i^\eta$ and, hence, $x_i^* \in V_i^\eta$ (i.e., x_i^* belongs to the upper contour set). Next, we establish a useful bound that holds in both the following (exhaustive) cases (i) x_i^* is in the range set \mathcal{R} of strong concavity (see Section 2) and (ii) x_i^* lies outside the range set \mathcal{R} . Note that u_i is α -strongly concave throughout \mathcal{R} , hence in case (i) we have

$$u_i(\omega_i) \leq u_i(x_i^*) + \nabla u_i(x_i^*)^T (\omega_i - x_i^*) - \frac{\alpha}{2} \|\omega_i - x_i^*\|^2 \quad (2)$$

¹⁰The definition of Q_i^η s provides a component-wise lower bound as well: $z_i \geq -\omega_i$, for each vector $z_i \in Q_i^\eta$.

Since $\|\omega_i - x_i^*\| = \|z_i^*\| > \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$, inequality (2) reduces to

$$\nabla u_i(x_i^*)^T x_i^* < \nabla u_i(x_i^*)^T \omega_i + (u_i(x_i^*) - u_i(\omega_i)) - \frac{\alpha}{2} \frac{2(\lambda\ell\delta+1)}{\alpha} \quad (3)$$

This bound holds in case (ii) as well, i.e., holds even if x_i^* lies outside \mathcal{R} . Here, we can consider a convex combination of ω_i and x_i^* , say vector \tilde{x}_i , that is at a distance r (as specified in Section 2) away from ω_i . The assumption on \mathcal{R} ensures that $\tilde{x}_i \in \mathcal{R}$ and $\frac{\alpha r^2}{2} \geq \lambda\delta\ell + 1$. Applying strong concavity with \tilde{x}_i we get $u_i(\omega_i) \leq u_i(\tilde{x}_i) + \nabla u_i(\tilde{x}_i)^T (\omega_i - \tilde{x}_i) - \frac{\alpha}{2} \|\omega_i - \tilde{x}_i\|^2 \leq u_i(\tilde{x}_i) + \nabla u_i(\tilde{x}_i)^T (\omega_i - \tilde{x}_i) - (\lambda\delta\ell + 1)$. Since $u_i(x_i^*) + \nabla u_i(x_i^*)^T (\omega_i - x_i^*) \geq u_i(\tilde{x}_i) + \nabla u_i(\tilde{x}_i)^T (\omega_i - \tilde{x}_i)$, inequality (3) follows. That is, this bound holds in general.

Next, observe that x_i^* must satisfy $u_i(x_i^*) = u_i(\bar{y}_i) + \eta < 1$:¹¹ if this is not the case (i.e., if we have $u_i(x_i^*) > u_i(\bar{y}_i) + \eta$), then by reducing a positive component¹² of $z_i^* = x_i^* - \omega_i$ we can ensure that z_i^* moves closer to $(-\delta)\mathbf{1}$ and at the same time z_i^* continues to be in Q_i^η (i.e., the following inequality continues to hold $u_i(x_i^*) \geq u_i(\bar{y}_i) + \eta$). Also, note that such a reduction maintains the containment of z_i^* s in Z ; specifically, the following inequality continues to hold $\sum_t \lambda_t^* z_t^* \leq (-\delta)\mathbf{1}$. A repeated application of this argument gives us $x_i^* = z_i^* + \omega_i$ with the property that $u_i(x_i^*) = u_i(\bar{y}_i) + \eta < 1$.

The utility function u_i is λ -Lipschitz, hence, its gradient at any point $x_i^* \in \mathbb{R}^\ell$ satisfies $\|\nabla u_i(x_i^*)\|_\infty \leq \lambda$. That is, $\|\nabla u_i(x_i^*)\|_1 \leq \lambda\ell$. Therefore, equation (3) gives us

$$\nabla u_i(x_i^*)^T x_i^* < \nabla u_i(x_i^*)^T (\omega_i + (-\delta)\mathbf{1}) \quad (4)$$

Now, we can apply Proposition 11 (stated and proved in Appendix A) with $x = x_i^*$, $w = (\omega_i + (-\delta)\mathbf{1})$, and inequality (4), to establish that there exists a positive $\mu \in (0, 1]$ with the property that $u_i(x_i^*) \leq u_i((1 - \mu)x_i^* + \mu(\omega_i + (-\delta)\mathbf{1}))$.

Rewriting, we get $u_i((1 - \mu)(x_i^* - \omega_i) + \mu(-\delta)\mathbf{1} + \omega_i) \geq u_i(x_i^*)$. Therefore,

$$\hat{z}_i := (1 - \mu)(x_i^* - \omega_i) + \mu(-\delta)\mathbf{1} = (1 - \mu)z_i^* + \mu(-\delta)\mathbf{1} \in Q_i^\eta,$$

as $x_i^* \in V_i^\eta$ and $z_i^* \in Q_i^\eta$.

¹¹Recall that the utility of any feasible allocation (and, hence, for \bar{y}_i) is normalized to be at most $1 - \bar{\eta}$. Also, we have $\eta < \bar{\eta}$.

¹²If all the components of z_i^* are negative, then $x_i^* \leq \omega_i$ and we get the desired bound $u_i(x_i^*) \leq u_i(\omega_i) \leq 1 - \bar{\eta}$. Here, the last inequality follows from the fact that the utility of any feasible allocation (and, hence, for ω_i) is normalized to be at most $1 - \bar{\eta}$.

The vector \widehat{z}_i is itself a convex combination of z_i^* and $(-\delta)\mathbf{1}$. This ensures that there exists a convex combination, $(\widehat{\lambda}_t)_{t \in H^*}$, such that

$$\widehat{\lambda}_i \widehat{z}_i + \sum_{t \in H^* \setminus \{i\}} \widehat{\lambda}_t z_t^* \leq (-\delta)\mathbf{1}.$$

Note that $\|\widehat{z}_i - (-\delta)\mathbf{1}\| = (1 - \mu)\|z_i^* - (-\delta)\mathbf{1}\| < \|z_i^* - (-\delta)\mathbf{1}\|$. This argument can be repeated for all other j s with $\|z_j^* - (-\delta)\mathbf{1}\| = \max\{\|z_t^* - (-\delta)\mathbf{1}\| \mid t \in H^*\}$, contradicting the definition (optimality) of z_t^* s.

Therefore, the desired upper bound on the norm holds: $\|z_t^*\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$, for all $t \in H^*$.

To complete the proof we will show that z_t^* s can be transformed into vectors $\widetilde{z}_t \in Q_t^\eta$ that satisfy the same norm bound and whose convex combination is equal to $(-\delta)\mathbf{1}$. Write $\varphi := \sum_t \lambda_t^* z_t^*$ and note that $\varphi \leq (-\delta)\mathbf{1}$. If component $a \in [\ell]$ of φ is strictly less than $-\delta$, then there exists a z_i^* such that its a th component is less than $-\delta$ as well: $z_{i,a}^* < -\delta$. We can increase $z_{i,a}^*$ till either it becomes equal to zero, or the a th component of φ reaches $-\delta$.¹³ Note that in this transformation while the a th component of z_i^* increases in value, it decreases in magnitude. Hence, the utility $u_i(z_i^* + \omega_i)$ increases and the norm of z_i^* decreases. Repeatedly applying this procedure gives us vectors $\widetilde{z}_t \in Q_t^\eta$ such that $\|\widetilde{z}_t\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$ and $\sum_t \lambda_t^* \widetilde{z}_t = (-\delta)\mathbf{1}$.

Overall, this implies that $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \widehat{Q}_i\right)$ and the stated claim follows. \square

3.2 Lemma for κ -Core Allocations

Recall the notation from Theorem 3 and write $\delta := \frac{\varepsilon}{h}$ along with $\gamma := \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$. Note that, in Theorem 3, the lower bound on κ is equal to $\frac{8\gamma^2}{\delta^2}$.

In this section we apply Lemma 4 to establish an important separation property satisfied by allocations in the κ -core of an economy. In contrast to the result developed in the previous section, the next lemma (Lemma 5) specifically addresses κ -core allocations.

In particular, given a κ -core allocation $\bar{x} = (\bar{x}_i)_{i \in [h]} \in \mathbb{R}_+^{h\ell}$, let $\bar{\eta} > 0$ be such that $u_i(\bar{x}_i) \in [0, 1 - \bar{\eta}]$ for all $i \in [h]$ —the normalization of the utilities (to lie in $[0, 1)$) ensures that such a $\bar{\eta}$ exists. Also, consider an arbitrarily small, but positive, parameter $\eta \in (0, \bar{\eta})$.

¹³Here, the convex coefficients, λ_i^* s, remain unchanged.

