The Competitive Core of Combinatorial Exchange

Simon Jantschgi* Thành Nguyen† Alexander Teytelboym‡

August 2024

Abstract

We consider combinatorial exchanges where agents have (possibly random) endowments and ordinal preferences over bundles of indivisible goods. We focus on the *ordinal core*—the set of all individually rational lottery assignments which cannot be blocked by a coalition whose members trade probability shares and receive first-order stochastically dominating lotteries. We show that the ordinal core is always non-empty and that any ordinal core allocation can be implemented as a lottery over near-feasible and ex-post efficient outcomes. In fact, we prove that there is always a set of allocations in the core, which we call the *competitive core*, that can be supported by prices arising from a novel competitive equilibrium foundation. When endowments are deterministic, competitive core allocations can be implemented over near-feasible ex-post core outcomes. Moreover, competitive core allocations are ordinally envy-free and ex-post envy-free up to one good (where envy is only justified if another agent's endowment is either smaller or worth less in equilibrium). A mechanism that selects a competitive core allocation is strategyproof in the large. Our framework can be used in many real-world market design applications, such as organ exchanges, tuition exchanges, time bank sharing, shift exchanges, and resource reallocation.

*Department of Economics and Nuffield College, University of Oxford; simon.jantschgi@economics.ox.ac.uk.

†Daniels School of Business, Purdue University; nguye161@purdue.edu.

‡Department of Economics, Institute for New Economic Thinking, and St. Catherine's College, University of Oxford; alexander.teytelboym@economics.ox.ac.uk.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 949699).

1 Introduction

Combinatorial exchanges are market institutions involving the exchange of bundles of indivisible goods without monetary transfer among agents who enter the market with an initial endowment and ordinal information of preferences.[1] Combinatorial exchanges are by now ubiquitous. Examples of combinatorial exchanges include time banks, in which participants exchange service time slots (Andersson et al. 2021); tuition exchange programs, in which university employees exchange tuition benefits for their children (Dur and Ünver 2019); organ exchange programs in which patient-donor groups exchange organs based on compatibility (Roth et al. 2004); humanitarian aid allocations, where NPOs trade resources (Huang et al. 2023); and shift reallocations, where workers exchange shifts to accommodate their schedules (Manjunath and Westkamp 2021). Combinatorial exchanges are a generalization of combinatorial assignment problems (e.g., course allocation (Budish 2011, Budish and Cantillon 2012), allocation of food to food banks (Prendergast 2017, 2022), and refugee resettlement (Ahani et al. 2023, Delacrétaz et al. 2023)) because they allow participants to start with initial endowment. Indeed, in many settings, such as refugee resettlement (Delacrétaz et al. 2023), the possibility of an initial endowment can be a decision for the market designer.

While our understanding of combinatorial assignment (in which goods are owned by the designer) has progressed substantially in recent years, very little is known about properties of practical combinatorial exchange mechanisms. Beyond the unit-demand setting in which Gale's celebrated Top Trading Cycles (TTC) is a strategyproof mechanism that selects the unique outcome in the core (Roth and Postlewaite 1977, Shapley and Scarf 1974), most results on combinatorial exchanges emphasize the impossibility of a deterministic mechanism achieving desirable properties like strategyproofness, individual rationality, efficiency, and fairness simultaneously (Klaus and Miyagawa 2002, Pápai 2007, Sönmez 1999).

Random mechanisms provide a way to obtain positive results and are a key tool to achieve fairness in the allocation of indivisible goods. However, when a mechanism uses randomization, it is not enough to focus either on ex-post or on ex-ante properties. Naively randomizing over ex-post efficient mechanisms, as in a random serial dictatorship, might lead to ex-ante inefficiencies (Bogomolnaia and Moulin 2001), while individual lottery assignments might not be implementable as a lotteries over desirable economy-wide allocations. Hence, the designer must balance both ex-post and ex-ante desiderata.

[1]The restriction on monetary transfers often arises from ethical concerns about placing market values on certain goods or from fairness considerations (Roth 2007). Resource (re)allocation based on prices tends to benefit the wealthiest rather than those most in need, which can be seen as unfair in many cases.

In this paper, we study random mechanisms for combinatorial exchange and show that desirable ex-ante and ex-post efficiency, individual rationality and fairness properties can be achieved alongside large-market incentive properties. We consider the *ordinal core*, which naturally generalizes the standard notion of the deterministic core for random assignments without making any assumptions about the cardinal representation of preferences (Bogomolnaia and Moulin 2001).

We show that the ordinal core of a combinatorial exchange problem is always non-empty and can be realized as a lottery over near-feasible and ex-post efficient allocations. However, our contributions extend beyond merely establishing the existence of ordinal core allocations. Our key contribution is to define a new notion of competitive equilibrium that yields allocations in the ordinal core. We refer to these allocations as the *competitive core*, which inherits many desirable efficiency, fairness and incentive properties that one might expect from a competitive equilibrium allocation. Building on this foundation, we introduce our novel *Competitive Core Mechanism*, which harnesses these properties to offer a new solution for real-world market design challenges.

Concretely, we consider a general setting of combinatorial exchange in which traders have *weak* ordinal preferences over bundles, and are endowed with goods according to some, possibly *random*, initial allocation. Allowing preferences with ties (in the combinatorial setting) and random endowments opens up many important applications. Ordinal preferences with ties are desirable for many applications because they allow agents to submit preferences using criteria rather than the full ranking list, thus reducing the complexity of preference elicitation. Random endowments, on the other hand, naturally arise in practice, such as in housing allocations through lottery tickets (Abdulkadiroğlu and Sönmez 1998) or fractional allocations (Athanassoglou and Sethuraman 2011). This is especially relevant when there are more agents with entitlements than available goods or when a market designer assigns initial allocations based on fairness criteria.

We define the *ordinal core* as the set of feasible random allocations that satisfy the following conditions: (1) the allocation of each agent stochastically dominates their initial endowment (individually rational); and (2) it cannot be blocked by any subset of agents trading probability shares from their endowments, such that all involved traders prefer the new allocation in the sense of first-order stochastic dominance.

Our first contribution is to prove that the ordinal core is always non-empty, and that any ordinal core allocation is realizable as a lottery over near-feasible ex-post efficient allocations (Theorem 1). For deterministic endowments, the ordinal core that we construct is realizable as a lottery over near-feasible ex-post core allocations (Theorem 2). For unit demand, we get exact realizability. For multi-unit demand, the violation is upper-bounded by the maximum size of acceptable bundles minus one,

and thus independent of the number of agents or goods. The realization of the lottery assignment in the ordinal core would be done by the designer who can adjust supply of goods if necessary. Our adjustment bounds are realistic and can be used in many applications such as course allocation, student exchange, and refugee resettlement.

Moreover, we provide a novel competitive market foundation for a set of allocations in the ordinal core. The connection between the ordinal core and competitive equilibrium turns out to be insightful, but is non-trivial. First, due to the high degree of non-convexity, a straightforward notion of competitive equilibrium might not exist. Second, because of satiated preferences, a competitive equilibrium might not be efficient.

Our mechanism, therefore, involves two novel steps. First, we add a new good called artificial money and randomly relax an agent's budget by assigning them with an additional random fractional endowment. We then find a *approximate competitive equilibrium from random endowments* (A-CERE), which includes a price vector and a random optimal consumption, in which the market clears in expectation, up to a small error term concerning the random budget relaxation (Proposition 1). Each realization of the consumption corresponds to the optimal consumption for the budget obtained from each realization of the random endowment. Second, akin to classical results from general equilibrium theory, we show that any A-CERE is in the *approximate ordinal core* (Proposition 2). Third, we take that random budget perturbation to zero to guarantee exact market-clearing. We show that since each A-CERE is in the approximate ordinal core, the limit allocation is also in the ordinal core of the original economy (Proposition 3). We call this limit allocation, called the *competitive core* and its existence proves the non-emptiness of the ordinal core of the original economy.

Competitive core allocations have several desirable fairness properties. In combinatorial exchange problems with endowments, ensuring envy-freeness for all pairs of agents while maintaining individual rationality is generally impossible. With endowments (or priorities), however, it is often possible to eliminate *justified envy*, which in our case occurs when agents with "more valuable" endowments envy those with "less valuable" ones. Justified envy then requires defining what makes one endowment more or less "valuable" than another. Allocations in the competitive core eliminate two types of justified envy. First, in terms of goods, larger endowments are generally more valuable, and we demonstrate that agents with larger endowments do not envy those with smaller endowments. Second, our competitive market framework attaches prices to goods reflecting agents' rates of marginal substitution over probability shares. We show that, at equilibrium prices, agents with more valuable endowments do not envy those with less valuable ones because agents with higher budgets will not envy those with equal or lower budgets. We formally prove that our mechanism eliminates ex-ante justified

envy based on set-inclusion and equilibrium-valuations (Theorem 3), as well as ex-post justified envy up to one good (Theorem 4).

Finally, while our mechanism for selecting the competitive core is not incentive compatible (since in a finite market agents' reports can influence prices), it nevertheless has good asymptotic incentive properties. We show that our mechanism is strategyproof in the large (SP-L), a notion of asymptotic incentive compatibility introduced by Azevedo and Budish (2019): in a large-market limit economy where agents are price takers, it becomes a dominant strategy for agents to report their preferences truthfully. Azevedo and Budish (2019) argue that, across various practical domains, the mechanisms that are most successful are often those that are SP-L.

Related Work

Our paper contributes to a growing body of research on exchange economies involving discrete goods and their applications in market design. Aziz (2015) introduced and proved the non-emptiness of the ordinal core in the unit-demand exchange setting.[2] In this paper, we define and prove the non-emptiness of the ordinal core in the very general combinatorial exchange setting. The ordinal core is a natural extension of the core concept to random allocation settings and some of its variations have been explored in the literature (see Table 1). There are two key types of models of setting involving exchange of indivisible goods.

The first type of models considers ordinal preferences. The seminal works by Roth and Postlewaite (1977), Shapley and Scarf (1974) introduced the housing market model with unit demand and Gale's Top Trading Cycles (TTC) mechanism. TTC selects the unique outcome in the core, and is thus individually rational and efficient, as well as strategyproof (Roth 1982). Since then, a substantial body of literature has emerged, studying TTC, extending it, and applying it to real-world scenarios including house allocation (Abdulkadiroğlu and Sönmez 1999), kidney exchange (Roth et al. 2004), and school choice (Abdulkadiroğlu and Sönmez 2003).[3] However, as discussed in our introduction, many real-life applications are not unit demand. For the general multi-unit combinatorial exchange problem, mostly impossibility results are known: no deterministic mechanism can achieve efficiency, strategyproofness and individual rationality, see, e.g., (Klaus and Miyagawa 2002, Konishi et al. 2001, Pápai 2003, 2007, Sönmez 1999). In such multi-unit settings, most papers avoid these impossibilities

[2]The ordinal core has also be defined for probabilistic two-sided matching models (Doğan and Yıldız 2016, Manjunath 2013).

[3]For a survey article of the history of TTC in theory and practice, see Morrill and Roth (2024).

and obtain positive results by significantly restricting the domain of preferences (e.g., dichotomous, responsive, or lexicographic), see Table 1 for an overview. This limitation is neccessary, because the basic building block of the TTC mechanism is a one-to-one exchange, and expanding the preference domain requires substantially extending TTC or an entirely different approach.

The second type of models use cardinal preferences and apply notions of competitive equilibrium (CE) in a pseudo-market. In this framework, an efficient allocation is determined by applying competitive equilibrium principles to fictitious prices assigned to goods. An agent's budget is equivalent to the value of their endowment. A key advantage of CE is that it is in the core. However, guaranteeing the existence of CE typically requires convexity assumptions, which necessitates a shift from ordinal to cardinal preferences. This approach was originally developed by Hylland and Zeckhauser (1979), Varian (1974) for assignment problems and has been further studied for exchange models recently by Echenique et al. (2019, 2021, 2023), Garg et al. (2020).

However, a key technical challenge remains: With endowments, even in a unit-demand setting, a pseudo-market equilibrium might not exist (Hylland and Zeckhauser 1979). Hence, most of the literature has focused on restricted preference domains (e.g., additive) and opted for approximate equilibria, which results in outcomes that only approximately the core.

Nevertheless, our paper demonstrates a method to combine the advantages of both approaches. We assume general ordinal preferences without further domain-specific assumptions, yet we are still be able to use the competitive equilibrium framework. Our approach to applying competitive equilibrium with ordinal preferences by introducing a continuous perturbation to each agent's endowment and gradually reducing this perturbation to zero, is new to combinatorial exchange.[4] It allows us to show the existence of ordinal core by the limit of CEs. Our result avoids the traditional nonexistence issue by allowing randomization and relaxing feasibility for deterministic outcomes, similar to Budish (2011) and Nguyen et al. (2016).

The competitive equilibrium approach we follow is particularly useful to study fairness. In his seminal work, Varian (1974) proposed *competitive equilibrium from equal incomes* to find an envy-free allocation. However, in exchange models, it is impossible to rule out envy without violating property rights. Hence, in the presence of property rights or priorities, it has been argued that equity, instead of meaning envy-freeness, should be characterized by the absence of *justified envy*. Our paper offers two attractive notions of justified envy-freeness for both random allocation and ex-post deterministic

[4] A similar perturbation technique was employed by Nguyen et al. (2024) in the context of combinatorial assignment without endowments, but their analysis is simpler and does not involve taking limits.

allocations based on the size and value of agent's endowments.

Three recent papers study different notions of justified envy in pseudo-markets with endowments. Echenique et al. (2019) consider price-dependent incomes, similar to endowments, and study justified envy, for which a remedy would not violate any agent's participation constraint, a concept first introduced by Yılmaz (2010) for a housing market with existing tenants.[5] Garg et al. (2020) prove the existence of an approximate pseudo-market equilibrium that is approximately equal-type envy free. Echenique et al. (2023) consider balanced pseudo-market-equilibria equilibria that trade-off individual and aggregate endowments, and study justified envy based on market prices, which is closest to our definition of justified envy based on equilibrium valuations.

A related class of problems is combinatorial assignment, where agents do not initially own any goods. Similar to combinatorial exchange, approaches based on Probabilistic Serial (Bogomolnaia and Moulin 2001, Nguyen et al. 2016) and competitive equilibrium (Budish 2011, Budish et al. 2013, Gul et al. 2022, Kornbluth and Kushnir 2021, Miralles and Pycia 2021, Nguyen et al. 2024, Nguyen and Vohra 2024) have been extensively studied in this context.[6] Exchange problems, however, are significantly more complex than assignment problems.[7] While assignment problems primarily focus on efficiency, exchange problems necessitate two additional constraints: individual rationality and no-coalition blocking. As a result, methods developed for assignment problems often fall short when applied to exchange settings, while our Competitive Core mechanism is specifically developed for exchange markets. We hope that our mechanism has the potential to bring about significant improvements in exchange economies, much like the impact that the aforementioned mechanisms have had on combinatorial assignment.

[5]This notion views an assignment as unfair if an agent does not prefer their consumption to another agent's and swapping their consumptions respects the latter's individual rationality. Similar notions of no-justified envy with priorities appear in the literature on school choice, see, e.g. Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu et al. (2020).

[6]In a pure assignment setting, pseudo-markets are used where agents are assigned token budgets, which serve as their only source of income. Two notable extensions to the assignment environment include He et al. (2018), who incorporate priority structures, and Echenique et al. (2021), who introduce constraints into pseudo-markets with cardinal preferences.

[7]Inbetween the combinatorial assignment and exchange problems are assignment models with priorities or property rights, which allow to study individual rationality in assignment models. In the unit demand setting, Pápai (2000) introduced an extension of TTC, called hierarchical exchanges, which was later extended by Pycia and Ünver (2017) to mechanisms called Trading cycles. Furthermore, Balbuzanov and Kotowski (2019) introduced the exclusion core, an extension of the core where endowments are viewed as exclusion rights. Similarly, Abdulkadiroğlu and Sönmez (1999), Yılmaz (2010) study housing markets with existing tenants, where some agents do not have an endowment.