For each consumer $i \in [h]$, write $U_i^\eta := \{x \in \mathbb{R}_+^\ell \mid u_i(x) \geq u_i(\bar{x}_i) + \eta\}$ to denote the upper contour set with respect to the allocated bundle \bar{x}_i and margin η . Also, write P_i^η to denote the set of trades from the endowment ω_i that render consumer i better off than the allocated bundle, $P_i^\eta := \{z \in \mathbb{R}^\ell \mid u_i(z + \omega_i) \geq u_i(\bar{x}_i) + \eta\}$. By definition, $z \in P_i^\eta$ iff $(z + \omega_i) \in U_i^\eta$.

The next lemma provides an important characterization in terms of the convex hull of the P_i^η s. This lemma shows that under a κ -core allocation beneficial trades (i.e., vectors in P_i^η s) cannot be combined (as a convex combination) across consumers to obtain $(-\delta)\mathbf{1}$; recall that $\delta := \varepsilon/h$. At a high level, this corresponds to the fact that, under a κ -core allocation, mutually beneficial redistributions are not possible across consumers.

Lemma 5. *For any allocation $\bar{x} = (\bar{x}_i)_{i \in [h]}$ that belongs to the κ -core of an economy \mathcal{E}^n , we have*

$$(-\delta)\mathbf{1} \notin \text{cvh} \left(\bigcup_{i=1}^h P_i^\eta \right).$$

Proof. Write $\nu := (-\delta)\mathbf{1}$. Suppose, towards a contradiction, that $\nu \in \text{cvh} \left(\bigcup_{i=1}^h P_i^\eta \right)$. Then, applying Lemma 4, we get that there exist vectors $z_t^* \in P_t^\eta$ such that their norm is bounded, $\|z_t^*\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}} = \gamma$, and their convex hull contains ν :

$$\nu \in \text{cvh} (z_t^*)_{t \in H^*}$$

Here, H^* is a subset of $[h]$ with the property that all t in H^* have strictly positive weight in the convex combination that yields ν .

Therefore, via the approximate version of Carathéodory's theorem (see, e.g., Theorem 2 in Barman (2015)), we get that, for any integer $k \geq \frac{8\gamma^2}{\delta^2}$, there exists a vector $\nu' \in \text{cvh} (z_t^*)_{t \in H^*}$ which is δ close to ν (i.e., $\|\nu' - \nu\| < \delta$) and it satisfies

$$\nu' = \sum_{t \in H^*} \frac{\beta_t}{k} z_t^*,$$

with integers $\beta_t \in \mathbb{Z}_+$ summing up to k : $\sum_{t \in H^*} \beta_t = k$.

Choose an integer $k \geq \frac{8\gamma^2}{\delta^2}$ with $k \leq \kappa$. Then, ν' is δ -close to ν and it is a k -uniform convex combination of the z_t^* s. Given that the Euclidean distance between ν' and $\nu = (-\delta)\mathbf{1}$ is strictly smaller than δ , the vector ν' is componentwise strictly negative: $\nu' \ll \mathbf{0}$.

Now, consider a coalition S in \mathcal{E}^n defined by including $\beta_t \in \mathbb{Z}_+$ copies of consumers of type $t \in H^*$. Note that $|S| = k \leq \kappa$.

For each member of S of type t , let consumption bundle $y_{i,t}^* := z_t^* + \omega_t$. Since $z_t^* \in P_t^\eta$, we have

$$u_t(y_{i,t}^*) \geq u_t(\bar{x}_t) + \eta \quad \text{for all } (i, t) \in S.$$

Also, the fact that ν' is a k -uniform convex combination of the z_t^* s gives us

$$\frac{1}{k} \left(\sum_{(i,t) \in S} y_{i,t}^* - \sum_{(i,t) \in S} \omega_t \right) = \nu' \ll \mathbf{0}$$

Therefore, $\sum_{(i,t) \in S} y_{i,t}^* \ll \sum_{(i,t) \in S} \omega_t = \sum_{(i,t) \in S} \omega_{i,t}$. As a consequence, the coalition S , which has cardinality $k \leq \kappa$, blocks the given allocation $\bar{x} = (\bar{x}_i)_{i \in [h]}$, in the replica economy \mathcal{E}^n . This contradicts that \bar{x} belongs to the κ -core of \mathcal{E}^n . \square

3.3 Proof of Theorem 3

For the given κ -core allocation $\bar{x} = (\bar{x}_i)_{i \in [h]} \in \mathbb{R}_+^{h\ell}$ of the economy \mathcal{E}^n , we have $U_i^\eta := \{x \in \mathbb{R}_+^\ell \mid u_i(x) \geq u_i(\bar{x}_i) + \eta\}$ and $P_i^\eta := \{z \in \mathbb{R}^\ell \mid u_i(z + \omega_i) \geq u_i(\bar{x}_i)\}$, for each $i \in [h]$.

Lemma 5 guarantees that $(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h P_i^\eta\right)$. Hence, there exists a hyperplane between $(-\delta)\mathbf{1}$ and the convex hull; an implication of the hyperplane separation theorem.

In particular, let $p \in \mathbb{R}^\ell$ specify such a separating hyperplane: $p \cdot ((-\delta)\mathbf{1}) \leq p \cdot \text{cvh}\left(\bigcup_{i=1}^h P_i^\eta\right)$. The upward closure of P_i^η s (a consequence of strict monotonicity of the consumers' utility functions), and $p \neq \mathbf{0}$, ensures that $p > \mathbf{0}$. Hence, by scaling, we can assume that $p \in \Delta$.

We will show that p —as a price vector—certifies that \bar{x} is an ε -Walrasian allocation, i.e., (p, \bar{x}) is an ε -Walrasian equilibrium.

Recall that, in an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$, a pair $(q, \bar{y}) \in \Delta \times \mathbb{R}_+^{h\ell}$ (with price vector q and allocation \bar{y}) is deemed to be an ε -Walrasian equilibrium iff the following two conditions hold for every consumer i : (i) $|q^T \bar{y}_i - q^T \omega_i| \leq \varepsilon$ and (ii) for any bundle x , with $u_i(x) > u_i(\bar{y}_i)$, we have $q^T x > q^T \omega_i - \varepsilon/h$. To complete the proof, we will show that (p, \bar{x}) satisfies these conditions.

Note that, for each $i \in [h]$ and any $z = (x - \omega_i) \in P_i^\eta$, the separation by $p \in \Delta$ implies

$$p^T(x - \omega_i) = p^T z \geq p^T((-\delta)\mathbf{1}) = -\delta.$$

Hence, for any bundle $x \in U_i^\eta$ (i.e., for any bundle x that satisfies $u_i(x) \geq u_i(\bar{x}_i) + \eta$) the expenditure is at least the income (minus δ): $p^T x \geq p^T \omega_i - \delta$; recall that $\delta = \varepsilon/h$.

Here, the analysis holds for any $\eta > 0$, however small. That is, for a bundle $x \in \mathbb{R}_+^\ell$, with the property that, $u_i(x) > u_i(\bar{x}_i)$, we have

$$p^T x \geq p^T \omega_i - \delta \tag{5}$$

Therefore, (p, \bar{x}) satisfies the second condition in the definition of an ε -Walrasian equilibrium.

Finally, we will show that even under the allocated bundle \bar{x}_i , the expenditure is close to the income. Using the continuity of the utilities and a small enough η , we can apply inequality (5) to obtain $p^T \bar{x}_j \geq p^T \omega_j - \delta$ for all consumers $j \in [h]$.

Allocation $\bar{x} = (\bar{x}_j)_{j \in [h]}$ satisfies the equal treatment property (Lemma 1), hence $\sum_{j \in [h]} \bar{x}_j = \sum_{j \in [h]} \omega_j$. Therefore, for each consumer i , we have $(\bar{x}_i - \omega_i) = \sum_{j \in [h] \setminus \{i\}} (\omega_j - \bar{x}_j)$. Taking inner product with $p \in \Delta$, we get that (p, \bar{x}) satisfies the first condition in the definition of an ε -Walrasian equilibrium as well:

$$p^T (\bar{x}_i - \omega_i) = \sum_{j \in [h] \setminus \{i\}} p^T (\omega_j - \bar{x}_j) \leq (h - 1)\delta < \varepsilon.$$

Overall, we get that (p, \bar{x}) is an ε -Walrasian equilibrium and this completes the proof.

4 Testing Algorithm for Walrasian Allocations

This section develops a polynomial-time algorithm that efficiently determines whether a given allocation is an ε -Walrasian allocation, or not. Specifically, given an exchange economy¹⁴ $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ —with h consumers and ℓ goods—along with an allocation $\bar{y} = (\bar{y}_i)_{i \in [h]} \in \mathbb{R}^{h\ell}$, the developed algorithm efficiently finds a price vector $p \in \Delta$ (if one exists) such that (p, \bar{y}) is an ε -Walrasian equilibrium in \mathcal{E} . If no such price vector exists (i.e., \bar{y} is not an ε -Walrasian allocation), then the algorithm correctly reports as such.

The developed algorithm runs in time that is polynomial in the number of consumers h . The algorithm applies, in particular, to completely heterogeneous economies, in which all the consumers can be of different type. In other words, our algorithmic results are not confined to the replica-economy framework.

¹⁴Our algorithm only requires oracle access to the underlying utilities u_i s and their gradients. In particular, our algorithmic result will hold even in the absence of an explicit (say, a closed form) description of the utility functions.

The testing algorithm builds upon Lemma 4. As before, for any given allocation $\bar{y} = (\bar{y}_i)_{i \in [h]}$ (which might or might not be approximately Walrasian) and each consumer $i \in [h]$, we define the set $Q_i := \{z \in \mathbb{R}^\ell \mid u_i(z + \omega_i) \geq u_i(\bar{y}_i)\}$.¹⁵ We will prove (in Lemma 6 below) that the (non) containment of $(-\delta)\mathbf{1}$ in the convex hull of the Q_i s characterizes Walrasian equilibria; recall that $\delta := \varepsilon/h$.

This geometric characterization is interesting in its own right. However, given that the sets Q_i s are unbounded, this characterization (in terms of Q_i s), does not, in and of itself, translate into an efficient testing algorithm. Specifically, in order to apply the Ellipsoid method and test whether a vector is contained in a specific convex set, one requires the set to be bounded. The quantitative treatment developed in this paper—in particular, Lemma 4—enables us to bypass this issue. Specifically, we will show that it suffices to work with the a bounded subset of Q_i . Towards that end, we define

$$\widehat{Q}_i := Q_i \cap \left\{ z \in \mathbb{R}^\ell : \|z\| \leq \sqrt{\frac{2(\lambda\ell\delta + 1)}{\alpha}} \right\} \quad (6)$$

Here, α is the strong-concavity parameter of the utilities and λ is their Lipschitz constant.

As observed before, the set Q_i is closed and convex. Furthermore, the subset \widehat{Q}_i is compact (closed and bounded) and convex.¹⁶ Also, note that any vector z in Q_i of norm less than $\sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$, belongs to \widehat{Q}_i as well.

Lemma 6. *lemmalemmanoncontainment An allocation \bar{y} is an ε -Walrasian allocation in an economy \mathcal{E} iff*

$$(-\delta)\mathbf{1} \notin \text{cvh} \left(\bigcup_{i=1}^h \widehat{Q}_i \right).$$

Here, for each consumer $i \in [h]$, the set \widehat{Q}_i is as defined in equation (6).

The proof of this lemma is deferred to Appendix C. In light of Lemma 6, testing whether an allocation \bar{y} is approximately Walrasian, or not, reduces to determining

¹⁵That is, in terms of the notation developed in Section 3.1, we are considering Q_i^η with $\eta = 0$.

¹⁶We can enlarge \widehat{Q}_i to ensure that it always has a nonempty interior. For instance, we can set the the radius of the intersecting Euclidean ball to be, say, $\max \left\{ \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}, 2\|\bar{y}_i\| \right\}$; note that $\bar{y}_i \in Q_i$ and, hence, with this redefinition we have $\bar{y}_i \in \widehat{Q}_i$.

whether the vector $(-\delta)\mathbf{1}$ is contained in the convex hull of the \widehat{Q}_i s. Below, in Theorem 7, we develop an efficient separation oracle for the convex hull of the \widehat{Q}_i s—the testing algorithm is obtained by simply applying this separation oracle onto $(-\delta)\mathbf{1}$.

For designing the efficient oracle, we use the equivalence of optimization and separation Grötschel, Lovász, and Schrijver (2012). We show that (linear) optimization problems can be solved in polynomial time over the convex hull of the \widehat{Q}_i s. Therefore, we obtain the desired separation oracle. Note that this is a somewhat atypical application of the optimization-separation equivalence—we start with an optimization algorithm to obtain a separating one.

The running time of our algorithm is polynomial in the input size; in particular, the running time is polynomial in the bit complexity of the underlying parameters (including ε). Furthermore, the algorithm only requires oracle access to the utilities and their gradients. A proof of the following theorem appears in Appendix D.

Theorem 7. *Let \mathcal{E} be an exchange economy with monotonic, continuously differentiable, and strongly concave utilities. Then, there exists a polynomial-time algorithm that, given an allocation \bar{y} in \mathcal{E} , determines whether \bar{y} is an ε -Walrasian allocation.*

4.1 Testing Algorithm for Economies with Piecewise-Linear Concave Utilities

In this section we consider economies, $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$, in which consumers' utilities are piecewise-linear concave (PLC). In this PLC setting, for every agent, the utility of each consumption bundle is obtained by taking a minimum over a set of linear functions. Specifically, PLC utilities have the form $u_i(x) := \min_k \left\{ \sum_j U_{i,j}^k x_j + T_i^k \right\}$, here x_j is the amount of good j in the consumption bundle $x \in \mathbb{R}_+^\ell$ and the nonnegative parameters $U_{i,j}^k \in \mathbb{R}_+$ and $T_i^k \in \mathbb{R}_+$ define the k th linear function for agent i . These parameters are given as input to specify each agent's utility.

While PLC utilities are concave, they are not strongly concave. Hence, under PLC utilities, one cannot directly apply Theorem 7. However, we show that the ideas developed in the previous section can be adapted to obtain a polynomial-time algorithm for testing whether a given allocation is approximately Walrasian in a PLC economy. Finding equilibria (exact and approximate) is known to be computationally hard under PLC utilities (Chen, Dai, Du, and Teng (2009); Garg, Mehta, Vazirani, and Yazdanbod (2017)). Hence, the result in this section identifies an interesting dichotomy between testing and finding a Walrasian equilibrium.

First, we will establish a containment result analogous to Lemma 4. Using this containment result (and arguments similar to the ones developed in Theorem 7) we develop an efficient, testing algorithm for PLC economies in Theorem 10.

In the PLC context, a key observation (established below in Lemma 8) is that the containment property can be obtained by considering vectors of norm at most

$$\Lambda := \max_{i \in [h], x \in \mathbb{R}_+^\ell} \left\{ \|x - \omega_i\| : u_i(x) \leq u_i \left(\sum_{t=1}^h \omega_t \right) \right\} \quad (7)$$

Under PLC utilities, the bit complexity of Λ is polynomially bounded: specifically, for any consumption bundle $x \in \mathbb{R}_+^\ell$, with the property $u_i(x) \leq u_i(\sum_t \omega_t)$, we have, for each component $a \in [\ell]$, $x_a \leq \frac{u_i(\sum_t \omega_t)}{\min_k U_{i,a}^k} \leq \frac{u_i(\sum_t \omega_t)}{\min_{k,j} U_{i,j}^k}$.¹⁷ The bit complexity of this upper bound is polynomially large and, hence, Λ is also polynomially upper bounded, in terms of bit complexity.

Note that, this bound does not require the utilities to be normalized. In fact, the results developed in this section hold for any exchange economy wherein the utilities are concave and Λ (as defined in (7)) is appropriately bounded.

For any given allocation $\bar{y} = (\bar{y}_i)_{i \in [h]}$, consider the bundle $\bar{y}_i \in \mathbb{R}^\ell$ allocated to consumer i and, as before, write $Q_i := \{z \in \mathbb{R}^\ell \mid z + \omega_i \in \mathbb{R}_+^\ell \text{ and } u_i(z + \omega_i) \geq u_i(\bar{y}_i)\}$. In addition, we define a bounded subset of Q_i

$$\tilde{Q}_i := Q_i \cap \{z \in \mathbb{R}^\ell \mid \|z\| \leq \Lambda\} \quad (8)$$

For each consumer i , the subset \tilde{Q}_i is compact, convex, and has a nonempty interior.

Lemma 8. *Let \bar{y} be an allocation in an exchange economy \mathcal{E} with PLC utilities. Suppose that the sets Q_i and \tilde{Q}_i , for $i \in [h]$, are as defined above. Then, with parameter $\delta > 0$, we have*

$$(-\delta)\mathbf{1} \in \text{cvh} \left(\bigcup_{i=1}^h Q_i \right) \quad \text{iff} \quad (-\delta)\mathbf{1} \in \text{cvh} \left(\bigcup_{i=1}^h \tilde{Q}_i \right).$$

Proof. The proof of this claim is almost identical to that of Lemma 4—the difference being that here we use the bound provided by Λ , instead of relying on strong concavity.