Mechanism	Pref.	Multi-Unit	Full Domain	Random endow.	Stability	Fairness	Strategy-proof
TTC	Ord.	No	Yes	No	ex-p. Core	n/a	Yes
reTTC	Ord.	No	Yes	Yes	ex-p. Core	ord. EF	Yes
feTTC	Ord.	No	Yes	Yes	ord. Core	n/a	No
CC	Ord.	No	Yes	Yes	ord. IR +Eff.	ord. JEF	No
STC	Ord.	Yes	No	No	ex-p. IR	n/a	Yes
cTTC	Ord.	Yes	No	No	ex-p. IR+Eff.	n/a	Yes
RS	Ord.	Yes	No	No	ex-p. IR+Eff.	n/a	Yes
aTTC	Ord.	Yes	No	No	ex-p. Core	n/a	No
RP	Ord.	Yes	No	No	ord. IR+Eff.	ord. EF	Yes
IRP	Ord.	Yes	No	No	ex-p. IR+Eff.	n/a	Yes
ϵ*HZEQ*	Card.	Yes	No	Yes	ϵ-ex-a. Core	ex-a. ETOE	Asymp.
bHZEQ	Card.	Yes	No	Yes	ϵ-ex-a. IR+Eff.	ex-a. JEF	Asymp.
***CCM* (this paper)**	Ord.	Yes	Yes	Yes	ord. Core	ord. JEF+ex-p. JEF1	SP-L

Table 1: **Summary of Existing Mechanisms and Their Properties.** For ordinal preferences, *Top Trading Cycles* **(TTC)** was introduced by Shapley and Scarf (1974) for unit demand settings and deterministic endowments. Abdulkadiroğlu and Sönmez (1998) extended TTC to random endowments **(reTTC)**, and Aziz (2015) extended TTC to fractional endowments **(feTTC)**. Beyond TTC mechanisms, in markets with unit demand and fractional endowments, Athanassoglou and Sethuraman (2011) introduce *Controlled Consumption* **(CS)**. For multi-type housing markets with strictly separable preferences Wako (2005) and Feng et al. (2024) study the componentwise extension **(cTTC)**. For multi-unit housing markets with monotone and responsive preferences, there are *Segmented Trading Cycles* **(STC)** (Pápai 2003), *circulation TTC* **(cTTC)** (Biró et al. 2022a), and *Random Serial rules* **(RS)** (Biró et al. 2022b). For multi-unit housing markets with conditionally lexicograpic preferences, Fujita et al. (2018) propose *augmented TTC* **(aTTC)**. For markets with multi-unit demand, Andersson et al. (2021) develop *Random Priority* **(RP)** for dichotomous preferences, while Manjunath and Westkamp (2021) design *Individually Rational Priority* **(IRP)** for trichotomous preferences. For linear cardinal utilities, Garg et al. (2020) prove the existence of an approximate pseudo-market equilibrium ϵ**HZEQ**, and Echenique et al. (2023) consider balanced pseudo-market-equilibria equilibria **(bHZEQ)** for continuous, monotone, and concave utilities. Our *Competitive Core Mechanism* **(CCM)** is the first to address combinatorial exchange with multi-unit demand, random endowments, and unrestricted ordinal preferences, while still providing strong guarantees of stability, fairness, and incentives.
Abbreviations. ex-p. = ex-post, ord. = ordinal, and ex-a. = ex-ante, IR=individual rationality, Eff. = efficiency, (J)EF(1) = (justified) envy freeness (up to one good), ETOE = equal treatment of equals, asymp. = asymptotically, SP-L = strategyproof in the large
Ranking of Properties. Core $\Rightarrow$ IR + Eff., EF $\Rightarrow$ JEF $\Rightarrow$ ETOT, and JEF $\Rightarrow$ JEF1. For stability, ex-a. $\Rightarrow$ ord. $\Rightarrow$ ex-p. For fairness, ex-p. $\Rightarrow$ ord. $\Rightarrow$ ex-a.

2 Example

Before we establish the general model, we illustrate our key solution concept, the *ordinal core*, using a well-known example. There are four agents, $\{1, 2, 3, 4\}$, and four goods, $\{a, b, c, d\}$. Agents have unit demand. Agents 1 and 2 prefer $a \succ b \succ c \succ d$, while agents 3 and 4 prefer $b \succ a \succ d \succ c$. Endowments are drawn uniformly at random. We call this allocation in Figure 1a the *random endowment*. For any realization of the random endowment, we can employ Gale's Top Trading Cycles (TTC) algorithm (Shapley and Scarf 1974) to find an allocation in the core. We call the resulting allocation the *TTC allocation from random endowment*, see Figure 1b. As shown by Abdulkadiroğlu and Sönmez (1998), this allocation is equivalent to the one from a random serial dictatorship. However, Bogomolnaia and Moulin (2001) showed that this allocation is ordinally inefficient since every agent strictly prefers their lottery in Figure 1c. That is, for all agents, probability mass is shifted from less preferred to more preferred goods. Indeed, the allocation in Figure 1c is ordinally efficient, meaning there does not exist another feasible random allocation in which each agent receives a first-order stochastically dominating lottery. Moreover, this allocation is in the *ordinal core*: it cannot be blocked by any subset of agents, i.e., they cannot feasibly exchange probability shares from their endowments such that all blocking traders weakly prefer the blocking allocation, with at least one strictly preferring it. Since the allocation in Figure 1b is not ordinally efficient, it is not in the ordinal core. Indeed, this example shows that the TTC from random endowments does not produce an outcome in the ordinal core.

(a) *Random Endowment*

	a	b	c	d
1	1/4	1/4	1/4	1/4
2	1/4	1/4	1/4	1/4
3	1/4	1/4	1/4	1/4
4	1/4	1/4	1/4	1/4

(b) *Random Endowment + TTC*

	a	b	c	d
1	5/12	1/12	5/12	1/12
2	5/12	1/12	5/12	1/12
3	1/12	5/12	1/12	5/12
4	1/12	5/12	1/12	5/12

(c) *Ordinal Core Allocation*

	a	b	c	d
1	1/2	0	1/2	0
2	1/2	0	1/2	0
3	0	1/2	0	1/2
4	0	1/2	0	1/2

Figure 1: **Three Random Allocations.** The *random endowment* (Left) is dominated by the *TTC allocation from random endowment* (Middle), which can be realized as a lottery over ex-post efficient allocations. However, it is dominated by the *ordinal core allocation* (Right).

Note that the allocation in Figure 1c is the solution of the *Probabilistic Serial mechanism* developed by (Bogomolnaia and Moulin 2001) for the assignment setting.

While this mechanism always results in an ordinally efficient allocation in the assignment setting, it may not be individually rational for some endowments and might fail to produce an ordinal core allocation, as it neglects initial endowments. In the case of unit demand, and fractional endowments, Aziz (2015)'s *fractional TTC* finds an ordinal core allocation.

In the rest of the paper, we show that it is always possible to find an allocation in the ordinal core for the general combinatorial exchange problem.

3 The Combinatorial Exchange Problem

There is a finite set N of *agents* with $|N| = n$, and a finite set M of *goods* with $|M| = m$. Each good j has a finite *capacity* $c^j \in \mathbb{N}^{>0}$.

A *bundle* is an integral vector $\mathbf{x} \in \mathbb{N}^m$. The bundle consumed by agent i is denoted by $\mathbf{x}_i$. There might be further constraints on consumption; let $\Psi_i \subset \mathbb{N}^m$ denote the set of *acceptable bundles* for agent i. We do not assume free disposal. If a bundle $\mathbf{x}_i$ is acceptable to agent i, it is possible that $\mathbf{x}_i^{'} \leq \mathbf{x}_i$ is not acceptable. We assume that the maximum size of an acceptable bundle for any agent is at most Δ.[8] A *deterministic allocation* $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is a list of acceptable bundles, one for each agent. We allow random allocations. Let $\mathcal{L}(\Psi_i)$ denote the set of all lotteries over the set of acceptable bundles Ψ_i. Let $\tilde{\mathbf{x}}_i \in \mathcal{L}(\Psi_i)$ denote a lottery obtained by agent i and $\mathbb{E}[\tilde{\mathbf{x}}_i]$ the expectation of this lottery. A *random allocation*, $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n) \in \mathcal{L}(\Psi_1) \times ... \times \mathcal{L}(\Psi_n)$, is a list of lotteries over the acceptable bundles, one for each agent.

Each agent i enters the market with an *endowment* $\tilde{\mathbf{e}}_i \in \mathcal{L}(\Psi_i)$, meaning they initially owns a lottery over their acceptable bundles. Let $\tilde{\mathbf{e}} = (\tilde{\mathbf{e}}_1, ..., \tilde{\mathbf{e}}_n) \in \mathcal{L}(\Psi_1) \times ... \times \mathcal{L}(\Psi_n)$ be the vector of all endowment. A lottery with probability one on a bundle corresponds to the classic case of *deterministic endowments*. We assume that endowments are *feasible* with respect to capacities, that is, $\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{e}}_i] = \mathbf{c}$.

Each agent i has weak ordinal preferences $\succeq_i$ over their set of acceptable bundles Ψ_i. We denote the preferences of all agents by $\succeq := (\succeq_i)_{i \in N}$. If preferences are strict, we write $\succ_i$ instead of $\succeq_i$. If agent i is indifferent between bundles, that is $\mathbf{x}_i \succeq_i \mathbf{y}_i$ and $\mathbf{y}_i \succeq_i \mathbf{x}_i$, we write $\mathbf{x}_i \sim_i \mathbf{y}_i$. We use $\succeq_{-i}$ to denote the preference profile excluding agent i. For random allocations, we use the partial order over lotteries, which is induced by first-order stochastic dominance. Agent i prefers the random allocation $\tilde{\mathbf{x}}$ to $\tilde{\mathbf{y}}$, if

[8] In many real-world applications, such as organ transplantation, complementary time slots, and refugee services, the maximum value of Δ varies. While some scenarios require strictly unit demand, others may not. However, for most applications, 6 is typically an upper bound.

$\tilde{\mathbf{x}}_i$ (first-order) stochastically dominates $\tilde{\mathbf{y}}_i$.[9] A lottery $\tilde{\mathbf{x}}_i$ stochastically dominates $\tilde{\mathbf{y}}_i$, denote by $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{y}}_i$, if for every acceptable bundle $\mathbf{z}$ it holds that $\mathbb{P}(\tilde{\mathbf{x}} \succeq_i \mathbf{z}) \geq \mathbb{P}(\tilde{\mathbf{y}} \succeq_i \mathbf{z})$. We say $\tilde{\mathbf{x}}_i$ strictly stochastically dominates $\tilde{\mathbf{y}}_i$, denote by $\tilde{\mathbf{x}}_i \succ_i^{sd} \tilde{\mathbf{y}}_i$, if for at least one bundle $\mathbf{z}$, the inequality is strict.

An *economy* is defined as $(\tilde{\mathbf{e}}, \succeq)$, which fully specifies the tuple $(N, M, \mathbf{c}, \tilde{\mathbf{e}}, \Psi, \succeq)$.

4 The Ordinal Core

Allocations should adhere to capacity constraints, ensuring that the quantity of goods allocated does not exceed the available supply in the economy.

Definition 1 (Feasibility)**.** A deterministic allocation $\mathbf{X}$ is *feasible* with respect to capacities $\mathbf{c}$ if $\sum_{i \in N} \mathbf{x}_i \leq \mathbf{c}$. A random allocation $(\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *feasible* with respect to capacities $\mathbf{c}$ if $\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i] \leq \mathbf{c}$. An allocation (deterministic or random) is Δ-*near feasible* with respect to capacities $\mathbf{c}$, if there exists $\mathbf{c}'$ with $\mathbf{c}' \leq \mathbf{c} + \Delta \cdot \mathbf{1}$ such that the allocation is feasible with respect to $\mathbf{c}'$.

A random allocation specifies a lottery for each agent but does not require that they can be simultaneously realized as a lottery over feasible allocations.

Definition 2 (Realizability)**.** A random allocation $\tilde{\mathbf{X}}$ is *realizable* over a set of deterministic allocations if it can be implemented as a lottery over that set.

As agents come to the market with an endowment, we need to guarantee that they weakly prefer their allocation to their endowment to guarantee participation.

Definition 3 (Individual Rationality)**.** For deterministic endowments $\mathbf{e} = (\mathbf{e}_1, ..., \mathbf{e}_n)$, a deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is *individually rational*, if for every agent i, it holds that $\mathbf{x}_i \succeq_i \tilde{\mathbf{e}}_i$. For random endowments $\tilde{\mathbf{e}} = (\tilde{\mathbf{e}}_1, ..., \tilde{\mathbf{e}}_n)$, a random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *individually rational*, if for every agent i, it holds that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{e}}_i$.

A key market design criterion for an allocation is *efficiency*, meaning no other allocation weakly improves all agents and strictly improves at least one.

Definition 4 (Efficiency)**.** A feasible deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is *efficient* with respect to capacities $\mathbf{c}$, if there is no other feasible allocation $\mathbf{Y} = (\mathbf{y}_1, ..., \mathbf{y}_n)$ such that for all agents i $\mathbf{y}_i \succeq_i \mathbf{x}_i$, and for at least one agent j, the preference is strict, that is, $\mathbf{y}_j \succ_j \tilde{\mathbf{x}}_j$. A deterministic allocation is Δ-*near efficient* with respect to capacities

[9] This definition is equivalent to stating that agent i would prefer a particular lottery for all von Neumann-Morgenstern utility functions consistent with their ordinal preferences.

$\mathbf{c}$ if it is efficient with respect to capacities $\mathbf{c}'$ with $\mathbf{c}' \leq \mathbf{c} + \Delta \cdot \mathbf{1}$. A feasible random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *ordinally efficient* with respect to capacities $\mathbf{c}$ if there is no other feasible random allocation $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_1, ..., \tilde{\mathbf{y}}_n)$ such that for all agents i $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i$, and for at least one agent j, the preference is strict, that is, $\tilde{\mathbf{y}}_j \succ_j^{sd} \tilde{\mathbf{x}}_j$.

In an exchange economy, individual rationality and efficiency are special cases of *core stability* (Shapley and Scarf 1974), which ensures that no coalition has an incentive to block a proposed allocation. A coalition's partial reallocation must be feasible within their own endowments. Individual rationality means no single member forms a blocking coalition, while efficiency means the grand coalition is not blocking.

Definition 5 (Deterministic Core)**.** For deterministic endowments $\mathbf{e} = (\mathbf{e}_1, ..., \mathbf{e}_n)$, and a deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$, a *blocking coalition* is a subset of traders $S \subseteq N$ and a partial allocation $\mathbf{Y} = (\mathbf{y}_i)_{i \in S}$ for traders in S such that $\sum_{i \in S} \mathbf{y}_i \leq \sum_{i \in S} \mathbf{e}_i$, and $\mathbf{y}_i \succeq_i \mathbf{x}_i$ for all $i \in S$, and $\mathbf{y}_i \succ_i \mathbf{x}_i$ for at least an agent $i \in S$. A deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is in the *core*, if it is feasible with respect to the total endowment, and has no blocking coalitions. A deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is in the Δ-*near feasible core*, if it is Δ-near feasible with respect to the total endowment, $\sum_i \mathbf{e}_i$, and has no blocking coalitions.

We extend the definition of the core to random endowments and allocations. We do so by extending the concept of *ordinal efficiency* (Bogomolnaia and Moulin 2001) to the *ordinal core*, by considering blocking coalitions with respect to stochastic dominance.