¹⁷We can address the case in which one of the nonnegative coefficients $U_{i,j}^k$ is equal to zero. In particular, adapting the arguments in Lemma 8 one can obtain $x_a \leq \frac{u_i(\sum_t \omega_t)}{\min_{\{k,j: U_{i,j}^k \neq 0\}} U_{i,j}^k}$.

To begin with, note that the reverse direction of the claim is direct, since $\tilde{Q}_i \subset Q_i$ for all $i \in [h]$.

For the forward direction, we have vectors $z_i \in Q_i$ and a convex combination $\lambda_i \geq 0$, for $i \in [h]$, such that $\sum_{i=1}^h \lambda_i = 1$ and

$$\sum_{i=1}^h \lambda_i z_i = (-\delta)\mathbf{1} \quad (9)$$

Let $R := \max_i \{\|z_i\| : i \in [h]\}$. By definition, the z_i s are contained in the (closed) Euclidean ball $B(R)$ of radius R and center $\mathbf{0}$. Note that, for each $i \in [h]$, the intersection $Q_i \cap B(R)$ is a compact set.

Let Z denote the collection of all tuples

$$(z'_1, z'_2, \dots, z'_h) \in (Q_1 \cap B(R)) \times (Q_2 \cap B(R)) \times \dots \times (Q_h \cap B(R))$$

for which there exists convex coefficients λ'_i 's such that $\sum_i \lambda'_i z'_i \leq (-\delta)\mathbf{1}$, i.e., there exists a convex combination of z'_i 's which is component-wise upper bounded by $(-\delta)\mathbf{1}$.¹⁸

From (9), we know that Z is nonempty. Given that the sets $(Q_i \cap B(R))$ s are compact, one can show that Z is compact as well. Hence, the problem of minimizing $\max \{\|z'_i - (-\delta)\mathbf{1}\| \mid (z'_i)_i \in Z\}$ admits an optimal solution, say $(z_t^*)_t$. Note that, by definition of Z , there exists convex coefficients $(\lambda_t^*)_{t \in H^*}$ that satisfy $\sum_{t \in H^*} \lambda_t^* z_t^* \leq (-\delta)\mathbf{1}$; here, subset $H^* \subseteq [h]$ is selected to ensure that $\lambda_t^* > 0$ for all $t \in H^*$.

Next, we will prove that $\|z_t^*\| \leq \Lambda$, for all $t \in H^*$. Subsequently, we will show that using z_t^* s we can obtain vectors $\tilde{z}_t \in Q_t$ that satisfy the same norm bound ($\|\tilde{z}_t\| \leq \Lambda$) and whose convex combination is equal to $(-\delta)\mathbf{1}$. This norm bound implies that $\tilde{z}_t \in \tilde{Q}_t$ and, hence, leads to the desired containment: $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \tilde{Q}_i\right)$.

The bundle $x_i^* := z_i^* + \omega_i$ must satisfy $u_i(x_i^*) = u_i(\bar{y}_i) \leq u_i(\sum_t \omega_t)$: if this is not the case (i.e., we have a strict inequality $u_i(x_i^*) > u_i(\bar{y}_i)$), then by reducing a positive component¹⁹ of $z_i^* = x_i^* - \omega_i$ we can ensure that z_i^* moves closer to $(-\delta)\mathbf{1}$ and at the same time z_i^* continues to be in Q_i (i.e., the following inequality continues to hold $u_i(x_i^*) \geq u_i(\bar{y}_i)$). Also, note that such a reduction maintains the containment of z_i^* s in Z ; specifically, the following inequality continues to hold $\sum_t \lambda_t^* z_t^* \leq (-\delta)\mathbf{1}$. A repeated application of this argument gives us $x_i^* = z_i^* + \omega_i$ with the property that

¹⁸The definition of Q_i s provide a component-wise lower bound as well: $z_i \geq -\omega_i$, for each vector $z_i \in Q_i$.

¹⁹If all the components of z_i^* are negative, then $x_i^* \leq \omega_i$ and we get the desired bound $u_i(x_i^*) \leq u_i(\omega_i) \leq u_i(\sum_t \omega_t)$.

$u_i(x_i^*) = u_i(\bar{y}_i) \leq u_i(\sum_t \omega_t)$, for all i . Therefore, using the definition of Λ (see (7)), we obtain the stated bound $\|z_i^*\| \leq \Lambda$.

To complete the proof we will show that z_t^* s can be transformed into vectors $\tilde{z}_t \in Q_t$ that satisfy the same norm bound and whose convex combination is equal to $(-\delta)\mathbf{1}$. Write $\varphi := \sum_t \lambda_t^* z_t^*$ and note that $\varphi \leq (-\delta)\mathbf{1}$. If component $a \in [\ell]$ of φ is strictly less than $-\delta$, then there exists a z_i^* such that its a th component is also less than $-\delta$: $z_{i,a}^* < -\delta$. We can increase $z_{i,a}^*$ till either it becomes equal to zero, or the a th component of φ reaches $-\delta$.²⁰ Note that in this transformation while the a th component of z_i^* increases in value, it decreases in magnitude. Hence, the utility $u_i(z_i^* + \omega_i)$ increases and the norm of z_i^* decreases. Repeatedly applying this procedure gives us vectors $\tilde{z}_t \in Q_t$ such that $\|\tilde{z}_t\| \leq \Lambda$ and $\sum_t \lambda_t^* \tilde{z}_t = (-\delta)\mathbf{1}$.

Overall, this implies that $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \tilde{Q}_i\right)$ and the stated claim follows. \square

Lemma 8 leads to the following characterization for approximate equilibria in PLC economies.

Lemma 9. *An allocation \bar{y} is an ε -Walrasian allocation in a PLC economy \mathcal{E} iff*

$$(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h \tilde{Q}_i\right).$$

Here, for each consumer $i \in [h]$, the set \tilde{Q}_i is as defined above.

The proof of this result is identical to that of Lemma 6 and is omitted; one has to simply use Lemma 8, instead of Lemma 4.

We now establish the main result of this section.

Theorem 10. *There exists a polynomial-time algorithm that—given an allocation $\bar{y} = (\bar{y}_i)_{i \in [n]}$ in an exchange economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [n]}$ with PLC utilities—determines whether \bar{y} is an ε -Walrasian allocation, or not.*

Proof. The design of the testing algorithm for PLC utilities is quite similar to the method developed in the proof of Theorem 7. We present the details for the PLC setting for completeness.

As a direct consequence of Lemma 9 we have: testing for approximately Walrasian allocation corresponds to determining whether the vector $(-\delta)\mathbf{1}$ is contained in the convex hull of the \tilde{Q}_i s; see (8) for the definition of these sets.

²⁰Here, the convex coefficients, λ_i^* s, remain unchanged.

Write $\tilde{\mathcal{Q}} := \text{cvh} \left(\bigcup_{i=1}^h \tilde{Q}_i \right)$. We will develop an efficient algorithm, ALG, for solving linear optimization problems over $\tilde{\mathcal{Q}}$, i.e., for solving problems of the form

$$\max c^T z \quad \text{subject to } z \in \tilde{\mathcal{Q}} \quad (10)$$

Here, $c \in \mathbb{R}^\ell$ is an input vector. The equivalence of optimization and separation (see, e.g., Grötschel, Lovász, and Schrijver (2012)) implies that ALG can be used to design a polynomial-time algorithm SEP that provides a separation oracle for $\tilde{\mathcal{Q}}$. That is, using SEP we can perform the desired test of determining whether $(-\delta)\mathbf{1} \in \tilde{\mathcal{Q}}$, or not.

In order to apply the optimization-separation equivalence we need to ensure that $\tilde{\mathcal{Q}}$ is compact, convex, and has a nonempty interior. These properties are satisfied by \tilde{Q}_i s individually, hence they hold for $\tilde{\mathcal{Q}}$ as well. Therefore, we can evoke the equivalence (via an application of the Ellipsoid method over the *polar* of $\tilde{\mathcal{Q}}$) and obtain the algorithm SEP.

To develop the algorithm, ALG, that efficiently solves linear optimization problems of the form (10), we note that the feasible set $\tilde{\mathcal{Q}}$ is a convex hull of the \tilde{Q}_i s. Hence, for any $c \in \mathbb{R}^\ell$, an optimal solution of (10) can be obtained by solving

$$\max_{i \in [h]} \left(\max c^T z_i \quad \text{subject to } z_i \in \tilde{Q}_i \right) \quad (11)$$

Here, for each i , the decision variable $z_i \in \mathbb{R}^\ell$ lies in the set \tilde{Q}_i . Below we will provide, for each i , a polynomial-time algorithm, ALG_i , that solves the linear optimization problem over \tilde{Q}_i , i.e., ALG_i efficiently solves $\max c^T z_i$ subject to $z_i \in \tilde{Q}_i$. Hence, ALG can be obtained by directly taking a maximum over the (optimal) solutions obtained by the ALG_i s.