Definition 6 (The Ordinal Core)**.** For a random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$, a *ordinal blocking coalition* is a subset of traders $S \subseteq N$ and a partial allocation $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_i)_{i \in S}$ for traders in S such that: $\sum_{i \in S} \mathbb{E}[\tilde{\mathbf{y}}_i] \leq \sum_{i \in S} \mathbb{E}[\tilde{\mathbf{e}}_i]$, $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i$ for all $i \in S$, and $\tilde{\mathbf{y}}_i \succ_i^{sd} \tilde{\mathbf{x}}_i$ for at least one agent $i \in S$. A random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is in the *ordinal core*, if it is feasible with respect to the total endowment, $\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{e}}_i]$, individually rational and has no blocking coalitions.

For deterministic allocations, individual rationality follows as coalitions of size 1 are not blocking. For random allocations, this means an agent's random allocation is not stochastically dominated by their endowment. Because stochastic dominance is a partial order, no-blocking of size 1 is weaker than individual rationality. Thus, in the definition of ordinal core, we impose the stronger condition that a trader's random allocation weakly dominates their endowment.

Existence and Realizability

We can now directly state our first main results: The non-emptiness of the ordinal core, and its approximate realizability. Recall that Δ is the maximum bundle size.

Theorem 1. *For any economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$, the ordinal core is non-empty. Moreover, any ordinal core allocation is realizable over $(\Delta - 1)$-near feasible and efficient deterministic allocations.*

The existence-proof of Theorem 1 builds on a novel competitive market foundation, which will be developed in detail in Section 5. In particular, the existence will follow directly as a corollary from Proposition 3. The realizability result relies on tools from the approximation of random allocations (Nguyen et al. 2016). The full details of the proof can be found in Section 5.2.

We note that the bound in this theorem is tight. For example, when $\Delta = 1$, that is, agents only have unit demand, we recover the Birkhoff-von Neumann theorem, and get that every feasible random allocation is realisable with respect to original capacities.

Deterministic Endowments. Studying the ex-post core involves comparing an ex-post endowment with an ex-post allocation. Since we study the ordinal core, where agents trade probability shares, random endowments are not realized before the final allocation. Therefore, in this subsection, we focus on deterministic endowments, arguably the most important special case of our general model.

For deterministic endowments, the following theorem strengthens Theorem 1 and proves that there exists a random allocation in the ordinal core that is realizable over near-feasible deterministic core-allocations. This also implicitly proves the existence of a near-feasible deterministic core-allocation.

Theorem 2. *For any economy $\mathcal{E} = (\mathbf{e}, \succeq)$ with deterministic endowments, there exists an ordinal core allocation that can be realized over $(\Delta - 1)$-near feasible core allocations.*

The proof of Theorem 2 again builds on the competitive market foundation developed in Section 5. In particular, in Section 6, we provide a mechanism that selects a 'competitive' ordinal core allocation, which will be allocations that satisfy Theorem 2. Theorem 2 will follow directly from Theorem 4, where we prove the realizability of the competitive core over near-feasible deterministic core allocations, along with additional fairness guarantees.

5 Competitive Market Foundation

To show that the ordinal core is non-empty, we use an established economic insight: the connection between the core and competitive equilibrium. In general equilibrium theory with perfectly divisible goods, any competitive equilibrium is known to be in the core.[10]

However, in our market environment with indivisible goods and non-convex preferences, several issues arise. First, a competitive equilibrium typically does not exist.

Example 1 (Non-Existence of CE)**.** There are three agents, $\{1, 2, 3\}$, and three goods, $\{a, b, c\}$. Agent 1 owns one unit each of goods a and b and wants to consume as much of good c as possible. Agents 2 and 3 each own one unit of good c and prefer $a \succ b \succ c$. A competitive equilibrium (CE) cannot exist because if the price of good c is at least that of good a, both agents 2 and 3 will demand good a, exceeding its supply. Conversely, if the price of a exceeds that of c, no one can afford a, reducing its price to zero, which makes all prices zero, violating the conditions for a CE.

Key Idea 1 (Random Budget Relaxation). To address this, we introduce random budget relaxations, ensuring agents 2 and 3 have different budgets. This allows setting prices so each agent can afford good a half of the time and good b the other half. Symmetric perturbations ensure that the allocation is fair ex-ante.

However, the next example illustrates a nuanced issue with efficiency that arises due to the non-satiation of budgets in the presence indivisible goods.

Example 2 (Inefficiency due to Non-Satiation)**.** There are three agents, $1, 2, 3$, and a single good g with a capacity of $\mathbf{c}^g = 6$. Each agent deterministically owns two units of the good. Agent 1 is indifferent between consuming one or two units. Agents 2 and 3 can consume either two or three units and strictly prefer the latter. In a core allocation, agent 1 would consume one unit, while either agent 2 or 3 consumes three units, and the other consumes two. A 'fair' ordinal core allocation would flip a coin to randomize between these options. However, these allocations do not align with any competitive equilibrium. Regardless of the price $p^g > 0$ for good g, even with small random budget perturbations, agents 2 and 3 cannot afford three units, so their random demand is almost surely two units. Agent 1, owning two units, can always choose their preferred bundle of one unit. Thus, at any price $p^g > 0$, the allocation resulting from random demand and budget relaxation is one unit for agent 1 and two units each for agents 2 and 3, which is not ordinally efficient.

[10] See, e.g., Aumann (1964), Debreu and Scarf (1963), Hildenbrand (1982).

Key Idea 2 (Artificial Currency). To resolve this issue, we introduce an artificial currency, called money, with each agent receiving the same endowment. Preferences are refined so that agents break ties based on the amount of money they have. Agent 1 can exchange the extra unit of good g for money, while agents 2 and 3 can exchange money for the additional unit of good g. In equilibrium, the price of the artificial currency is set so that, with random budget relaxations, agents 2 and 3 can afford the additional unit of the good, ensuring the probabilities sum to one and the good is always allocated. With fair budget relaxations, agents 2 and 3 receive two units with probability 0.5 and three units with probability 0.5.

Outline. In our exchange economy, we introduce an artificial currency, called money, and consider small random relaxations of agents' budget constraints. Money ensures that, in cases of indifference, agents select their favorite expenditure-minimizing bundle and fully use up their budgets. Preferences are extended so that agents break ties based on the amount of money they possess, without altering their fundamental ordering of bundles. Small random budget relaxations help resolve ties among agents with identical endowments. We use budget relaxations instead of perturbations to guarantee individual rationality, allowing agents to always repurchase their own endowments.

In Section 5.1, we introduce a novel solution concept called *approximate competitive equilibrium from random endowments*. Here, given prices for goods, agents consume according to their random demand, based on their endowments and budget relaxations, and the market approximately clears for all goods in expectation. In Section 5.1.1, we demonstrate that for continuous random budget relaxations, random demand is upper hemi-continuous. This allows us to use standard fixed-point arguments to guarantee the existence of an approximate equilibrium (Proposition 1), where the market-clearing error is bounded by the budget relaxation size δ. In Section 5.1.2, we use traditional economic reasoning, similar to the classical case of perfectly divisible goods, to show that any equilibrium allocation is in the δ-ordinal core (Proposition 2).

Finally, in Section 5.2, we achieve a feasible random allocation by reducing the size of the random budget relaxation δ to zero. The resulting limit allocation, called the *asymptotic competitive equilibrium from random endowments*, exists and will always be in the core and will be exactly feasible (Proposition 3). This proves the non-emptiness of the ordinal core (Theorem 1). Moreover, this approach will yield a limit price vector that, while not representing a competitive equilibrium, corresponds to the marginal rates of substitution for the corresponding ordinal core allocation.

5.1 Approximate Competitive Equilibrium from Random Endowments

Artificial Currency. First, we introduce an auxiliary perfectly divisible good called *money*. We augment deterministic allocations $\mathbf{x}_i = (\mathbf{x}_i^G, \mathbf{x}_i^M) \in \Psi_i \times \mathbb{R}_{\geq 0}$, where $\mathbf{x}_i^M$ denotes the allocation of money. Thus, random allocations $\mathbf{x}_i$ are extended from lotteries over bundles to probability distributions over $\Psi_i \times \mathbb{R}_{\geq 0}$. There is a fixed amount of money in the economy, denoted by the capacity $\mathbf{c}^M > 0$. Initially, money is distributed evenly among agents: Hence, each agent's endowment is a probability distribution over $\Psi_i \times \mathbb{R}_{\geq 0}$, such that the endowments are feasible with respect to $\mathbf{c}^M > 0$, and that the marginal distribution over Ψ corresponds to an agent's initial lottery over bundles. For simplicity, we denote the extended endowment again by $\tilde{\mathbf{e}}_i$. We extend the agent's original preferences over bundles to preferences over bundles and money as follows: Agent i strictly prefers (x, m_1) over (y, m_2), write $(x, m_1) \succ_i (y, m_2)$ if either $x \succ_i y$ or $x \sim_i y$ and $m_1 > m_2$. Agent i is indifferent between (x, m_1) and (y, m_2), write $(x, m_1) \sim_i (y, m_2)$ if $x \sim_i y$ and $m_1 = m_2$. That is, no consumption of money can change the preference ordering over acceptable bundles, and given bundles, for which an agent is indifferent, they strictly prefer to have more money.

Budget Relaxation. We consider a random vector $(\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n)$ uniformly distributed over the n-dimensional δ-simplex, such that $\sum_{i=1}^n \tilde{\mathbf{b}}_i = \delta$ almost surely and $\tilde{\mathbf{b}}_i \geq 0$ for all i.[11] Let $\mathbf{b}_i$ denote a realization of $\tilde{\mathbf{b}}_i$. We call $\mathbf{b}_i$ the *random budget relaxation* for agent i.

Random Demand. Next, consider a *price vector* $\mathbf{p} = (\mathbf{p}^G, \mathbf{p}^M) \in \mathbf{R}_+^{m+1}$ that specifies a non-negative price for each good. We construct the *random demand of agent* i, denoted by $\chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i)$, as follows: First, agents receive a realized endowment from their initial probability distribution of goods and money, that is $\mathbf{e}_i \sim \tilde{\mathbf{e}}_i$. Second, their budget is randomly relaxed by $\mathbf{b}_i \sim \tilde{\mathbf{b}}_i$ times the value of the entire economy. Third, given their realized random budget, agents demand their favourite affordable bundles. Note that, given the extended preferences, agent i demands their favourite affordable bundle of goods $\mathbf{x}_i^G \in \Psi_i$, and the amount of money $\mathbf{x}_i^M \in \mathbb{R}_+$ to fully satiate their budget.[12] If an agent is indifferent between multiple affordable bundles, they will choose the bundle that minimizes their spending due to their preference for more money.

[11] For $\delta = 1$, this distribution is also known as the flat Dirichlet distribution.

[12] If the price of money is zero and agents don't spend their entire budget on their preferred bundle of goods, we allow for infinite money consumption, see Appendix A.1.

Therefore, random demand can only be set-valued if an agent is indifferent between bundles with the exact same price, and all strictly preferred bundles are unaffordable within their budget. More formally, random demand is the following random variable:

$$\chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) := \max_{\succeq_i} \left\{ \mathbf{x}_i = (\mathbf{x}_i^G, \mathbf{x}_i^M) \in \Psi_i \times \mathbb{R}_{\geq 0} \text{ and } \mathbf{p}\cdot\mathbf{x} \leq \mathbf{p}\cdot\mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}\cdot\mathbf{c} | \mathbf{e}_i \sim \tilde{\mathbf{e}}_i \wedge \mathbf{b}_i \sim \tilde{\mathbf{b}}_i \right\}.$$

Denote agent i's *expected demand* by $\mathbb{E}[\chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i]$, where the expectation is over the realization of the initial endowment and the random budget relaxation.

Next, we introduce a novel version of competitive equilibrium in this environment: When agents consume their random demand given their endowments and budget relaxation, the market for all goods clear in expectation (up to a constant δ).

Definition 7 (Approximate Competitive Equilibrium from Random Endowments)**.** Given an economy $(\tilde{\mathbf{e}}, \succeq)$, and random budget relaxations $(\tilde{\mathbf{b}}_1, ..., \tilde{\mathbf{b}}_n)$, a price vector $\mathbf{p} = (\mathbf{p}^G, \mathbf{p}^M)$ and the random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ comprise a *δ-approximate competitive equilibrium from random endowments (δ-CERE)* if

1. $\tilde{\mathbf{x}}_i \in \chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i)$, and
2. $\sum_{i=1}^n \mathbb{E}[\tilde{\mathbf{x}}_i]^j \leq (1+\delta)\mathbf{c}^j$ for every good j, and $\sum_{i=1}^n \mathbb{E}[\tilde{\mathbf{x}}_i]^j = (1+\delta)\mathbf{c}^j$, if $p^j > 0$.

For $\delta = 0$, we call the price vector and the allocation a *competitive equilibrium from random endowments (CERE)*. Note that a probability distribution over bundles and money induces a lottery over bundles, by considering the marginal distribution. We call this allocation the corresponding *δ-CERE allocation*.

5.1.1 Existence

In our non-convex market with indivisible goods and ordinal preferences, achieving a competitive equilibrium from random endowments without random budget relaxations is often impossible. For instance, with deterministic initial endowments, demand is not upper hemicontinuous, as small price changes can abruptly make bundles affordable. However, our next theorem establishes that for any $\delta > 0$, there exists a δ-CERE given that random budget relaxations are continuous. The key insight is that continuous random budget relaxations smooth out random demand in the space of random allocations, allowing us to leverage a standard fixed point argument. This result is significant as existence theorems in non-convex environments are rare, and it marks a novel contribution in the context of random allocations. It is important to note that

smoothing random demand through continuous budget relaxations incurs a trade-off by violating exact feasibility, as a δ-CERE is only δ-feasible.[13]

Proposition 1. *Consider any economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$. For every $\delta > 0$, a δ-CERE exists.*

Proof Sketch (Proposition 1). The first observation is that with continuous budget relaxations, random demand is upper hemicontinuous within the space of probability distributions over money and bundles of goods.[14] This allows us to define an upper hemicontinuous correspondence on the space of prices and allocations, where a 'price player' adjusts prices by increasing them for overdemanded goods and decreasing them for underdemanded ones, thereby maximizing the value of excess demand.[15] We then apply the Kakutani-Glicksberg-Fan fixed point theorem to show that this correspondence has a fixed point. We then show that this fixed point corresponds to a δ-CERE. The full proof is relegated to Appendix A.1.

5.1.2 Welfare Theorems

The fundamental theorems of welfare economics are a cornerstone of general equilibrium theory with divisible goods: Any competitive equilibrium is efficient (First Welfare Theorem) and any efficient allocation can be supported as a competitive equilibrium for some endowments (Second Welfare Theorem). Additionally, any competitive equilibrium is in the core. We extend these results to our market with indivisible goods and random allocations. As a δ-CERE is not feasible, we introduce a relaxation of the ordinal core, called the *δ-ordinal core,* which considers approximately feasible blocking coalitions.

Definition 8 (The δ-Ordinal Core)**.** A *δ-ordinal blocking coalition* is a subset of traders $S \subseteq N$ and a partial allocation $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_i)_{i \in S}$ for traders in S such that: $\sum_{i \in S} \mathbb{E}[\tilde{\mathbf{y}}_i] \leq (1+\delta) \sum_{i \in S} \mathbb{E}[\tilde{\mathbf{e}}_i]$, $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i$ for all $i \in S$, and $\tilde{\mathbf{y}}_i \succ_i^{sd} \tilde{\mathbf{x}}_i$ for at least one agent $i \in S$. A random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is in the *δ-ordinal core*, if $\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i] \leq (1+\delta)\mathbf{c}$, it is individually rational, and has no δ-blocking coalitions.

We can now formulate our welfare theorems as follows:

Proposition 2. *Consider any economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$. For any $\delta \geq 0$, any δ-CERE allocation is in the δ-ordinal core. Conversely, any allocation in the δ-ordinal core can be supported as a δ-CERE allocation for some initial δ-feasible endowments.*

[13] Random budget perturbations, which allow budget restrictions, could achieve exact feasibility but at the cost of losing individual rationality, as agents might no longer be able to afford their own endowment.

[14] Allowing infinite money consumption at price zero requires compactification to apply a fixed-point theorem, but in equilibrium, the price of money must be positive.