We will now complete the chain of arguments mentioned above by designing the optimization algorithm ALG_i . This algorithm is itself based on the Ellipsoid method. As detailed below, the gradients of the utility function u_i (at different points) can be used to separate \tilde{Q}_i from vectors that are not contained in it. Hence, with this separation technique in hand, we can apply the Ellipsoid method over \tilde{Q}_i to obtain ALG_i .²¹

Given a query vector $q \in \mathbb{R}^\ell$, it is easy to test if $q \in \tilde{Q}_i := Q_i \cap \{z \in \mathbb{R}^\ell : \|z\| \leq \Lambda\}$. We directly verify (a) $u_i(q + \omega_i) \geq u_i(\bar{y}_i)$ (to ensure that $q \in Q_i$) and (b) $\|q\| \leq \Lambda$.²²

²¹Recall that, \tilde{Q}_i s are compact, convex, and have a nonempty interior. Hence, the Ellipsoid method is applicable over these sets.

²²We do not have to compute Λ exactly. Here, for the algorithm, an upper bound (with polynomial bit complexity) of Λ suffices.

Consider the case in which $q \notin \tilde{Q}_i$. To run the Ellipsoid method (that underlies ALG_i), we need a separating hyperplane for such a $q \in \mathbb{R}^\ell$. There are two complementary (though, nonexclusive) cases either (i) $q \notin Q_i$ (i.e., $u_i(q + \omega_i) < u_i(\bar{y}_i)$) and (ii) $\|q\| > \Lambda$.

In case (i), the subgradient at $\chi := q + \omega_i$ provides the separating hyperplane: in particular, for the PLC utility u_i , let $\bar{k} \in \operatorname{argmin}_k \left\{ \sum_j U_{i,j}^k \chi_j + T_i^k \right\}$, i.e., $u_i(q + \omega_i) = \sum_j U_{i,j}^{\bar{k}} \chi_j + T_i^{\bar{k}}$. The (subgradient) vector $\pi := (U_{i,j}^{\bar{k}})_{j \in [\ell]}$ provides the desired separation. This follows from the fact that, for any $z \in \mathbb{R}^\ell$ we have

$$\begin{aligned} u_i(z + \omega_i) &\leq \pi^T (z + \omega_i) + T_i^{\bar{k}} \\ &= \pi^T (q + \omega_i) + T_i^{\bar{k}} + \pi^T (z - q) \\ &= u_i(q + \omega_i) + \pi^T (z - q) \end{aligned} \tag{12}$$

Specifically, if $z \in Q_i$, then $u_i(z + \omega_i) \geq u_i(\bar{y}_i) > u_i(q + \omega_i)$. Hence, using (12), we get the desired separation

$$\pi^T q < \pi^T z \quad \text{for all } z \in \tilde{Q}_i \subset Q_i.$$

In case (ii), the vector $\pi := -\frac{q}{\Lambda\|q\|}$ suffices. Note that

$$\pi^T q = -\frac{\|q\|}{\Lambda} < -1 \tag{13}$$

For any $z \in \tilde{Q}_i$ we have $\|z\| \leq \Lambda$. Now, the Cauchy-Schwartz inequality gives us $|\pi^T z| \leq \|\pi\| \|z\| = \frac{1}{\Lambda} \|z\| \leq 1$. This inequality along with (13) shows that π is indeed a separating hyperplane: $\pi^T q < \pi^T z$ for all $z \in \tilde{Q}_i$.

Overall, we observe that separation with respect to the \tilde{Q}_i s can be performed efficiently. Hence, via the Ellipsoid method, we obtain, for each i , the algorithm ALG_i that optimizes over \tilde{Q}_i .

Combining ALG_i s we get the optimization algorithm (over \tilde{Q}) ALG , which, in turn, leads to SEP (the desired algorithm that separates with respect to \tilde{Q}). \square

References

ANDERSON, R. M. (1978): “An elementary core equivalence theorem,” *Econometrica*, pp. 1483–1487.

- ANDERSON, R. M. (1982): “A Market Value Approach to Approximate Equilibria,” *Econometrica*, 50(1), 127–136.
- ANDERSON, R. M. (1992): “The core in perfectly competitive economies,” *Handbook of game theory with economic applications*, 1, 413–457.
- ANDERSON, R. M., M. A. KHAN, AND S. RASHID (1982): “Approximate equilibria with bounds independent of preferences,” *The Review of Economic Studies*, 49(3), 473–475.
- ARROW, K. J., AND F. H. HAHN (1971): *General competitive analysis*. Holden San Francisco.
- AUMANN, R. J. (1964): “Markets with a continuum of traders,” *Econometrica*, pp. 39–50.
- BARMAN, S. (2015): “Approximating nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory’s theorem,” in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 361–369. ACM.
- BOYD, S., AND L. VANDENBERGHE (2004): *Convex optimization*. Cambridge university press.
- CHEN, X., D. DAI, Y. DU, AND S.-H. TENG (2009): “Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities,” in *FOCS ’09: Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pp. 273–282, Washington, DC, USA. IEEE Computer Society.
- CODENOTTI, B., A. SABERI, K. VARADARAJAN, AND Y. YE (2006): “Leontief economies encode nonzero sum two-player games,” in *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 659–667. Society for Industrial and Applied Mathematics.
- DEBREU, G., AND H. SCARF (1963): “A limit theorem on the core of an economy,” *International Economic Review*, 4(3), 235–246.
- DENG, X., C. PAPADIMITRIOU, AND S. SAFRA (2002): “On the complexity of equilibria,” in *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pp. 67–71. ACM.
- EDGEWORTH, F. Y. (1881): *Mathematical psychics: An essay on the application of mathematics to the moral sciences*, vol. 10. Kegan Paul.

- (1884): “The rationale of exchange,” *Journal of the Statistical Society of London*, pp. 164–166.
- GARG, J., R. MEHTA, V. V. VAZIRANI, AND S. YAZDANBOD (2017): “Settling the complexity of Leontief and PLC exchange markets under exact and approximate equilibria,” in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 890–901. ACM.
- GRODAL, B. (1972): “A second remark on the core of an atomless economy,” *Econometrica*, pp. 581–583.
- GRÖTSCHEL, M., L. LOVÁSZ, AND A. SCHRIJVER (2012): *Geometric algorithms and combinatorial optimization*, vol. 2. Springer Science & Business Media.
- HILDENBRAND, W. (1974): *Core and Equilibria of a Large Economy*. Princeton university press.
- HILDENBRAND, W., D. SCHMEIDLER, AND S. ZAMIR (1973): “Existence of approximate equilibria and cores,” *Econometrica*, pp. 1159–1166.
- MAS-COLELL, A. (1978): “A note on the core equivalence theorem: How many blocking coalitions are there?,” *Journal of Mathematical Economics*, 5(3), 207–215.
- (1979): “A refinement of the core equivalence theorem,” *Economics Letters*, 3(4), 307–310.
- MAS-COLELL, A., M. D. WHINSTON, J. R. GREEN, ET AL. (1995): *Microeconomic theory*. Oxford University Press.
- SCHMEIDLER, D. (1972): “A remark on the core of an atomless economy,” *Econometrica*, 40(3), 579.
- STARR, R. M. (1969): “Quasi-equilibria in markets with non-convex preferences,” *Econometrica*, pp. 25–38.
- VAZIRANI, V. V., AND M. YANNAKAKIS (2010): “Market Equilibrium under Separable, Piecewise-Linear, Concave Utilities,” in *Proc. Innovations in Computer Science*, pp. 156–165, Beijing, China.
- VIND, K. (1972): “A third remark on the core of an atomless economy,” *Econometrica*, 40(3), 585.

A Supplementary Results for Section 3

Here, we state and establish the proposition used in the proof of Lemma 4

Proposition 11. *For a continuously differentiable and concave function $u : \mathbb{R}^\ell \mapsto \mathbb{R}$, let vector $x, w \in \mathbb{R}^\ell$ satisfy the inequality $\nabla u(x)^T x < \nabla u(x)^T w$. Then, there exists a positive $\mu \in (0, 1]$ such that $u(x) \leq u((1 - \mu)x + \mu w)$.*

Proof. Let $\zeta := \nabla u(x)^T(w - x) > 0$. Also, for parameter $\mu \in (0, 1]$, denote by x_μ the convex combination $(1 - \mu)x + \mu w$. As μ tends to zero, x_μ tends to x .