[15] A similar version of a price player was already used in Arrow and Debreu (1954).

Proof Sketch (Proposition 2). For an allocation to be supported by a δ-CERE, a price vector $\mathbf{p}^G$ must exist such that for every agent i, if they consume a bundle $\mathbf{x}$ with positive probability, any more preferred bundle $\mathbf{y} \succ_i \mathbf{x}$ must have a strictly higher cost. Otherwise, agent i would prefer $\mathbf{y}$. Conversely, if such a price vector exists, initial endowments can be redistributed so that each random allocation aligns with the optimal consumption under that endowment. δ-CERE prices are scale-invariant and can be scaled such that the cost difference between bundles $\mathbf{y}$ and $\mathbf{x}$ is at least 1. Thus, a random allocation can be sustained by a δ-CERE if and only if there is a price vector $\mathbf{p}^G \geq 0$ such that $\mathbf{p}^G \cdot (\mathbf{y} - \mathbf{x}) \geq 1$ for every $\mathbf{y} \succ_i \mathbf{x}$ and $\mathbb{P}(\tilde{\mathbf{x}}_i = \mathbf{x}) > 0$. This can be formulated as a linear program, and Farkas' lemma provides a characterization for the existence of a solution through the dual of the linear program. We show that the dual condition is satisfied if and only if there are no blocking coalitions. Moreover, any δ-CERE is individually rational, as agents can always afford to consume their own endowment, and thus only shift probability mass to more preferred bundles. The full proof is relegated to Appendix A.2.

5.2 Asymptotic Competitive Equilibrium from Random Endowments (Proof of Theorem 1).

The existence of δ-CEREs, as established in Proposition 1, comes with the caveat that these equilibria are only δ-feasible. In this subsection, we consider limit allocations that are constructed by letting the size of the budget relaxation δ go to zero.

We construct limits using the standard product topology: The price simplex is a metric space as a subspace of $\mathbb{R}^m$. The space of lotteries over k bundles is also a subspace of $\mathbb{R}^k$.

Definition 9 (Asymptotic Competitive Equilibrium from Random Endowments)**.** Given an economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$, a random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ and price vector $\mathbf{p} = (\mathbf{p}^G, \mathbf{p}^M)$ are an *asymptotic competitive equilibrium from random endowments*, if there exists a sequence $(\delta_k)_{k\geq 1}$ with $\lim_{k\to\infty} \delta_k = 0$, such that $(\tilde{\mathbf{X}}, \mathbf{p})$ is the limit of a sequence of δ_k-CEREs $(\tilde{\mathbf{X}}^{\delta_k}, \mathbf{p}^{\delta_k})$. We call the marginal distribution over bundles the corresponding *asymptotic CERE allocation*.

The following theorem demonstrates that such a limit allocation exists, it is exactly feasible, and, like δ-CEREs, it is in the ordinal core.

Proposition 3. *Consider any economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$. There exists an asymptotic CERE. Moreover, any asymptotic CERE allocation is in the ordinal core.*

Proof Sketch (Proposition 3). First, we show that the space of prices and allocations is sequentially compact, ensuring any sequence has a convergent subsequence, which proves the existence of a limit. Along the sequence, each allocation is δ_k-feasible: $\sum_{i\in N}\mathbb{E}[\tilde{\mathbf{x}}_i^{\delta_k}] \leq (1+\delta_k)\mathbf{c}$. Using a standard continuity and compactness argument, we can interchange limits and expectations $\sum_{i\in N}\mathbb{E}[\lim_{k\to\infty}\tilde{\mathbf{x}}_i^{\delta_k}] = \lim_{k\to\infty}\sum_{i\in N}\mathbb{E}[\tilde{\mathbf{x}}_i^{\delta_k}]$ to prove feasibility. Next, assume, for contradiction, that there exists a blocking coalition $S \subset N$ with a feasible partial allocation $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_i)_{i\in S}$ such that $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i^{\infty}$ for all $i \in S$ (strict for at least one agent). Stochastic dominance can be expressed in terms of cumulative distribution functions (CDFs): $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i$ if and only if $F_{\tilde{\mathbf{y}}_i}(\mathbf{z}) \leq F_{\tilde{\mathbf{x}}_i}(\mathbf{z})$ for all $\mathbf{z}$. Convergence in the lottery space implies pointwise convergence of these CDFs. If the limit allocation were stochastically dominated by a blocking coalition's allocation, then by continuity, one could construct a δ_k-ordinal blocking coalition for sufficiently small δ_k and $\tilde{\mathbf{X}}^{\delta_k}$. This contradicts the fact that $\tilde{\mathbf{X}}^{\delta_k}$ is in the ordinal δ_k-core along the subsequence. Hence, the limit allocation must be in the ordinal core. Individual rationality follows similarly. The full proof is relegated to Appendix A.1.

Proof of Theorem 1. Proposition 3 implies the first part of our main result in Theorem 1, the non-emptiness of the ordinal core, as any asymptotic CERE allocation is in the ordinal core. Thus, to finish the proof of Theorem 1, we need to show that any random allocation in the ordinal core is realizable over $(\Delta-1)$-near feasible deterministic efficient allocations. However, as any random allocation in the ordinal core is ordinally efficient, this results follows directly from Nguyen et al. (2024, Theorem 4.4):

> If the maximum bundle is of size at most Δ, then any ordinally efficient random allocation is realizable over $(\Delta-1)$-near feasible efficient allocations.

6 The Competitive Core

In light of the competitive market foundations developed in Section 5, we introduce the following terminology for asymptotic CERE allocations that we will use thoughout the rest of the paper.

Definition 10 (The Competitive Core)**.** Given an economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$, a random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is a *competitive core allocation*, if there exists a price vector $\mathbf{p}$, such that $\tilde{\mathbf{X}}$ is an asymptotic CERE allocation with associated price vector $\mathbf{p}$.

A competitive core allocation is therefore an allocation in the ordinal core that can be supported by asymptotic CERE prices. It turns out that competitive core allocation have several desirable properties.

6.1 Ordinal envy-freeness guarantees

For random (re)allocation mechanisms, standard measures of fairness are concerned with ex-ante envy, that is, agents comparing their assigned lotteries. In our model with ordinal preferences, we study *ordinal envy*, in which agents compare lotteries using first-order stochastic domination.

In terms of ordinal envy, the strongest notion of fairness is *ordinal envy-freeness* (Bogomolnaia and Moulin 2001), that is, no agent prefers the allocation of another agent to their own allocation.

Definition 11 (Envy-Freeness)**.** A random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *envy-free* if for all pairs of agents i, j, it holds that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$.

While this property is desirable in an assignment economy, it is too strong for an exchange economy, as envy-freeness is incompatible with individual rationality.

Example 3 (Unjustified Envy)**.** There are two agents 1 and 2 and two goods a and b. Agent 1 owns good a and agents two owns good b. Both agents prefer good a over good b. The only envy-free and efficient random allocation is that both agents get good a half the time, and good b half the time. However, this allocation is not individually rational for agent 1, as they are endowed with good a initially.

In this example, agent 2's envy should be viewed as *unjustified*, as agent 1 already owned the most preferred bundle. Thus, in an exchange setting, it becomes crucial to consider, when envy is justifiable, and when it is not. A bare-minimum fairness requirement for a random allocation is *equal treatment of equals*. That is, no agent envies the lottery of another agent, who comes to the market with the exact same endowment.

Definition 12 (Equal Treatment of Equals)**.** A random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ respects *equal treatment of equals*, if for any pair of agents i, j with the same endowment $\tilde{\mathbf{e}}_i = \tilde{\mathbf{e}}_j$, it holds that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_{i'}$.

Note that equal treatment of equals only allows to compare agents with the exact same endowment, and thus only yield very limited comparability.

However, there are other natural cases, in which it is clear that an agent should not envy another agent. In general, as we consider goods, a strictly bigger endowment is more valuable. Hence, an agent with a strictly bigger endowment should not envy another agent with a strictly smaller endowment.

We formalize this via set-inclusion: We define the following partial order $\geq$ on deterministic bundles $\mathbf{x}, \mathbf{y} \in \mathbb{N}^m$: $\mathbf{x} \geq \mathbf{y} \iff x^j \geq y^j \quad \forall j \in \{1, \ldots, m\}$, and

$\mathbf{x} > \mathbf{y} \iff \mathbf{x} \geq \mathbf{y}$ and $\exists j \in \{1, \ldots, m\}$ such that $x^j > y^j$. We extend this idea to random allocations: For a random allocation $\tilde{\mathbf{x}}$, let $\mathbb{P}(\tilde{\mathbf{x}} \geq \mathbf{z})$ be the probability that the lottery results in a bundle that contains bundle $\mathbf{z}$. We define a partial order $\succeq^{si}$ on random allocations via $\tilde{\mathbf{x}} \succeq^{si} \tilde{\mathbf{y}} \iff \mathbb{P}(\tilde{\mathbf{x}} \geq \mathbf{z}) \geq \mathbb{P}(\tilde{\mathbf{y}} \geq \mathbf{z})$ for all bundles $\mathbf{z}$. This definition captures the idea that the lottery $\tilde{\mathbf{x}}$ assigns higher probability weight to strictly larger bundles (in the sense of $\geq$) than the lottery $\tilde{\mathbf{y}}$, and there is at least one bundle for which the probability weight is strictly higher.

Definition 13 (Justified envy-freeness based on set-inclusion)**.** A random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *justified envy-free based on set-inclusion*, if for any pair of agents i, j such that $\tilde{\mathbf{e}}_i \succeq^{si} \tilde{\mathbf{e}}_j$, it holds that $\tilde{\mathbf{x}_i} \succeq_i \tilde{\mathbf{x}_j}$.

However, justified envy-freeness based on set inclusion is still restrictive, as it does not allow us to compare agents, who are endowed with bundles of different goods.

Ideally, we would like to compare agent's endowments based on their *value*. However, in general, the value of an endowment depends on the preferences of other traders. Thus, it is not possible to compare two endowments without additional information. Suppose that we find a price vector that corresponds to the the correct marginal rates of substitution for all goods.

For a price vector $\mathbf{p} \in \mathbb{R}^m_{\geq 0}$, the $\mathbf{p}$-*valuation* of a deterministic bundle $\mathbf{x}_i$ is $\mathbf{v_p}(\mathbf{x}) = \mathbf{x} \cdot \mathbf{p} \in \mathbb{R}_{\geq 0}$. We extend the notion of value to a random allocation, and let $\mathbf{v_p}(\tilde{\mathbf{x}})$ be the corresponding random variable of valuations, given the realization of the random allocation. We define a partial order $\succeq^{\mathbf{p}}$ over random bundles as follows: $\tilde{\mathbf{x}} \succeq^{\mathbf{p}} \tilde{\mathbf{y}} \iff \mathbb{P}(\mathbf{v_p}(\tilde{\mathbf{x}}) \geq \mathbf{z}) \geq \mathbb{P}(\mathbf{v_p}(\tilde{\mathbf{y}}) \geq z)$ for all $z \in \mathbb{R}$. That is, $\tilde{\mathbf{x}} \succeq^{\mathbf{p}} \tilde{\mathbf{y}}$ if and only if the random valuation of $\tilde{\mathbf{x}}$ first-order stochastically dominates the random valuation of $\tilde{\mathbf{y}}$.

Definition 14 (Justified envy-freeness based on valuations)**.** Given a price-vector $\mathbf{p}$, a random allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ is *justified envy-free based on* $\mathbf{p}$*-valuations* if for any pair of agents i, j such that $\tilde{\mathbf{e}}_i \succeq^{\mathbf{p}} \tilde{\mathbf{e}}_j$, it holds that $\tilde{\mathbf{x}}_i \succeq^{sd}_i \tilde{\mathbf{x}}_j$.

Note that justified envy-freeness based on either set-inclusion or $\mathbf{p}$-valuations imply equal treatment of equals.

The following theorem shows that any allocation in the competitive core indeed satisfies these ordinal fairness criteria:

Theorem 3. *Given any economy $\mathcal{E} = (\tilde{\mathbf{e}}, \succeq)$, every competitive core allocation is justified envy-free based on set-inclusion and $\mathbf{p}^{eq}$-valuations, where $\mathbf{p}^{eq}$ is the price vector from the corresponding asymptotic CERE.*

Proof Sketch (Theorem 3). Based on our competitive market framework, an agent is constrained only by their random budget, determined by their endowment and the random budget relaxation. First, we show that for δ-CEREs, if endowments are strictly larger in terms of set inclusion or more valuable with respect to a price vector, then the random budget of a trader is first-order stochastically dominant over one with a smaller or less valuable endowment, due to the identical distribution of random budget relaxations. With a dominant budget, a trader can always afford the same or better bundles, ensuring they do not envy another agent whose endowment leads to a dominated budget distribution. Second, we show that these fairness properties carry over to the limit allocation. The full proof is relegated to Appendix A.4.

6.2 Ex-post guarantees: envy and efficiency

In this subsection, we study the ex-post guarantees of competitive core allocations. Ex-post considerations about inclusion in the core or fairness require comparing an ex-post endowment with an ex-post allocation. Therefore, we restrict ourselves in this subsection to deterministic endowments.

First, we note that in our market environment with indivisible goods, it is not possible to guarantee the absence of ex-post justified envy:

Example 4 (Ex-post justified envy)**.** There are three agents, 1, 2, and 3, and two goods, a and b. Initially, agents 1 and 2 each own one unit of a, and agent 3 owns one unit of b. Agents 1 and 2 prefer good b, while agent 3 prefers good a. In any ordinal core allocation, agent 3 receives one unit of good a, and agents 1 and 2 receive a probability distribution over good b and the remaining unit of good a. An ordinal core allocation that treats equals equally and avoids justified envy assigns each agent 1 and 2 either good a or b with probability 0.5. However, when this allocation is realized, the agent receiving good b envies the agent receiving good a, with justified envy due to their identical initial endowments.

We introduce a relaxed version of *envy-freeness up to one good (EF1)* (Budish 2011), called *justified envy-freeness up to one good (JEF1)*.

Definition 15 (Justified envy-freeness up to one good)**.** For deterministic endowments $(\mathbf{e}_1, ..., \mathbf{e}_n)$, a deterministic allocation $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_n)$ is *justified envy-free up to one good (JEF1)* based on

- *set-inclusion* , if for any pair of agents i, j with $\mathbf{e}_i \geq \mathbf{e}_j$, there exists a good k, such that $\mathbf{x}_i \succeq_i \left(\mathbf{x}_j - e^k\right)^+$.

- $\mathbf{p}$-*valuations*, if for a price vector $\mathbf{p}$ and any pair of agents i, j with $\mathbf{p} \cdot \mathbf{e}_i \geq \mathbf{p} \cdot \mathbf{e}_j$, there exists a good k, such that $\mathbf{x}_i \succeq_i \left(\mathbf{x}_j - e^k\right)^+$.

Both versions of JEF1 imply *equal treatment of equals up to one good*: For any pair of agents i, j with $\mathbf{e}_i = \mathbf{e}_j$, there exists a good k, such that $\mathbf{x}_i \succeq_i \left(\mathbf{x}_j - e^k\right)^+$.

The following results summarizes the realizability and ex-post core inclusion and fairness properties of our mechanism:

Theorem 4. *Given an economy $\mathcal{E} = (\mathbf{e}, \succeq)$ with deterministic endowments, every competitive core allocation can be realized over $(\Delta - 1)$-near-feasible core allocations that satisfy JEF1 based on set inclusion and $\mathbf{p}^{eq}$-valuations, where $\mathbf{p}^{eq}$ is the price vector from the corresponding asymptotic CERE.*

Proof Sketch (Theorem 4). Any deterministic allocation realized with positive probability in a competitive core allocation is a δ-CERE allocation for sufficiently small $\delta > 0$ and realized budget relaxations. Let $\mathbf{c}'$ be the aggregate consumption of this allocation. This allocation is feasible with respect to $\mathbf{c}'$ and, as a competitive equilibrium allocation (in the economy with artificial money), it is also in the ex-post core. If no blocking coalition exists using artificial money, then none exists without it, proving the allocation without artificial money is also in the deterministic core. Realizability then follows from Nguyen et al. (2016, Theorem 2.1):

> Any feasible random allocation with respect to $\mathbf{c}$ can be realized through deterministic allocations that are feasible with respect to $\mathbf{c} + (\Delta - 1) \cdot \mathbf{1}$.

For the fairness guarantee, consider a near-feasible deterministic allocation $\mathbf{X}$ realized from a competitive core allocation. Again, this is a δ-CERE allocation for small δ and small realizations of random budget relaxations. Assume agent i envies agent j, unresolved by removing one good, meaning every bundle from $\mathbf{x}_j$ minus one good is unaffordable for agent i, while $\mathbf{x}_j$ is affordable for agent j. Thus, agent i's budget is at most $\frac{m-1}{m}$ times that of agent j. If budget relaxations are small, this leads to a contradiction if agent i's endowment is greater or more valuable at equilibrium prices, ensuring a single good k exists to resolve justified envy. The full proof is in Appendix A.5.

Proof of Theorem 2. Finally, we note that the non-emptiness of the competitive core (Proposition 3), together with Theorem 4 implies our second main result Theorem 2: For any economy $\mathcal{E} = (\mathbf{e}, \succeq)$ with deterministic endowments, there exists an allocation in the ordinal core – any competitive core allocation – that can be realized over $(\Delta - 1)$-near

feasible deterministic core allocations. This also directly implies that there exists a $(\Delta - 1)$-near feasible deterministic core allocation.

7 The Competitive Core Mechanism

In this subsection, we define our *Competitive Core Mechanisms* that maps a revealed economy to a random allocation in the competitive core. We then show that this mechanism is *strategyproof in the large* (SP-L) (Azevedo and Budish 2019).

Let $\Psi \subset \mathbb{N}^m$ represent the finite set of all possible bundles, considering the finite capacities of goods. An agent's i *type* t_i includes their initial endowment $\tilde{\mathbf{e}}i$, their set of acceptable bundles $\Psi_i \subset \Psi$, and their preference relation $\succeq_i$ over these bundles. The *type space* is the product of initial endowments, acceptable bundles, and weak preference relations. Let $t = (t_i)_{i \in N}$ denote the vector of types for all agents, with $t = (t_i, t_{-i})$ specifying agent i's type and t_{-i} the types of all other agents. A *direct* mechanism Φ maps the type space T to the space of random allocations $\tilde{\mathbf{X}}$.

We assume that agents cannot misrepresent their endowments. A mechanism Φ is *semi-anonymous* if it restricts an agent's input to $T_{\tilde{\mathbf{e}}i}$, where $T\tilde{\mathbf{e}}i = \tilde{\mathbf{e}}i \times 2^{\Psi} \times \mathcal{P}(\Psi)$. Thus, agents are grouped based on their initial endowments.

We now define our direct and semi-anonymous mechanism, called the *Competitive Core Mechanism*.

Definition 16 (Competitive Core Mechanism)**.** Our mechanism operates as follows:

1. Each agent submits their initial endowment, a set of acceptable bundles, and their weak ordinal preferences over these bundles.
2. Based on the reported information, the mechanism computes an asymptotic competitive equilibrium from random endowments. This involves calculating a sequence of δ-CEREs (competitive equilibria from random endowments) derived from our fair random budget relaxations.
3. The mechanism then implements the lotteries for each agent, announces the resulting deterministic reallocation of goods, and executes this reallocation.

7.1 Incentive Properties

Consider a direct and semi-anonymous mechanism $\Phi^n(t)$ that maps a type vector with n agents to a random allocation. Let $\Phi_i^n(t)$ denote the random allocation assigned to agent i.

Let T^* be a subspace of the type space where there are only finitely many possible initial endowments $E^* = \tilde{\mathbf{e}}_1, \ldots, \tilde{\mathbf{e}}_{|E^*|}$.[16] As the set of acceptable bundles and the set of weak preference relations over these bundles are both finite, our restriction ensures that the product space of possible types is finite as well. Let $\overline{\mathcal{L}}(T^*)$ denote the space of all probability distributions with full support over T^*. We denote such full support distributions by m.

Next, we define the *expected random allocation* that agent i expects receiving, ex-interim, when reporting type t_i, assuming that the other $n-1$ agents report types independently and identically distributed according to distribution m.

Definition 17 (Expected random allocation)**.** For a direct and semi-anonymous mechanism Φ^n, the *expected random allocation* is

$$\mathbf{\Phi}^n(t_i, m) = \sum_{t_{-i} \in (T^*)^{n-1}} \Phi_i^n(t_i, t_{-i}) \mathbb{P}[t_{-i} \mid t_{-i} \sim \text{i.i.d.}(m)].$$

We extend *Strategyproofness in the Large* (SP-L) (Azevedo and Budish 2019) to our setting with weak ordinal preferences. Consider a trader with type t, including their weak ordinal preference relation $\succeq_i$ over the set of acceptable bundles Ψ_i. Let u_t be a von Neumann-Morgenstern utility function consistent with these preferences. By selecting one such utility function for each type, $\{u_t\}_{t \in T}$ forms a *cardinal representation* of weak ordinal preferences.

A mechanism is SP-L if, for any cardinal representation of ordinal preferencesfull-support and for any i.i.d. distribution of other agents' reports, reporting truthfully is approximately interim optimal in sufficiently large markets.[17]

Definition 18 (Strategyproofness in the large)**.** A direct, semi-anonymous mechanism Φ^n is *strategyproof in the large (SP-L)* if for any cardinal representation of preferences $(u_t)_{t \in T}$, and for all $m \in \overline{\mathcal{L}}(T^*)$ and $\epsilon > 0$, there exists n_0 such that for all $n \geq n_0$, $\tilde{\mathbf{e}} \in E^*$ and all $t_i, t_i' \in T^*_{\tilde{\mathbf{e}}}$, it holds that $u_{t_i}[\Phi_i^n(t_i, m)] \geq u_{t_i}[\Phi_i^n(t_i', m)] - \epsilon$.

We can now directly state our main result on incentives:

Theorem 5 (Incentive Properties of the Competitive Core Mechanism)**.** *The Competitive Core Mechanism is strategyproof in the large. For any cardinal representation of preferences, the maximum possible gains from misreporting converge to* 0 *at a rate of* $n^{-\frac{1}{2}+\epsilon}$*, where* n *is the number of agents.*

[16] This is satisfied, if endowments are deterministic. Hence, it only restricts the space of random endowments.

[17] As we consider semi-anonymous mechanisms, where agents cannot misreport their endowment, we use the extension of SP-L proposed in Azevedo and Budish (2019)[Appendix C].

Proof Sketch (Theorem 5). The Competitive Core Mechanism, $\Phi^{CCM}(\cdot)$, satisfies the equal treatment of equals condition (Theorem 3), meaning that for any two agents i and j with identical endowments $\tilde{\mathbf{e}}_i = \tilde{\mathbf{e}}_j$, it holds that $\Phi_i^{CCM}(t) \succeq_i^{sd} \Phi_j^{CCM}(t)$. This condition also ensures ex-ante envy-freeness for any cardinal representation of utilities within groups defined by equal endowments. Therefore we can use a result from Azevedo and Budish (2019) for envy-free semi-anonymous mechanisms, which states that $\Phi^{CCM}(\cdot)$ is strategyproof in the large, with potential gains from misreporting diminishing at a rate of $n^{-\frac{1}{2}+\epsilon}$ as the number of agents increases. The full proof is relegated to Appendix A.6.

8 Conclusion

This paper offers several positive results for combinatorial exchange with ordinal preferences. We show that the ordinal core is non-empty without imposing restrictions on endowments or preferences. The competitive equilibrium foundations of allocations in the competitive core prove exceptionally useful for studying their fairness, as they address envy in cases of less valuable initial endowments. Competitive core allocations eliminate ex-ante justified envy and ex-post justified envy up to one good, while the mechanism that selects the competitive core is strategyproof in the large. We also offer ex-post guarantees that are crucial for practical implementation: ordinal core allocations can be realized over near-feasible ex-post efficient allocations while competitive core allocations can be realized over near-feasible core allocations when endowments are deterministic.

Our framework is thus the first to provide strong ex-ante and ex-post guarantees for efficiency and fairness in general combinatorial exchange problems. The Competitive Core Mechanism can therefore be applied in combinatorial exchange far beyond unit-demand or restricted domain settings in lieu of TTC-like mechanisms.

There are multiple directions for future research. The first is to study the complexity of computing a competitive core allocation and to develop practical algorithms for the implementation of the CCM. The second is to examine how an implementation of CCM performs empirically relative to other existing mechanisms. The third is to theoretically study relevant extensions to the combinatorial exchange problem, e.g., to allow for a dynamic market in which agents enter or exist the market and resources need to be re-allocated over time.

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, "Random serial dictatorship and the core from random endowments in house allocation problems," *Econometrica*, 1998, *66* (3), 689–701.

____ and ____ , "House allocation with existing tenants," *Journal of Economic Theory*, 1999, *88* (2), 233–260.

____ and ____ , "School choice: A mechanism design approach," *American economic review*, 2003, *93* (3), 729–747.

Abdulkadiroğlu, Atila, Yeon-Koo Che, Parag A Pathak, Alvin E Roth, and Olivier Tercieux, "Efficiency, justified envy, and incentives in priority-based matching," *American Economic Review: Insights*, 2020, *2* (4), 425–442.

Ahani, Narges, Paul Gölz, Ariel D Procaccia, Alexander Teytelboym, and Andrew C Trapp, "Dynamic placement in refugee resettlement," *Operations Research*, 2023, *0* (0).

Andersson, Tommy, Ágnes Cseh, Lars Ehlers, and Albin Erlanson, "Organizing time exchanges: Lessons from matching markets," *American Economic Journal: Microeconomics*, 2021, *13* (1), 338–373.

Arrow, Kenneth J and Gerard Debreu, "Existence of an equilibrium for a competitive economy," *Econometrica*, 1954, *22* (3), 265–290.

Athanassoglou, Stergios and Jay Sethuraman, "House allocation with fractional endowments," *International Journal of Game Theory*, 2011, *40*, 481–513.

Aumann, Robert J, "Markets with a continuum of traders," *Econometrica: Journal of the Econometric Society*, 1964, *32* (1/2), 39–50.

Azevedo, Eduardo M and Eric Budish, "Strategy-proofness in the large," *Rev Econ Stud*, 2019, *86* (1), 81–116.

Aziz, Haris, "Generalizing top trading cycles for housing markets with fractional endowments," *arXiv:1509.03915*, 2015.

Balbuzanov, Ivan and Maciej H Kotowski, "Endowments, exclusion, and exchange," *Econometrica*, 2019, *87* (5), 1663–1692.

Biró, Péter, Flip Klijn, and Szilvia Pápai, "Balanced exchange in a multi-unit Shapley-Scarf market," Technical Report, Barcelona School of Economics, Working Paper 2022.

____ , ____ , and ____ , "Serial rules in a multi-unit Shapley-Scarf market," *Games and Economic Behavior*, 2022, *136*, 428–453.

Bogomolnaia, Anna and Hervé Moulin, "A new solution to the random assignment problem," *J Econ Theory*, 2001, *100* (2), 295–328.

Budish, Eric, "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," *J Polit Econ*, 2011, *119* (6), 1061–1103.

___ and Estelle Cantillon, "The multi-unit assignment problem: Theory and evidence from course allocation at Harvard," *American Economic Review*, 2012, *102* (5), 2237–2271.

___ , Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom, "Designing random allocation mechanisms: Theory and applications," *American economic review*, 2013, *103* (2), 585–623.

Debreu, Gerard and Herbert Scarf, "A Limit Theorem on the Core of an Economy," *International Economic Review*, 1963, *4* (3), 235–246.

Delacrétaz, David, Scott Duke Kominers, and Alexander Teytelboym, "Matching mechanisms for refugee resettlement," *American Economic Review*, 2023, *113* (10), 2689–2717.

Doğan, Battal and Kemal Yıldız, "Efficiency and stability of probabilistic assignments in marriage problems," *Games and Economic Behavior*, 2016, *95*, 47–58.

Dugundji, James and Andrzej Granas, *Fixed Point Theory*, New York: Springer, 2003.

Dur, Umut Mert and M Utku Ünver, "Two-sided matching via balanced exchange," *Journal of Political Economy*, 2019, *127* (3), 1156–1177.

Echenique, Federico, Antonio Miralles, and Jun Zhang, "Fairness and efficiency for probabilistic allocations with endowments," Technical Report, Working paper 2019.

___ , ___ , and ___ , "Constrained pseudo-market equilibrium," *American Economic Review*, 2021, *111* (11), 3699–3732.

___ , ___ , and ___ , "Balanced equilibrium in pseudo-markets with endowments," *Games and Economic Behavior*, 2023, *141*, 428–443.

Feng, Di, Bettina Klaus, and Flip Klijn, "Characterizing the typewise top-trading-cycles mechanism for multiple-type housing markets," *Games and Economic Behavior*, 2024, *146*, 234–254.

Fujita, Etsushi, Julien Lesca, Akihisa Sonoda, Taiki Todo, and Makoto Yokoo, "A complexity approach for core-selecting exchange under conditionally lexicographic preferences," *Journal of Artificial Intelligence Research*, 2018, *63*, 515–555.

Garg, Jugal, Thorben Tröbst, and Vijay V Vazirani, "One-sided matching markets with endowments: Equilibria and algorithms," *arXiv preprint arXiv:2009.10320*, 2020.

Gul, Faruk, Wolfgang Pesendorfer, Mu Zhang et al., "The Efficient Allocation of

Indivisible Goods," Technical Report, Working paper 2022.

He, Yinghua, Antonio Miralles, Marek Pycia, and Jianye Yan, "A pseudo-market approach to allocation with priorities," *American Economic Journal: Microeconomics*, 2018, *10* (3), 272–314.

Hildenbrand, Werner, "Core of an economy," *Handbook of mathematical economics*, 1982, *2*, 831–877.

Huang, Weixiao, Elise Jade Deshusses, Jennifer Ann Pazour, Yunus Doğan Telliel, Sarah Eliza Stanlick, and Andrew Christopher Trapp, "Empowering Collective Impact: Introducing SWAP for Resource Sharing," in "Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization" 2023, pp. 1–12.

Hylland, Aanund and Richard Zeckhauser, "The efficient allocation of individuals to positions," *J Polit Econ*, 1979, *87* (2), 293–314.

Klaus, Bettina and Eiichi Miyagawa, "Strategy-proofness, solidarity, and consistency for multiple assignment problems," *International Journal of Game Theory*, 2002, *30*, 421–435.

Konishi, Hideo, Thomas Quint, and Jun Wako, "On the Shapley–Scarf economy: the case of multiple types of indivisible goods," *Journal of Mathematical Economics*, 2001, *35* (1), 1–15.

Kornbluth, Daniel and Alexey Kushnir, "Undergraduate Course Allocation through Competitive Markets," Technical Report, Available at SSRN 3901146 2021.

Manjunath, Vikram, "Stability and the core of probabilistic marriage problems," *Available at SSRN 1809941*, 2013.

___ and Alexander Westkamp, "Strategy-proof exchange under trichotomous preferences," *Journal of Economic Theory*, 2021, *193*, 105197.

Miralles, Antonio and Marek Pycia, "Foundations of pseudomarkets: Walrasian equilibria for discrete resources," *Journal of economic theory*, 2021, *196*, 105303.

Morrill, Thayer and Alvin E Roth, "Top trading cycles," *Journal of Mathematical Economics*, 2024, *112*, 102984.

Munkres, James R., *Topology*, 2nd ed., Upper Saddle River, NJ: Prentice Hall, 2000.