Given that u is concave,

$$\begin{aligned} u(x) &\leq u(x_\mu) + \nabla u(x_\mu)^T(x - x_\mu) \\ &= u(x_\mu) - \mu \nabla u(x_\mu)^T(w - x) \end{aligned} \tag{14}$$

The Cauchy-Schwarz inequality gives us $|\langle \nabla u(x_\mu) - \nabla u(x), (w - x) \rangle| \leq \|\nabla u(x_\mu) - \nabla u(x)\| \cdot \|w - x\|$. Since u is continuously differentiable, there exists a small enough, but positive, μ such that the right-hand-side of the previous inequality is strictly less than ζ . For such a $\mu > 0$ we have $\nabla u(x_\mu)^T(w - x) > 0$. Therefore, using inequality (14), we get that there exists a $\mu \in (0, 1]$ for which $u(x) \leq u((1 - \mu)x + \mu w)$. \square

B Proof of Lemma 1

This section shows that, if the utilities of the consumers are strictly monotonic, continuous, and strictly concave, then allocations in the h -core of a replica economy satisfy the equal treatment property.

Let $\bar{x} = (\bar{x}_{1,1}, \dots, \bar{x}_{1,h}, \dots, \bar{x}_{n,1}, \dots, \bar{x}_{n,h})$ be an h -core allocation of the economy \mathcal{E}^n . We will show that if \bar{x} does not satisfy the equal treatment property, then a coalition of size h —consisting of the worse off consumers of each type—will block the allocation \bar{x} , contradicting the fact that it is in the h -core.

Say, towards a contradiction, that \bar{x} does not satisfy the equal treatment property. Since each consumer of type $t \in [h]$ has the same utility function, we can select the worse off consumer among the n ones that have type t ; in particular, for each $t \in [h]$, let index $j^*(t) \in [n]$ be such that $u_t(\bar{x}_{j^*(t),t}) \leq u_t(\bar{x}_{i,t})$ for all $i \in [n]$.

Consider the size- h coalition $S := \{(j^*(t), t)\}_{t \in [h]}$ and the S -allocation $(\hat{x}_t)_{(t \in [h])}$ defined as $\hat{x}_t := \frac{1}{n} \sum_{i=1}^n \bar{x}_{i,t}$.

Note that $(\hat{x}_t)_{(t \in [h])}$ is indeed an S -allocation: Since \bar{x} is an allocation in \mathcal{E}^n , it satisfies $\sum_{i=1}^n \sum_{t=1}^h \bar{x}_{i,t} = \sum_{i=1}^n \sum_{t=1}^h \omega_t = \sum_{t=1}^h n \omega_t$. Dividing by n gives us

$\sum_{t=1}^h \hat{x}_t = \sum_{t=1}^h \omega_t$. Therefore, the consumers in the coalition S can trade among themselves and, individually, obtain bundles \hat{x}_{tS} .

Furthermore, the strict concavity of u_t s ensures that $u_t(\hat{x}_t) \geq u_t(\bar{x}_{j^*(t),t})$ for all $(j^*(t), t) \in S$ —with one of the inequalities being strict. Hence, coalition S blocks the allocation \bar{x} and, by way of contradiction, the stated claim follows.

C Proof of Lemma 6

In this section we restate and prove Lemma 6.

Lemma 12. *lemmalemmanoncontainment* An allocation \bar{y} is an ε -Walrasian allocation in an economy \mathcal{E} iff

$$(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h \hat{Q}_i\right).$$

Here, for each consumer $i \in [h]$, the set \hat{Q}_i is as defined in equation (6).

Proof. First, we will consider the case wherein $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \hat{Q}_i\right)$ and show that this containment implies that \bar{y} is not an ε -Walrasian allocation. Then, we will address the complementary case and prove that if $(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h \hat{Q}_i\right)$, then \bar{y} is indeed an ε -Walrasian allocation.

The containment $(-\delta)\mathbf{1} \in \text{cvh}\left(\bigcup_{i=1}^h \hat{Q}_i\right)$ implies that there exists vectors $z_i \in \hat{Q}_i \subset Q_i$, with $i \in [h]$, along with convex coefficients λ_i s such that $\sum_i \lambda_i z_i = (-\delta)\mathbf{1}$. Here, $\lambda_i \geq 0$, for all i , and $\sum_i \lambda_i = 1$.

For each i , write $x_i = z_i + \omega_i$. Hence, we have

$$\sum_i \lambda_i (x_i - \omega_i) = (-\delta)\mathbf{1} \tag{15}$$

Furthermore, given that $z_i \in Q_i$, the definition of Q_i implies $u_i(x_i) \geq u_i(\bar{y}_i)$, for each i . Say, towards a contradiction, that \bar{y} is an ε -Walrasian allocation. That is, there exists a price vector $p \in \Delta$ such that (p, \bar{y}) is a ε -Walrasian equilibrium. Taking inner product with p on both sides of the equation (15), we obtain $\sum_i \lambda_i p^T (x_i - \omega_i) = -\delta$. Since λ_i s are convex coefficients, there exists an index j such that $p^T x_j \leq p^T \omega_j - \delta$. With δ slightly greater than ε/h , we can in fact obtain a strict inequality $p^T x_j < p^T \omega_j - \varepsilon/h$. Also, note that $u_j(x_j) \geq u_j(\bar{y}_j)$. Hence, increasing x_j (componentwise) by a small amount leads to a bundle \hat{x}_j with the property that $u_j(\hat{x}_j) > u_j(\bar{y}_j)$

and $p^T \hat{x}_j \leq p^T \omega_j - \varepsilon/h$. The existence of \hat{x}_j contradicts the fact that (p, \bar{y}) is an ε -Walrasian equilibrium.

To complete the proof we now consider the complementary case: $(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h \widehat{Q}_i\right)$.

Using Lemma 4 (in contrapositive form, with $\eta = 0$) we get $(-\delta)\mathbf{1} \notin \text{cvh}\left(\bigcup_{i=1}^h Q_i\right)$.

Consider a hyperplane between $(-\delta)\mathbf{1}$ and this convex hull: $p \cdot ((-\delta)\mathbf{1}) \leq p \cdot \text{cvh}\left(\bigcup_{i=1}^h Q_i\right)$. The hyperplane separation theorem guarantees the existence of such a $p \in \mathbb{R}^\ell$. Since Q_i is upward closed, p is componentwise nonnegative. Also, note that $p \neq \mathbf{0}$. Therefore, by scaling, we can assume that $p \in \Delta$. Next, we will show that (p, \bar{y}) is an ε -Walrasian equilibrium.

For each consumer $j \in [h]$ and any $z \in Q_j$, consider the vector $y = z + \omega_j$. The definition of Q_j implies $u_j(y) \geq u_j(\bar{y}_j)$. The separating property of hyperplane p gives us $p^T(y - \omega_j) = p^T z \geq p^T(-\delta)\mathbf{1} = -\delta$. Therefore, for each $j \in [h]$, the second condition in the definition of an approximate Walrasian equilibrium is satisfied: for any y with the property that $u_j(y) > u_j(\bar{y}_j)$ we have

$$p^T y \geq p^T \omega_j - \delta \tag{16}$$

Inequality (16) holds, in particular, for \bar{y}_j s: $p^T \bar{y}_j \geq p^T \omega_j - \delta$. In addition, given that \bar{y} is an allocation in the economy \mathcal{E} we have that, under \bar{y} , the supply is equal to the demand: $\sum_{j=1}^h \bar{y}_j = \sum_{j=1}^h \omega_j$.

Hence, for any fixed i , $\bar{y}_i - \omega_i = \sum_{j \in [h] \setminus \{i\}} (\omega_j - \bar{y}_j)$. Multiplying both sides of this equality with p^T we establish the first condition that defines an ε -Walrasian equilibrium:

$$\begin{aligned} p^T(\bar{y}_i - \omega_i) &= \sum_{j \in [h] \setminus \{i\}} p^T(\omega_j - \bar{y}_j) \\ &\leq \sum_{j \in [h] \setminus \{i\}} \delta && \text{(using (16) with } y = \bar{y}_j) \\ &= (h-1)\delta < \varepsilon \end{aligned}$$

That is, $|p^T \bar{y}_i - p^T \omega_i| \leq \varepsilon$. Hence, (p, \bar{y}) is an ε -Walrasian equilibrium and the stated claim follows. \square

D Proof of Theorem 7

Here we restate and prove Theorem 7

Theorem 13. *Let \mathcal{E} be an exchange economy with monotonic, continuously differentiable, and strongly concave utilities. Then, there exists a polynomial-time algorithm that, given an allocation \bar{y} in \mathcal{E} , determines whether \bar{y} is an ε -Walrasian allocation.*

Proof. As a direct consequence of Lemma 6 we have: testing for approximately Walrasian allocation corresponds to determining whether the vector $(-\delta)\mathbf{1}$ is contained in the convex hull of the \widehat{Q}_i s; see equation (6) for the definition of these sets.