Nguyen, Thành, Ahmad Peivandi, and Rakesh Vohra, "Assignment problems with complementarities," *Journal of Economic Theory*, 2016, *165*, 209–241.

___ , Alexander Teytelboym, and Shai Vardi, "Efficient Combinatorial Assignment," Technical Report, Working paper 2024.

Nguyen, Thành and Rakesh Vohra, "(Near) Substitute Preferences and Equilibria with Indivisibilities," *Journal of Political Economy*, 2024.

Pápai, Szilvia, "Strategyproof assignment by hierarchical exchange," *Econometrica*,

2000, *68* (6), 1403–1433.

____ , "Strategyproof exchange of indivisible goods," *Journal of Mathematical Economics*, 2003, *39* (8), 931–959.

____ , "Exchange in a general market with indivisible goods," *Journal of Economic Theory*, 2007, *132* (1), 208–235.

Prendergast, Canice, "How food banks use markets to feed the poor," *Journal of Economic Perspectives*, 2017, *31* (4), 145–162.

____ , "The allocation of food to food banks," *Journal of Political Economy*, 2022, *130* (8), 1993–2017.

Pycia, Marek and M Utku Ünver, "Incentive compatible allocation and exchange of discrete resources," *Theoretical Economics*, 2017, *12* (1), 287–329.

Roth, Alvin E, "Incentive compatibility in a market with indivisible goods," *Economics letters*, 1982, *9* (2), 127–132.

____ , "Repugnance as a Constraint on Markets," *Journal of Economic perspectives*, 2007, *21* (3), 37–58.

____ **and Andrew Postlewaite**, "Weak versus strong domination in a market with indivisible goods," *Journal of Mathematical Economics*, 1977, *4* (2), 131–137.

____ **, Tayfun Sönmez, and M Utku Ünver**, "Kidney exchange," *The Quarterly journal of economics*, 2004, *119* (2), 457–488.

Shapley, Lloyd and Herbert Scarf, "On cores and indivisibility," *Journal of mathematical economics*, 1974, *1* (1), 23–37.

Sönmez, Tayfun, "Strategy-proofness and essentially single-valued cores," *Econometrica*, 1999, *67* (3), 677–689.

Varian, Hal R., "Equity, envy, and efficiency," *J Econ Theory*, 1974, *9*, 63–91.

Wako, Jun, "Coalition-proof Nash allocation in a barter game with multiple indivisible goods," *Mathematical Social Sciences*, 2005, *49* (2), 179–199.

Yılmaz, Özgür, "The probabilistic serial mechanism with private endowments," *Games and Economic Behavior*, 2010, *69* (2), 475–491.

A Proofs

A.1 Proof of Proposition 1

Proof. We prove this proposition in multiple steps. In step 1, we extend random demand to allow for an infinite consumption of money, if its price hits zero, and we provide necessary topological properties of the space of random allocations. In step 2, we define a correspondence on the product space of prices and random allocations, and show that this map admits a fixed point. Finally, in step 3, we show that the fixed points of our constructed correspondence are δ-CEREs.

Step 1. Extension of Random Demand. For any $\delta \geq 0$, and for any δ-CERE, we can always scale prices by a constant factor. Thus, we normalize prices such that they all sum up to 1. That is, we consider the simplex $\Delta = \{\mathbf{p} \in \mathbb{R}^{m+1}_{+} : \sum_{j=1}^{m+1} \mathbf{p}^j = 1\}$. We extend the definition of random demand to include demand for money, if the price of money is equal to 0: Recall that for a given price vector $\mathbf{p}$, the random demand of agent i is constructed as follows: First, agents receive a realized endowment from their initial probability distribution of goods and money, that is $\mathbf{e}_i \sim \tilde{\mathbf{e}}_i$. Second, their budget is randomly relaxed by $\mathbf{b}_i \sim \tilde{\mathbf{b}}_i$ times the value of the entire economy. Third, given their realized random budget, agents demand their favourite affordable bundle, with a tie-breaking rule if necessary. Given the constructed preferences, agent i demands their favourite bundle of goods $\mathbf{x}_i^G \in \Psi_i$ that is affordable, and the amount of money $\mathbf{x}_i^M \in \mathbb{R}_{\geq 0}$ to fully satiate their budget, if the price of money is strictly positive, and an infinite amount of money, if the price of money is zero. Thus, agents demand an infinite amount of money, if and only if the price of money is equal to zero. More formally, if $\mathbf{p}^M > 0$, then

$$\chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) := \max_{\succeq_i} \left\{ \mathbf{x}_i = (\mathbf{x}_i^G, \mathbf{x}_i^M) \in \Psi_i \times \mathbb{R}_{\geq 0} \text{ and } \mathbf{p}\cdot\mathbf{x} \leq \mathbf{p}\cdot\mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}\cdot\mathbf{c} | \mathbf{e}_i \sim \tilde{\mathbf{e}}_i \wedge \mathbf{b}_i \sim \tilde{\mathbf{b}}_i \right\},$$

where ties among expenditure-minimizing bundles are broken uniformly at random. If $\mathbf{p}^M = 0$, then

$$\chi_i^G(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) := \max_{\succeq_i} \left\{ \mathbf{x}_i^G \in \Psi_i \mid \mathbf{p}^G \cdot \mathbf{x}^G \leq \mathbf{p} \cdot \mathbf{e}_i + \beta_i \cdot \mathbf{p} \cdot \mathbf{c}, \text{ and minimize } \mathbf{p}^G \cdot \mathbf{x}^G \right\},$$

and $\chi_i^M(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) = \infty$, where ties among expenditure-minimizing bundles are broken uniformly at random.

Thus, for a given price vector $\mathbf{p}$, random demand of agent i induces a probability

distribution over the space $\Psi_i \times \overline{\mathbb{R}}_{\geq 0}$, where $\overline{\mathbb{R}}_{\geq 0} = \mathbb{R}_{\geq 0} \cup \{\infty\}$, denotes the one-point compactification of $\mathbb{R}_{\geq 0}$, see, e.g. Munkres (2000)[Chapter 29]. We denote the space of all such probability distributions as $\mathcal{P}(\Psi \times \overline{\mathbb{R}}_{\geq 0})$. In the following, we briefly outline the topological properties of this space, as we will define a fixed-point mapping on the product space in the next step.

Step 2. Topology of the Space of Probability Distributions $\mathcal{P}(\Psi_i \times \overline{\mathbb{R}}_{\geq 0})$ We first study the space $\Psi_i \times \overline{\mathbb{R}}_+$. $\Psi_i \subset \mathbb{N}^m$ is a finite and thus discrete set, which we can endow with the discrete topology. Recall that we denote by $\overline{\mathbb{R}}_+$ the one-point compactification of $\mathbb{R}_{\geq 0} = [0, \infty)$, and write $\overline{\mathbb{R}}_+ = [0, \infty) \cup \{\infty\}$. We endow this set with the Borel σ-algebra $\mathcal{B}(\overline{\mathbb{R}}_+)$ of the one-point compactification, which is generated by the open sets in $\overline{\mathbb{R}}_+$, including the sets of the form $(a, \infty] \cup \{\infty\}$ for $a \geq 0$. We consider the product space $\Psi \times \overline{\mathbb{R}}_+$ with the product topology, and the space of probability measures on this space, and show that it, when endowed with the weak topology, has desirable topological properties.

- **Compactness.** The set Ψ_i is trivially compact as a finite set, and the one-point compactification is compact by construction. For the set $\Psi \times \overline{\mathbb{R}}_+$, we use the product topology. The product of two compact spaces is compact with respect to the product topology (Tychonoff's theorem). Because the underlying space is compact, every probability measure on this space is tight. Prokhorov's theorem implies that the space $\mathcal{P}(\Psi \times \overline{\mathbb{R}}_+)$ is relatively compact with respect to the weak topology. Because the underlying space is also closed, we get that the space of probability measures is compact with respect to the weak topology.
- **Convexity.** The convex combination of two probability measures is again a probability measure, proving that the space is convex.
- **Subset of a Hausdorff locally convex topological vector space.** The space of probability measures on a compact space X can be embedded in the dual space of continuous functions $C(X)^*$, which is Hausdorff locally convex with the weak* topology.

Next, consider the product space $\mathcal{P}(\Psi_1 \times \overline{\mathbb{R}}_{\geq 0}) \times ... \times \mathcal{P}(\Psi_n \times \overline{\mathbb{R}}_{\geq 0})$ corresponding to the space of probability distributions for all agents. As a product space, it inherits the three desirable properties from above. Lastly, we define the space $Z = \Delta \times \mathcal{P}(\Psi_1 \times \overline{\mathbb{R}}_{\geq 0}) \times ... \times \mathcal{P}(\Psi_n \times \overline{\mathbb{R}}_{\geq 0})$, which also contains the price simplex. By the same reasoning, this space Z is compact, convex, and a subset of a Hausdorff locally convex topological vector space.

Step 3. Fixed point mapping. Next, we define a correspondence $\Phi : Z \to 2^Z$ and apply the Glicksberg-Fan extension of the Kakutani-fixed point theore (Dugundji and Granas 2003)[Chapter II, Section 5.8]:

> Let S be a non-empty, compact, and convex subset of a Hausdorff locally convex space. Let $\Phi : Z \to 2^Z$ be upper hemicontinuous and $\Phi(z)$ is non-empty, compact and convex for all $z \in Z$. Then, Φ has a fixed point.

The construction of Φ is componentwise and as follows: For each agent i, we define a map $f_i : Z \to \mathcal{P}(\Psi_i \times \overline{\mathbb{R}}_{\geq 0})$ as $f_i(\mathbf{p}, \tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n) = \chi_i(\mathbf{p}, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i)$. That is, the map only depends on the price vector $\mathbf{p} \in \Delta$ and maps the price vector to the random demand of agent i, which is a probability distribution over $\Psi_i \times \overline{\mathbb{R}}_{\geq 0}$. First, for every price vector, this function is well-defined. That is, because for each realization of the endowment and budget relaxation, agents demand a unique affordable bundle because of the tie-breaking rule. Thus, random demand induces a unique probability distribution on $\mathcal{P}(\Psi_i \times \overline{\mathbb{R}}_{+})$. Because this outcome is unique, this map, interpreted as a correspondence, is non-empty, convex, and compact. Finally, consider an agents random budget, that is, the random variable $\mathbf{p} \cdot \tilde{\mathbf{e}}_i + \tilde{\mathbf{b}}_i \cdot \mathbf{p} \cdot \mathbf{c}$ induced from realization from the random endowment vector $\mathbf{e}_i$ and the random variable $\tilde{\mathbf{b}}_i$ corresponding to the budget relaxation. For a fixed price vector $\mathbf{p}$, and any distribution of endowments $\mathbf{e}_i \sim \tilde{\mathbf{e}}_i$, the random variable is continuously distributed, because budget relaxation $\mathbf{b}_i \sim \tilde{\mathbf{e}}_i$ are also continuously distributed. Hence, as a function of $\mathbf{p}$, it follows that the map $f_i : Z \to \mathcal{P}(\Psi_i \times \overline{\mathbb{R}}_{\geq 0})$ is continuous, and thus, in particular, also upper hemicontinuous.

Next, we define an additional map $f_0 : Z \to \Delta$. The intuition behind this map is as follows: Given probability distributions on $\Psi_i \times \overline{\mathbb{R}}_{+}$ for all agents, that correspond to the random demand of all agents, we compute the aggregate expected demand as the sum of individual expected demand. If for some good, this demand exceeds the relaxed capacity of the good, decrease the price for the good. If for some good, the demand is less than the relaxed capacity, increase the price of the good.

For this, consider a random demand $\tilde{\mathbf{x}}_i \in \mathcal{P}(\Psi_i \times \overline{\mathbb{R}}_{\geq 0})$ of an agent i. This probability distribution induces a random vector $(\mathbf{x}_i^G, \mathbf{x}_i^M)$, where $\mathbf{x}_i^G \in \Psi_i$ and $\mathbf{x}_i^M \in \overline{\mathbb{R}}_{\geq 0} = [0, \infty) \cup \{\infty\}$. We define the expectation of this random vector component-wise as $\mathbb{E}[\tilde{\mathbf{x}}_i] = \left(\mathbb{E}\left[\tilde{\mathbf{x}}_i^G\right], \mathbb{E}\left[\tilde{\mathbf{x}}_i^M\right]\right)$. For the first component, the expectation of the random vector is well-defined, as any probability distribution $\tilde{\mathbf{x}}_i$ induces a lottery $\tilde{\mathbf{x}_i^G}$ over Ψ_i. For the second component, let μ_i be the underlying probability distribution on $\Psi_i \times \overline{\mathbb{R}}_{+}$. We define the expectation as the usual expectation, if the probability mass at infinity is

zero, and set the expectation to be infinite otherwise:

$$\mathbb{E}[\tilde{\mathbf{x}}_i^M] = \begin{cases} \int_0^\infty x d\mu_i(\tilde{\mathbf{x}}_i^M = x) & \text{if } \mu_i(\tilde{\mathbf{x}_i^M} = \infty) = 0 \\ \infty & \text{if } \mu_i(\tilde{\mathbf{x}_i^M} = \infty) > 0. \end{cases} \tag{1}$$

Moreover, let $\mathbb{E}[\tilde{\mathbf{x}}] = \sum_{i=1}^n \mathbb{E}[\tilde{\mathbf{x}}_i]$ be the aggregate expected demand, that is, the sum of indiviual expected demand. For the last component that corresponds to money, the aggregate expected demand is infinite, if and only if the individual demand is infinite for at least one agent.

We define a map $\mathbb{E} : Z \to \mathbb{R}^m \times \overline{\mathbb{R}}_+$ that maps the probability distributions $\tilde{\mathbf{x}}$ for all agents to the aggregate expected consumption. This map $\tilde{\mathbf{x}} \mapsto \mathbb{E}[\tilde{\mathbf{x}}]$ is continuous, and thus also upper hemicontinuous. Next, we define a map $f_0 : Z \to 2^\Delta$

$$f_0(\mathbb{E}[\tilde{\mathbf{x}}]) = \begin{cases} \arg\max_{\mathbf{p}' \in \Delta} \mathbf{p}' \cdot (\mathbb{E}[\tilde{\mathbf{x}}] - (1+\delta)\mathbf{c}) & \text{if } \mathbb{E}[\tilde{\mathbf{x}}^M] < \infty \\ (0, ..., 0, 1) & \text{if } \mathbb{E}[\tilde{\mathbf{x}}^M] = \infty. \end{cases} \tag{2}$$

For fixed $y \in \mathbb{R}^m \times \overline{\mathbb{R}}_+$, the correspondence $y \mapsto f_0(y)$ is by construction non-empty, compact, convex valued, and upper hemicontinuous. Thus, composing the map $\Psi(\cdot)$ yields a correspondence $f_0 : Z \to 2^\Delta$ that is non-empty, compact, convex valued, and upper hemicontinuous.

Finally, we construct the following product mapping $\Phi : Z \to Z$ via

$$\Phi(\mathbf{p}, \tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n) \mapsto (f_0(\mathbb{E}(\mathbf{p}, \tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)), f_1(\mathbf{p}, \tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n), ..., f_n(\mathbf{p}, \tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)). \tag{3}$$

This map satisfies all conditions for the Kakutani-Glicksberg-Fan fixed-point theorem. Thus, there exists a fixed point, denoted by $(\mathbf{p}^*, \tilde{\mathbf{x}}_1^*, ..., \tilde{\mathbf{x}}_n^*)$.

Step 4. Properties of the Fixed Point . As a last step, we show that any fixed point $(\mathbf{p}^*, \tilde{\mathbf{x}}_1^*, ..., \tilde{\mathbf{x}}_n^*)$ of the correspondence Φ corresponds to a δ-CERE.

First, we note that at this fixed point, the price of money must be strictly positive, that is, $\mathbf{p}^{*,M} > 0$. If $\mathbf{p}^{*,M} > 0$, then by construction, agents would demand an infinite amount, whenever they do not fully expend their budget on bundles. However, given the continuous budget relaxations, with probability one, agents will not fully expend their budget when consuming their favourite bundle. For a fixed price vector $\mathbf{p}$, prices of bundles are fixed. However, the distribution of an agent's budget is continuous. Hence, with probability one, all agents would demand an infinite amount of money. Thus, if the price of money was zero, it would hold that $\mathbb{E}[\tilde{\mathbf{x}}^{*,M}] = \infty$. In this case, the

price-adjusting map f_0 chooses the price $(0, ..., 0, 1)$. This shows that any price-vector with $\mathbf{p}^M = 0$ cannot be part of a fixed point, as the price of money changes from zero to 1. At this fixed-point, it not just holds that $\mathbb{E}[\tilde{x}] < \infty$, but for every realization of endowments, after demanding their favourite affordable bundle, agents only demand a finite amount of money to satiate their budget.

Next, we show that for the price vector $\mathbf{p}^*$ and allocation $(\tilde{\mathbf{x}}_1^*, ..., \tilde{\mathbf{x}}_n^*)$, we get the following version of the *Walras Law* for aggregate expected consumption:

$$\mathbf{p}^* \cdot \mathbb{E}[\tilde{\mathbf{x}}^*] = (1 + \delta) \cdot \mathbf{p}^* \cdot \mathbf{c}. \tag{4}$$

Since the price of money is positive, it follows from agent's preferences and the definition of random demand, that for every realization of an agent's initial endowment, they fully expend their budget. That is, for each realization of their random demand $\chi_i(\mathbf{p}^*, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) = (\mathbf{x}_i^g, \mathbf{x}_i^M)$ it holds that

$$\mathbf{p}^* \cdot \chi_i(\mathbf{p}^*, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i) = \mathbf{p}^* \cdot \mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}^* \cdot \mathbf{c}. \tag{5}$$

Taking the expectation on both sides and summing over all agents now yields Equation (4).

Moreover, because it holds that $\mathbb{E}[\tilde{\mathbf{x}}^{*,M}] < \infty$, we know from Equation (2) that

$$\mathbf{p}^* \in \arg\max_{\mathbf{p} \in \Delta} \mathbf{p} \cdot (\mathbb{E}[\tilde{\mathbf{x}}^*] - (1 + \delta)\mathbf{c}) \tag{6}$$

This implies that for all goods j, it must hold that $(\mathbb{E}[\tilde{\mathbf{x}}^{*,j}] - (1 + \delta)\mathbf{c}^j) \leq 0$. That is, because otherwise, setting $\mathbf{p}^j$=1 and all other prices to zero would lead to a strictly positive value, while it follows from the version of Walras Law that the total value at the fixed point is zero. Moreover, if $\mathbb{E}[\tilde{\mathbf{x}}^{*,j}] - (1 + \delta)\mathbf{c}^j < 0$, then it must hold that $\mathbf{p}^j = 0$. Thus, if $\mathbf{p}^j > 0$, then $\mathbb{E}[\tilde{\mathbf{x}}^{*,j}] - (1 + \delta)\mathbf{c}^j = 0$.

Finally, since each $\tilde{\mathbf{x}}_i^*$ is derived from $\chi_i(\mathbf{p}^*, \tilde{\mathbf{e}}_i, \tilde{\mathbf{b}}_i)$, the fixed point allocation corresponds to the random demand for all agents, satisfying the first condition of a δ-CERE, which finishes the proof. □

A.2 Proof of Proposition 2

Proof. We first prove that any δ-CERE allocation is individually rational. Next, we show that an allocation is a δ-CERE allocation if and only if it admits no δ-ordinal blocking coalitions.

Individual Rationality. Consider any δ-CERE allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$. Given any realization of endowments and random budget relaxations, agents can always afford to purchase the realization of their own endowment in a δ-CERE. That is, because for any price vector $\mathbf{p}$, an agent's random budget is only relaxed, and hence, their endowment is always affordable. Thus, for any realization, the random demand of agent i is either the realization from their initial endowment or a bundle that they prefer. Hence, in the resulting random allocation, probability mass is only shifted to more preferred bundles. Thus, the resulting allocation over bundles from a δ-CERE first order stochastically dominates an agent's initial endowment. That is, $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{e}}_i$.

No δ-ordinal blocking coaltions. Consider a δ-feasible allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$ of goods. If this allocation is induced by a δ-CERE in the economy without artificial money, there exist prices $\mathbf{p} \geq 0$ for goods that satisfies the following:

1. $\sum_{i=1}^n \mathbb{E}[\tilde{\mathbf{x}}_i]^j \leq (1+\delta)\mathbf{c}^j$ for every good j, and $\sum_{i=1}^n \mathbb{E}[\tilde{\mathbf{x}}_i]^j = (1+\delta)\mathbf{c}^j$, if $p^j > 0$.
2. For every agent i and a pair of deterministic bundles $\mathbf{y}' \succ_i \mathbf{y}$, where $\mathbf{y}$ is assigned with positive probability in $\tilde{\mathbf{x}}_i$, we have that $\mathbf{p} \cdot (\mathbf{y}' - \mathbf{y}) \geq 1$. If $\mathbf{y}' \sim_i \mathbf{y}$, we only have that $\mathbf{p} \cdot (\mathbf{y}' - \mathbf{y}) \geq 0$.

The second condition holds because if an agent consumes a bundle with a positive probability, this implies that there exists a realization of the random endowment such that it is the optimal and expenditure-minimizing bundle for the corresponding budget relaxation. Thus, for any more preferred bundle, the price must be strictly higher. By scaling the price, we can assume that the different is at least 1 for all such bundles. For two equivalent bundles, if one is assigned with positive probability, by expenditure-minimization, the price of the other bundle must be weakly greater.

Conversely, if there exists a set of prices, denoted as $\mathbf{p}$, satisfying equations (1) and (2), then $(\tilde{\mathbf{X}}, \mathbf{p})$ form a CERE in the economy, where endowments are re-distributed as follows: $\tilde{\mathbf{e}}_i \overset{d}{\sim} \tilde{\mathbf{x}}_i$, that is, each agent owns the resulting lottery as an initial endowment. Then, each agent can always afford the realization of their initial lottery, and every strictly preferred bundle of goods has a strictly higher price. Now choose the price of money sufficiently small, such that all strictly preferred bundles remain unaffordable. Note that this allocation is hence also individually rational, as all agents remain with their new initial endowment.

Hence, when provided with a δ-feasible allocation $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, ..., \tilde{\mathbf{x}}_n)$, we can determine, whether it is a δ-CERE allocation by using a linear program that identifies prices satisfying both of the conditions from above. Let $\{\mathcal{A}\mathbf{p} \geq b, \mathbf{p} \geq 0\}$ be the description of the linear program, where b is the vector with entry 1 for comparing bundles $\mathbf{y}' \succ_i \mathbf{y}$

with $\mathbf{p} \cdot (\mathbf{y}' - \mathbf{y}) \geq 1$ and with entry 0 for bundles $\mathbf{y}' \sim_i \mathbf{y}$ with $\mathbf{p} \cdot (\mathbf{y}' - \mathbf{y}) \geq 0$. By the Farkas Lemma, a solution $\mathbf{p}$ exists if and only if $\{\mathcal{A}^T\lambda \leq 0, \lambda \geq 0; \mathbf{b}^T\lambda > 0\}$ does not have a solution.

Finally, we show that the dual problem has a solution if and only if the δ-feasible allocation $\tilde{\mathbf{X}}$ admits a δ-ordinal blocking coalition.

First, an equivalent expression for strict stochastic dominance $\tilde{\mathbf{y}} \succ_i^{sd} \tilde{\mathbf{x}}$ is that $\tilde{\mathbf{y}}$ can be derived from $\tilde{\mathbf{x}}$ through a sequence of moves involving the reallocation of probabilities from less preferred to more preferred bundles, where at least one move shifts mass to a bundle that is strictly more preferred. Second, we note that the dual variable λ have the following interpretation: Each coordinate corresponds to a tuple $i, \mathbf{y}', \mathbf{y}$, where $\mathbf{y}' \succeq_i \mathbf{y}$ and $\mathbb{P}(\tilde{x}_i = \mathbf{y}) > 0$. $\mathcal{A}^T\lambda \leq 0$ implies that for every good j that is fully allocated

$$\sum_{i,\mathbf{y}',\mathbf{y}} \lambda_{i,\mathbf{y}',\mathbf{y}} \cdot y'_j - \sum_{i,\mathbf{y}',\mathbf{y}} \lambda_{i,\mathbf{y}',\mathbf{y}} \cdot y_j \leq 0. \tag{7}$$

By assumption, $\mathbb{P}(\tilde{\mathbf{x}}_i = \mathbf{y}) > 0$ holds. Thus, we can normalize λ such that $\lambda_{i,\mathbf{y}',\mathbf{y}} \leq \mathbb{P}(\tilde{\mathbf{x}}_i = \mathbf{y})$. We interpret $\lambda_{i,\mathbf{y}',\mathbf{y}}$ as the probability mass indicating the extent to which agent i shifts from bundle $\mathbf{y}$ to the strictly preferred bundle $\mathbf{y}'$. Equation (7) asserts that all agents can make such shifts without violating the resource constraints of goods, which are fully allocated.

For goods that are under-allocated in the allocation $\tilde{\mathbf{X}}$, Equation (7) may not hold, as the linear program contains additional conditions to satisfy condition (i). However, we can scale down λ by a constant factor, to ensure that the difference is within the limits of the remaining resources for that particular good:

$$\sum_{i,\mathbf{y}',\mathbf{y}} \lambda_{i,\mathbf{y}',\mathbf{y}} \cdot y'_j - \sum_{i,\mathbf{y}',\mathbf{y}} \lambda_{i,\mathbf{y}',\mathbf{y}} \cdot y_j \leq (1+\delta)c_j - \left(\sum_{i=1}^{n} \mathbb{E}[\tilde{\mathbf{x}}_i]\right)^j. \tag{8}$$

Finally, let $\tilde{\mathbf{z}}_i$ represent the resulting lottery for agent i after the probability shift according to λ. Because all agents shift probabilities to weakly preferred bundles, it holds for all i that $\tilde{\mathbf{z}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i$. Moreover, the condition $\mathbf{b}^T\lambda > 0$ indicates that there is at least one agent with a probability shift to a strictly preferred bundle that is strictly positive. Thus, at least one agent is strictly better off, that is, there exists j, such that $\tilde{\mathbf{z}}_j \succ_j^{sd} \tilde{\mathbf{x}}_j$. Equations (7) and (8) together imply that the resulting allocation $(\tilde{\mathbf{z}}_1, ..., \tilde{\mathbf{z}}_n)$ is feasible with respect to $(1+\delta)\mathbf{c}$. Hence, the allocation $\tilde{\mathbf{X}}$ is not δ-ordinally efficient, and thus, also not in the δ-ordinal core, as the grand coalition is a δ-ordinal blocking coalition.

Conversely, suppose that there exists a blocking coalition $S \subseteq N$. For these agents

i, let $\lambda_{i,\mathbf{y}',\mathbf{y}}$ be precisely the probability mass indicating the extent to which agent i shifts from bundle $\mathbf{y}$ to the more preferred bundle $\mathbf{y}'$ in the blocking coalition. For all other agents $i' \in N \setminus S$, let $\lambda_{i',\mathbf{y}',\mathbf{y}}$ be zero for all y and y'. The collection of all such λ's solves the dual problem. Hence, the dual problem has a solution if and only if there is no δ-ordinal blocking coalition.

Thus, there exists a solution to the primal problem, which provides a δ-CERE if and only if there are no δ-ordinal blocking coalitions. Together with the result on individual rationality, this finishes the proof. □

A.3 Proof of Proposition 3

Proof. We will prove this theorem through a series of steps. First, we show that the product space of prices and random allocations is sequentially compact, ensuring the existence of a limit. Next, we establish that the limit allocation is exactly feasible. In the third step, we show that the limit allocation does not allow for ordinal blocking coalitions. Finally, we prove that the limit is individually rational, which proves that the limit allocation is in the ordinal core.

Existence. We leverage a standard topological argument. Consider a strictly decreasing sequence $(\delta_k)_{k\geq 1}$, such that for all k, $\delta_k > 0$ and $\lim_{k\to\infty} \delta_k = 0$. Let $\tilde{\mathbf{X}}^{\delta_k}$ denote a δ_k-CERE allocation of bundles with the corresponding price vector $\mathbf{p}^{\delta_k}$ for goods (excluding artificial money), which exists by Proposition 1. We normalize the prices of goods so that they sum to 1, that is, the price vector $\mathbf{p}$ lies in $\Delta = \{\mathbf{p} \in \mathbb{R}^m_+ : \sum_{j=1}^m \mathbf{p}^j = 1\}$. Thus, for every $\delta_k > 0$, the δ_k-CERE allocation of goods (excluding artificial money) is an element of the product space formed by the m-dimensional price simplex and the space of all lotteries over acceptable bundles. With the standard product topology, this is a compact metric space: The price simplex is a metric space as a subspace of $\mathbb{R}^m$. The space of lotteries over k bundles is the k-dimensional simplex, and hence also a subspace of $\mathbb{R}^k$. Thus, the product space is metrizable and compact. By a standard result from topology, this space is sequentially compact, see, e.g., Munkres (2000). Therefore, the sequence of prices and allocations according to the δ_k-CEREs admits a convergent subsequence, where we denote the limit prices and allocation by $\mathbf{p}^\infty$ and $\tilde{\mathbf{X}}^\infty$, respectively. This limit allocation is now by construction an asymptotic CERE allocation for the convergent subsequence. To simplify notation for the rest of the proof, let $(\delta_k)_{k\geq 1}$ denote this subsequence .

Feasibility. Along the δ_k-sequence, all the random allocations $\tilde{\mathbf{X}}^{\delta_k}$ are δ_k-feasible (Proposition 2):

$$\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i^{\delta_k}] \leq (1 + \delta_k)\mathbf{c}. \tag{9}$$

Because the product space of allocations is compact and the expectation operator is continuous, we can interchange limit and expectation and get that

$$\lim_{k \to \infty} \sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i^{\delta_k}] = \sum_{i \in N} \mathbb{E}[\lim_{k \to \infty} \tilde{\mathbf{x}}_i^{\delta_k}] = \sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i^{\infty}]. \tag{10}$$

Moreover, δ_k-feasibility directly implies that

$$\lim_{k \to \infty} \sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i^{\delta_k}] \leq \lim_{k \to \infty} (1 + \delta_k)\mathbf{c} = \mathbf{c}. \tag{11}$$

Combining Equations (10) and (11) yields that $\sum_{i \in N} \mathbb{E}[\tilde{\mathbf{x}}_i^{\infty}] = \mathbf{c}$, which proves that the limit allocation $\tilde{\mathbf{X}}^{\infty}$ is indeed exactly feasible.

No Ordinal Blocking Coalitions. Assume for contradiction that the limit allocation $\tilde{\mathbf{X}}^{\infty}$ admits an ordinal blocking coalition. Then there exists a subset of traders $S \subset N$ with a partial allocation $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_i)_{i \in N}$ such that $\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i^{\infty}$ for all $i \in S$, and $\tilde{\mathbf{y}}_i \succ_i^{sd} \tilde{\mathbf{x}}_i^{\infty}$ for at least one agent i with

$$\sum_{i \in S} \mathbb{E}[\tilde{\mathbf{y}}_i] \leq \sum_{i \in S} \mathbb{E}[\tilde{\mathbf{e}}_i]. \tag{12}$$

We say that $\tilde{\mathbf{Y}}$ is feasible. We can construct a sequence of δ_k-feasible partial allocations that converge to the limit allocation in the product space. That is, we consider a sequence $\tilde{\mathbf{Y}}^{\delta_k} = (\tilde{\mathbf{y}}_i^{\delta_k})_{i \in S}$ such that $\tilde{\mathbf{Y}}^{\delta_k} \to \tilde{\mathbf{Y}}$.