Write $\widehat{Q} := \text{cvh}\left(\bigcup_{i=1}^h \widehat{Q}_i\right)$. We will develop an efficient algorithm, ALG, for solving linear optimization problems over \widehat{Q} , i.e., for solving problems of the form

$$\max c^T z \quad \text{subject to } z \in \widehat{Q} \quad (17)$$

Here, $c \in \mathbb{R}^\ell$ is an input vector.

The equivalence of optimization and separation (see, e.g., Grötschel, Lovász, and Schrijver (2012)) implies that ALG can be used to design a polynomial-time algorithm SEP that provides a separation oracle for \widehat{Q} . That is, using SEP we can perform the desired test of determining whether $(-\delta)\mathbf{1} \in \widehat{Q}$, or not.

In order to apply the optimization-separation equivalence we need to ensure that \widehat{Q} is compact, convex, and has a nonempty interior. These properties are satisfied by \widehat{Q}_i s individually, hence they hold for \widehat{Q} as well. Therefore, we can evoke the equivalence (via an application of the Ellipsoid method over the *polar* of \widehat{Q}) and obtain the algorithm SEP.

To develop the algorithm, ALG, that efficiently solves linear optimization problems of the form (17), we note that the feasible set \widehat{Q} is a convex hull of the \widehat{Q}_i s. Hence, for any $c \in \mathbb{R}^\ell$, an optimal solution of (17) can be obtained by solving

$$\max_{i \in [h]} \left(\max c^T z_i \quad \text{subject to } z_i \in \widehat{Q}_i \right) \quad (18)$$

Here, for each i , the decision variable $z_i \in \mathbb{R}^\ell$ lies in the set \widehat{Q}_i . Below we will provide, for each i , a polynomial-time algorithm, ALG_i , that solves the linear optimization problem over \widehat{Q}_i , i.e., ALG_i efficiently solves $\max c^T z_i$ subject to $z_i \in \widehat{Q}_i$. Hence, ALG can be obtained by directly taking a maximum over the (optimal) solutions obtained by the ALG_i s.

We will now complete the chain of arguments mentioned above by designing the optimization algorithm ALG_i . This algorithm is itself based on the Ellipsoid method. As detailed below, the gradients of the utility function u_i (at different points) can be used to separate \widehat{Q}_i from vectors that are not contained in it. Hence, with this

separation technique in hand, we can apply the Ellipsoid method over \widehat{Q}_i to obtain ALG_i .²³

Given a query vector $q \in \mathbb{R}^\ell$, it is easy to test if $q \in \widehat{Q}_i := Q_i \cap \left\{ z \in \mathbb{R}^\ell : \|z\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}} \right\}$. We directly verify (i) $u_i(q + \omega_i) \geq u_i(\bar{y}_i)$ (to ensure that $q \in Q_i$) and (ii) $\|q\| \leq \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$.

Consider the case in which $q \notin \widehat{Q}_i$. To run the Ellipsoid method (that underlies ALG_i), we need a separating hyperplane for such a $q \in \mathbb{R}^\ell$. There are two complementary (though, nonexclusive) cases either (i) $q \notin Q_i$ (i.e., $u_i(q + \omega_i) < u_i(\bar{y}_i)$) and (ii) $\|q\| > \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$.

In case (i), the gradient at $q + \omega_i$ (i.e., $\nabla u_i(q + \omega_i) \in \mathbb{R}_+^\ell$) provides the separating hyperplane: utility u_i is concave, hence $u_i(z + \omega_i) \leq u_i(q + \omega_i) + \nabla u_i(q + \omega_i)^T(z + \omega_i - q - \omega_i)$, for any $z \in \mathbb{R}^\ell$. Specifically, if $z \in Q_i$, then $u_i(z + \omega_i) \geq u_i(\bar{y}_i) > u_i(q + \omega_i)$. Using the previous two inequalities we get the desired separation, via $\pi := \nabla u_i(q + \omega_i)$

$$\pi^T q < \pi^T z \quad \text{for all } z \in \widehat{Q}_i \subset Q_i.$$

In case (ii), the vector $\pi := -\frac{q}{\gamma\|q\|}$ suffices; here $\gamma = \sqrt{\frac{2(\lambda\ell\delta+1)}{\alpha}}$. Note that

$$\pi^T q = -\frac{\|q\|}{\gamma} < -1 \tag{19}$$

For any $z \in \widehat{Q}_i$ we have $\|z\| \leq \gamma$. Now, the Cauchy-Schwartz inequality gives us $|\pi^T z| \leq \|\pi\| \|z\| = \frac{1}{\gamma} \|z\| \leq 1$. This inequality along with (19) shows that π is indeed a separating hyperplane: $\pi^T q < \pi^T z$ for all $z \in \widehat{Q}_i$.

Overall, we observe that separation with respect to the \widehat{Q}_i s can be performed efficiently. Hence, via the Ellipsoid method, we obtain, for each i , the algorithm ALG_i that optimizes over \widehat{Q}_i .

Combining ALG_i s we get the optimization algorithm (over $\widehat{\mathcal{Q}}$) ALG , which, in turn, leads to SEP (the desired algorithm that separates with respect to $\widehat{\mathcal{Q}}$). \square

Remark 1. The proof of Theorem 7 shows that if allocation \bar{y} is an ε -Walrasian allocation, then the hyperplane separating the vector $\nu := (-\delta)\mathbf{1}$ from $\mathcal{Q} := \text{cvh}(\cup_i Q_i)$ provides the equilibrium prices $p \in \Delta$. That is, if vector $p \in \Delta$ satisfies $p \cdot \nu \leq p \cdot \mathcal{Q}$,

²³Recall that, \widehat{Q}_i s are compact, convex, and have a nonempty interior. Hence, the Ellipsoid method is applicable over these sets.

then (p, \bar{y}) is an ε -Walrasian equilibrium. Note that such a price vector can be obtained by considering $\mathcal{P}_{\mathcal{Q}}(\nu) \in \mathcal{Q}$, the projection (under Euclidean distance) of ν onto \mathcal{Q} . In particular, via the variational characterization of convex projections, we have $(z - \mathcal{P}_{\mathcal{Q}}(\nu))^T(\nu - \mathcal{P}_{\mathcal{Q}}(\nu)) \leq 0$ for all $z \in \mathcal{Q}$. That is, $(\mathcal{P}_{\mathcal{Q}}(\nu) - \nu)^T z \geq (\mathcal{P}_{\mathcal{Q}}(\nu) - \nu)^T \mathcal{P}_{\mathcal{Q}}(\nu) \geq (\mathcal{P}_{\mathcal{Q}}(\nu) - \nu)^T \nu$ for all $z \in \mathcal{Q}$. Hence, with $p = \frac{\mathcal{P}_{\mathcal{Q}}(\nu) - \nu}{\|\mathcal{P}_{\mathcal{Q}}(\nu) - \nu\|_1}$, we get the desired separation.

The norm of $\mathcal{P}_{\mathcal{Q}}(\nu)$ is polynomially bounded: note that $\|\mathcal{P}_{\mathcal{Q}}(\nu) - \nu\| \leq \min_i \|\bar{y}_i - \nu\|$, since $\bar{y}_i \in \mathcal{Q}$. Another relevant observation is that Lemma 4 is not confined to $(-\delta)\mathbf{1}$ —we can establish such a containment result for any vector $q \in \mathbb{R}^\ell$ as long as we take obtain \widehat{Q}_i s by intersection Q_i s with a large enough ball. The radius of the ball just has to be polynomially large in $\|q\|$. Hence, with large enough \widehat{Q}_i s, we can ensure that $\mathcal{P}_{\mathcal{Q}}(\nu) \in \widehat{\mathcal{Q}} := \text{cvh}\left(\cup_i \widehat{Q}_i\right)$.

Additionally, the set containment $\widehat{\mathcal{Q}} \subset \mathcal{Q}$ implies the projection of ν onto $\widehat{\mathcal{Q}}$, say $\mathcal{P}_{\widehat{\mathcal{Q}}}(\nu) \in \mathbb{R}^\ell$, is the same as the desired vector $\mathcal{P}_{\mathcal{Q}}(\nu)$.

In the proof of Theorem 7 we have developed a polynomial-time separation oracle for $\widehat{\mathcal{Q}}$. Therefore, via the Ellipsoid method, the projection $\mathcal{P}_{\widehat{\mathcal{Q}}}(\nu)$ can be computed efficiently and, hence, we can find the equilibrium prices $p = \frac{\mathcal{P}_{\widehat{\mathcal{Q}}}(\nu) - \nu}{\|\mathcal{P}_{\widehat{\mathcal{Q}}}(\nu) - \nu\|_1}$.

E Illustrative Example of Strongly Concave Utilities

This section provides a simple example to illustrate the interplay of strong concavity and other parameters related to the utility functions. Consider an economy $\mathcal{E} = ((u_i, \omega_i))_{i \in [h]}$ (with h consumers and ℓ goods) wherein the utility of each consumer is

$$u(x) := \frac{1}{N} \sum_{j=1}^{\ell} \sqrt{x_j + \theta}.$$

Here, x_j is the amount of good j (present in the consumption bundle $x \in \mathbb{R}_+^\ell$), $\theta > 0$ is a fixed constant, and $N > 0$ is a normalization term. Here, dividing by a large enough, but fixed, parameter N ensures that for each feasible bundle $x \in \mathbb{R}_+^\ell$ in \mathcal{E} , we have $u(x) \in (0, 1)$. In particular, let $\bar{\omega}_j$ denote the total amount of good $j \in [\ell]$ present in the economy ($\bar{\omega}_j := \sum_i \omega_{i,j}$). Then, $N = \sum_{j=1}^{\ell} \sqrt{\bar{\omega}_j + \theta}$.

Also, note that the additive shift of θ ensures that the gradient of u is bounded for all $x \in \mathbb{R}_+^\ell$ and the Lipschitz constant of the utilities $\lambda = \frac{1}{2\sqrt{\theta}}$.

One can show that a function $u : \mathbb{R}^\ell \mapsto \mathbb{R}$ is α -strongly concave in a set $\mathcal{R} \subset \mathbb{R}^\ell$ iff $f(x) := u(x) + \frac{\alpha}{2}\|x\|^2$ is concave within \mathcal{R} . That is, for any $\alpha > 0$, if the Hessian of $f(x)$ is negative semidefinite for all $x \in \mathcal{R}$, then u is α strongly concave in \mathcal{R} .

The utility function u is separable across the ℓ goods, hence its Hessian (at point $x \in \mathbb{R}_+^\ell$) is a diagonal matrix with the (diagonal) entries being $\frac{-1}{4N(x_j+\theta)^{3/2}}$ for each $j \in [\ell]$. That is, for an $\alpha > 0$ and $x \in \mathbb{R}_+^\ell$, the Hessian of f is again a diagonal matrix with entries $\frac{-1}{4N(x_j+\theta)^{3/2}} + \alpha$, with $j \in [\ell]$.

Therefore, within the Euclidean ball of radius r (and center $\mathbf{0}$), the function u is $\left(\frac{1}{4N(r+\theta)^{3/2}}\right)$ -strongly concave; this value of α ensures that f is negative semidefinite throughout the ball.

As stated in Section 2, we require strongly concavity to hold within an appropriately large set around the endowments.²⁴ Specially, we need $\alpha r^2 \geq \frac{2\varepsilon\lambda\ell}{h} + 2$ (see Section 2). Using the above-mentioned expression for the modulus of strong concavity, α , the required inequality translates to

$$\frac{\sqrt{r}}{4N(1 + \theta/r)^{3/2}} \geq \frac{2\varepsilon\lambda\ell}{h} + 2$$

The left-hand-side of the previous inequality is an increasing function of r . Hence, this inequality holds for an appropriately large r (which only depends on the parameters of \mathcal{E} and not on the replication factor n). With this r and the corresponding value of α in hand one can apply the results developed in this work.

F Other Notions of Approximate Equilibria

In this section we will show that the notion of approximate Walrasian equilibrium considered in this work relates to other notions studied in computer science and economics.

A common alternate (to the one considered in this work) definition of approximation seeks to relax how exactly the consumers optimize, see, e.g., Deng, Papadimitriou, and Safra (2002) and Garg, Mehta, Vazirani, and Yazdanbod (2017). This notion requires that the consumers are approximately maximizing their utilities, though it assumes that the demand is approximately equal to the supply. Specifically, under this complementary notion, a pair $(p, \bar{x}) \in \Delta \times \mathbb{R}_+^{h\ell}$, with price vector $p \in \Delta$ and allocation \bar{x} , is said to be an $\hat{\varepsilon}$ -equilibrium (in an economy $\mathcal{E} = ((u_i, \omega_i))_i$

²⁴For ease of presentation, here we assume that the endowments are $\mathbf{0}$. Shifting all the vectors by $-\omega_i$, directly provide the arguments for general endowment vectors.

with h consumers and ℓ goods) iff the following conditions hold for all consumers $i \in [h]$:²⁵

(C1) For any bundle $y \in \mathbb{R}_+^\ell$, with the property that $p^T y \leq p^T \omega_i$, we have $u_i(y) \leq u_i(\bar{x}_i) + \hat{\varepsilon}$.

(C2) $|p^T \bar{x}_i - p^T \omega_i| \leq \hat{\varepsilon}$.

Note that, (C1) corresponds to the following requirement: $u_i(\bar{x}_i) \geq \left(\max_{y \in \mathbb{R}_+^\ell} \{u_i(y) \mid p^T y \leq p^T \omega_i\} \right) - \hat{\varepsilon}$.

We will show that an ε -approximate Walrasian equilibrium (as defined in Section 2) satisfies conditions (C1) and (C2) with $\hat{\varepsilon} = \frac{\varepsilon \lambda \sqrt{\ell}}{h}$; here, λ is the Lipschitz constant of the utilities u_i s.

By definition (see Section 2), an ε -approximate Walrasian equilibrium, say (p, \bar{x}) , satisfies (C2).²⁶

To establish the claim, next we will prove the contrapositive form of (C1) holds for (p, \bar{x}) . Consider a bundle $y \in \mathbb{R}_+^\ell$ with high utility $u_i(y) > u_i(\bar{x}_i) + \hat{\varepsilon} = u_i(\bar{x}_i) + \frac{\varepsilon \lambda \sqrt{\ell}}{h}$.

Observe that the Euclidean ball $B\left(y, \frac{\varepsilon \sqrt{\ell}}{h}\right)$, with center y and radius $\frac{\varepsilon \sqrt{\ell}}{h}$, is entirely contained in the upper contour set $\{x \in \mathbb{R}_+^\ell \mid u_i(x) \geq u_i(\bar{x}_i)\}$. Say, towards a contradiction, that this is not the case. Then, there exists a bundle $y' \in \mathbb{R}_+^\ell$ on the boundary of the upper contour set that is at most a distance $\frac{\varepsilon \sqrt{\ell}}{h}$ away from y , i.e., a boundary point y' that satisfies $\|y - y'\| \leq \frac{\varepsilon \sqrt{\ell}}{h}$. Since y' is at the boundary of the upper contour set, $u_i(y') = u_i(\bar{x}_i)$. This fact contradicts the high utility of y : $|u_i(y) - u_i(y')| \leq \lambda \|y - y'\| \leq \frac{\varepsilon \lambda \sqrt{\ell}}{h}$.

Therefore, in particular, the vector $\left(y - \frac{\varepsilon \sqrt{\ell}}{h} \frac{p}{\|p\|}\right) \in B\left(y, \frac{\varepsilon \sqrt{\ell}}{h}\right)$ belongs to the upper contour set. That is, $u_i\left(y - \frac{\varepsilon \sqrt{\ell}}{h} \frac{p}{\|p\|}\right) \geq u_i(\bar{x}_i)$.

Hence, our definition of ε -Walrasian equilibria ensures that

$$p^T \left(y - \frac{\varepsilon \sqrt{\ell}}{h} \frac{p}{\|p\|}\right) \geq p^T \omega_i - \varepsilon/h \quad (20)$$

Since the price vector $p \in \Delta$, we have $\|p\| \geq \frac{1}{\sqrt{\ell}}$. Rearranging inequality (20) shows that y satisfies (the contrapositive form of) the condition (C1): $p^T y \geq p^T \omega_i$.

²⁵Note that (C1) and (C2) are a strengthening of the alternate formulation mentioned above, since here we insist that the demand is exactly equal to the supply.

²⁶Here, assume that $\hat{\varepsilon} \geq \varepsilon$, otherwise we can consider $\hat{\varepsilon} = \max\left\{\frac{\varepsilon \lambda \sqrt{\ell}}{h}, \varepsilon\right\}$.

Supply vs demand and average budget gap. In contrast to some of the other formulations, the definition of approximate equilibria considered in this work strictly enforces that the supply is equal to the demand. Also, we consider a bound on the budget gap, $|p^T \bar{x}_i - p^T \omega_i|$, for every consumer i . Hence, our approximate equilibrium satisfies the average budget gap requirement considered in Mas-Colell (1979).