We define the cumulative distribution functions (CDFs) of $\tilde{\mathbf{x}}^{\infty}$ and $\tilde{\mathbf{y}}$ as

$$F_{\tilde{\mathbf{x}}^{\infty}}(\mathbf{z}) = \mathbb{P}(\tilde{\mathbf{x}}^{\infty} \preceq_i \mathbf{z}) \text{ and } F_{\tilde{\mathbf{y}}}(\mathbf{z}) = \mathbb{P}(\tilde{\mathbf{y}} \preceq_i \mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{13}$$

We can express first-order stochastic dominance in terms of these CDFs as

$$\tilde{\mathbf{y}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_i^{\infty} \text{ if and only if } F_{\tilde{\mathbf{y}}_i}(\mathbf{z}) \leq F_{\tilde{\mathbf{x}}_i^{\infty}}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{14}$$

Convergence of the random allocations in the product space implies the pointwise convergence of the CDFs. That is, for the sequence $\tilde{\mathbf{x}}^{(\delta_k)} \to \tilde{\mathbf{x}}$, it holds for all agents i

that $F_{\tilde{\mathbf{x}}_i^{\delta_k}}(\mathbf{z})$ converges pointwise to $F_{\tilde{\mathbf{x}}_i^{\infty}}(\mathbf{z})$:

$$\lim_{k\to\infty} F_{\tilde{\mathbf{x}}^{(\delta_k)}}(\mathbf{z}) = F_{\tilde{\mathbf{x}}^{\infty}}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{15}$$

Similarly, for our δ_k-feasible sequence of partial allocations $\tilde{\mathbf{Y}}^{\delta_k} = (\tilde{\mathbf{y}}_i^{\delta_k})_{i\in S}$ that converge to $\tilde{\mathbf{Y}}$, we get that $F_{\tilde{\mathbf{y}}^{(\delta_k)}}(\mathbf{z})$ converges pointwise to $F_{\tilde{\mathbf{y}}}(\mathbf{z})$:

$$\lim_{k\to\infty} F_{\tilde{\mathbf{y}}^{(\delta_k)}}(\mathbf{z}) = F_{\tilde{\mathbf{y}}}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{16}$$

Since $\tilde{\mathbf{y}} \succeq_i^{sd} \tilde{\mathbf{x}}^{\infty}$ for all agents and strictly for at least one agent i, it holds that $F_{\tilde{\mathbf{y}}_i}(\mathbf{z}) \leq F_{\tilde{\mathbf{x}}_i^{\infty}}(\mathbf{z})$ for all bundles z and there exists an acceptable bundle $\mathbf{z}'$ such that $F_{\tilde{\mathbf{y}}_i}(\mathbf{z}') \leq F_{\tilde{\mathbf{x}}_i^{\infty}}(\mathbf{z}')$.

The pointwise convergence of CDFs and the characterization of stochastic dominance via CDFs imply that for sufficiently small δ_k, the sequence of δ_k-feasible partial allocations $\tilde{\mathbf{Y}}^{\delta_k}$ also stochastically dominates the respective original allocation $\tilde{\mathbf{X}}^{\delta_k}$ for agents in S: That is $\tilde{\mathbf{y}}_i^{\delta_k} \succ_i^{sd} \tilde{\mathbf{x}}_i^{\delta_k}$ for all $i \in N$, and $\tilde{\mathbf{y}}_i^{(\delta_k)} \succ_i^{sd} \tilde{\mathbf{x}}_i^{(\delta_k)}$ for at least one agent i.

This contradicts the fact that $\tilde{\mathbf{X}}^{\delta_k}$ is in the ordinal δ_k-core, and thus not admit a δ_k-feasible blocking coalition. Thus, the limit allocation $\tilde{\mathbf{X}}^{\infty}$ does not admit an ordinal blocking coalition.

Individual Rationality. It follows from Proposition 2 that any δ-CERE is individually rational. That is, along the sequence, it holds for every agent i that $\tilde{\mathbf{x}}_i^{\delta} \succeq_i^{sd} \tilde{\mathbf{e}}_i$. An analogous convergence argument, similar to above, proves that the limit $\tilde{x}_i$ must be individually rational as well. Otherwise individual rationality would be violated for sufficiently small $\delta_k > 0$, using again the pointwise convergence of CDFs. □

A.4 Proof of Theorem 3

Proof. We show that any allocation in the competitive core satisfies the condition of no justified envy based on set-inclusion and equilibrium valuations. Both of these properties imply the equal treatment of equals. Our proof proceeds in two steps: Any competitive core allocation is an asymptotic competitive equilibrium from random endowments, which is itself a limit of a sequence of δ_k-CEREs, denoted as $\tilde{\mathbf{X}}^{\delta_k}$, with corresponding price vectors $\mathbf{p}^{\delta_k}$. First, we prove that each δ_k-CERE satisfies the corresponding fairness property. Second, we show that this property extends to the limit.

Justified Envy-Freeness based on Set-Inclusion. Consider any δ-CERE $\tilde{\mathbf{X}}^{\delta} = (\tilde{\mathbf{x}}_1^{\delta}, ..., \tilde{\mathbf{x}}_n^{\delta})$. Let $(\tilde{\mathbf{x}}_i^{\delta}, \tilde{\mathbf{x}}_j^{\delta})$ be the random bundles consumed by agents i and j,

when they have endowments $\tilde{\mathbf{e}}_i$ and $\tilde{\mathbf{e}}_j$ respectively, with $\tilde{\mathbf{e}}_i \succeq^{si} \tilde{\mathbf{e}}_i$. We show that it must hold that $\tilde{\mathbf{x}}_i^\delta \succeq_i^{sd} \tilde{\mathbf{x}}_j^\delta$. For a given price vector $\mathbf{p}$, let $\tilde{\mathbf{B}}_i$ denote the *random budget of agent* i, given their random endowment and budget relaxation:

$$\tilde{\mathbf{B}}_i = \tilde{\mathbf{e}}_i \cdot \mathbf{p} + \tilde{\mathbf{b}}_i \cdot \mathbf{c} \cdot \mathbf{p}. \tag{17}$$

Similarly, let $\tilde{\mathbf{B}}_j$ denote the random budget of agent j, Because it holds that $\tilde{\mathbf{e}}_i \succeq^{si} \tilde{\mathbf{e}}_i$ and random budget relaxations are identically distributed, the random variable $\tilde{\mathbf{B}}_i$ first order stochastically dominates $\tilde{\mathbf{B}}_j$. Let $\mathbf{B}_i$ be a realization of the random budget.

Let u_i be any cardinal valuation of bundles that is consistent with the ordinal preference relation over bundles $\succeq_i$. Denote by $\mathbf{x}_{(\mathbf{p},\mathbf{B}_i,\succeq_i)}$ one of the best affordable bundles under price $\mathbf{p}$, realization of the random budget $\mathbf{B}_i$, and preference relation $\succeq_i$. The following two conditions trivially hold for realizations $\mathbf{B}_i' \geq \mathbf{B}_i$:

$$u_i(\mathbf{x}_{(\mathbf{p},\mathbf{B}_i',\succeq_i)}) \geq u_i(\mathbf{x}_{(\mathbf{p},\mathbf{B}_i,\succeq_i)}) \tag{18}$$

$$u_i(\mathbf{x}_{(\mathbf{p},\mathbf{B}_i,\succeq_i)}) \geq u_i(\mathbf{x}_{(\mathbf{p},\mathbf{B}_i,\succeq_j)}) \tag{19}$$

Combining these two observations with the first order stochastic dominance of the random budget $\tilde{\mathbf{B}}_i$ over $\tilde{\mathbf{B}}_j$ yields the following: Taking the expectation over the realization of initial endowments and budget realizations yields that

$$\mathbb{E}_{\mathbf{x}\sim\tilde{\mathbf{x}}_i^\delta}[u_i(\mathbf{x}])] \geq \mathbb{E}_{\mathbf{x}\sim\tilde{\mathbf{x}}_j^\delta}[u_i(\mathbf{x})]. \tag{20}$$

Thus, the expected utility of $\tilde{\mathbf{x}}_i^\delta$ is at least the one of $\tilde{\mathbf{x}}_j^\delta$ for every cardinal valuation consistent with $\succeq_i$, which is equivalent to first order stochastic dominance $\tilde{\mathbf{x}}_i^\delta \succeq_i^{sd} \tilde{\mathbf{x}}_j^\delta$. Hence, for $\delta > 0$, $\tilde{\mathbf{X}}^\delta$ leads to no justified envy based on set-inclusion.

Next, we want to show that this property extends from a sequence of δ_k-CEREs $\tilde{\mathbf{X}}^{\delta_k}$ to the limit, the competitive core allocation $\tilde{\mathbf{X}}$. Suppose again that $\tilde{\mathbf{e}}_i \succeq^{si} \tilde{\mathbf{e}}_j$. We want to show that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$ in the limit. Recall that stochastic dominance $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$ can be expressed in terms of CDFs, see Equation (14). Since $\tilde{\mathbf{x}}_i^\delta \to \tilde{\mathbf{x}}_i$ as $\delta \to 0$, the CDFs $F_{\tilde{\mathbf{x}}_i^{\delta_k}}(\mathbf{z})$ converge pointwise to $F_{\tilde{\mathbf{x}}_i}(\mathbf{z})$, see Equation (15). Given that we just proved that $\tilde{\mathbf{x}}_i^{\delta_k} \succeq_i^{sd} \tilde{\mathbf{x}}_j^{\delta_k}$ for all $\delta_k > 0$, we have

$$F_{\tilde{\mathbf{x}}_i^{\delta_k}}(\mathbf{z}) \leq F_{\tilde{\mathbf{x}}_j^{\delta_k}}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{21}$$

Taking the limit $k \to \infty$ as $\delta_k \to 0$, we get

$$\lim_{k\to\infty} F_{\tilde{\mathbf{x}}_i^{\delta_k}}(\mathbf{z}) \leq \lim_{k\to\infty} F_{\tilde{\mathbf{x}}_j^{\delta_k}}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{22}$$

By the pointwise convergence of the CDFs, this implies

$$F_{\tilde{\mathbf{x}}_i}(\mathbf{z}) \leq F_{\tilde{\mathbf{x}}_j}(\mathbf{z}) \text{ for all bundles } \mathbf{z}. \tag{23}$$

Therefore, $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$ for all $i, j \in N$ with $\tilde{\mathbf{e}}_i \succeq^{si} \tilde{\mathbf{e}}_j$. Thus, any allocation in the competitive core ensures no justified envy based on set-inclusion.

No justified envy based on equilibrium-valuations. Consider any δ-CERE $\tilde{\mathbf{X}}^\delta = (\tilde{\mathbf{x}}_1^\delta, ..., \tilde{\mathbf{x}}_n^\delta)$ with price vector $\mathbf{p}^\delta$. Let $(\tilde{\mathbf{x}}_i^\delta, \tilde{\mathbf{x}}_j^\delta)$ be the random bundles consumed by agents i and j. Suppose that endowments $\tilde{\mathbf{e}}_i$ and $\tilde{\mathbf{e}}_j$ are such that $\mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}^\delta} \mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_j)$. We show that it must hold that $\tilde{\mathbf{x}}_i^\delta \succeq_i^{sd} \tilde{\mathbf{x}}_j^\delta$. For the equilibrium price vector $\mathbf{p}^\delta$, let $\tilde{\mathbf{B}}_i$ and $\tilde{\mathbf{B}}_j$ denote again the random budgets of agents i and j, see Equation (17). Because it holds that $\mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}^\delta} \mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_j)$ and random budget relaxations are identically distributed, $\tilde{\mathbf{B}}_i$ first-order stochastically dominates $\tilde{\mathbf{B}}_j$. Hence, the exact same argument as above for justified envy based on set-inclusion yields that the expected utility of $\tilde{\mathbf{x}}_i^\delta$ is at least the one of $\tilde{\mathbf{x}}_j^{\delta,f}$ for every cardinal valuation consistent with $\succeq_i$. This implies that $\tilde{\mathbf{x}}_i^\delta \succeq_i^{sd} \tilde{\mathbf{x}}_j^\delta$. Thus, for every $\delta > 0$, $\tilde{\mathbf{X}}^\delta$ leads to no justified envy based on $\mathbf{p}^\delta$-valuations.

Next, we want to show that this property extends from a sequence of δ_k-CEREs $\tilde{\mathbf{X}}^{\delta_k}$ to the limit, the competitive core allocation $\tilde{\mathbf{X}}$ with associated price vector $\mathbf{p}$. Suppose that endowments $\tilde{\mathbf{e}}_i$ and $\tilde{\mathbf{e}}_j$ are such that $\mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}} \mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_j)$. We show that it must hold that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$. First, we show that $\mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}} \mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_j)$ implies that for sufficiently small δ_k, we also have that

$$\mathbf{v}_{\mathbf{p}^{\delta_k}}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}^{\delta_k}} \mathbf{v}_{\mathbf{p}^{\delta_k}}(\tilde{\mathbf{e}}_j). \tag{24}$$

Since $\tilde{\mathbf{e}}_i \succeq^{\mathbf{p}} \tilde{\mathbf{e}}_j$, we have:

$$\mathbb{P}(\mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i) \geq z) \geq \mathbb{P}(\mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_j) \geq z) \text{ for all bundles } \mathbf{z}. \tag{25}$$

The continuity of valuations with respect to the price vector $\mathbf{p}$ implies that as $\mathbf{p}^{\delta_k} \to \mathbf{p}$, the random valuations converge in distribution:

$$\mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_i) \xrightarrow{d} \mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i) \text{ and } \mathbf{v}_{\mathbf{p}^\delta}(\tilde{\mathbf{e}}_j) \xrightarrow{d} \mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_j). \tag{26}$$

Therefore, for sufficiently small $\delta_k > 0$, we have

$$\mathbb{P}(\mathbf{v}_{\mathbf{p}^{\delta_k}}(\tilde{\mathbf{e}}_i) \geq z) \geq \mathbb{P}(\mathbf{v}_{\mathbf{p}^{\delta_k}}(\tilde{\mathbf{e}}_j) \geq z) \text{ for all bundles } \mathbf{z}, \tag{27}$$

from which Equation (24) follows. Let $\tilde{\mathbf{x}}_i^\delta$ and $\tilde{\mathbf{x}}_j^\delta$ be the corresponding allocations for sufficiently small $\delta_k > 0$. We know from the previous step that: $\tilde{\mathbf{x}}_i^{\delta_k} \succeq_i^{sd} \tilde{\mathbf{x}}_j^{\delta_k}$. Moreover, the allocations $\tilde{\mathbf{x}}_i^\delta$ and $\tilde{\mathbf{x}}_j^\delta$ converge to $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_j$ respectively, as $\delta_k \to 0$. The exact same convergence argument using CDFs that we used to show that the limit allocation leads to no justified envy based on set-inclusion now shows that $\tilde{\mathbf{x}}_i \succeq_i^{sd} \tilde{\mathbf{x}}_j$. Because $\mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i) \succeq^{\mathbf{p}} \mathbf{v}_{\mathbf{p}}(\tilde{\mathbf{e}}_i)$, it follows that any allocation in the competitive core ensures no justified envy based on $\mathbf{p}$-valuations. □

A.5 Proof of Theorem 4

Proof. Suppose that endowments are deterministic. By definition, any allocation in the competitive core is a marginal distribution over bundles from an asymptotic competitive equilibrium from random endowments in the augmented economy with artificial money.

For the remainder of the proof, we use endowments, allocations, and price vectors in the augmented economy. For example, the deterministic endowments including money will be denoted by $\mathbf{e} = (\mathbf{e}_1, ..., \mathbf{e}_n)$, a deterministic allocation will be denoted by $\mathbf{x} = (\mathbf{x}_1, ..., \mathbf{x}_n)$, and a price vector will be denoted as $\mathbf{p}$. Let $\mathbf{c}$ be the capacities of goods and money. We will interpret all of these as vectors in $\mathbb{R}^{m+1}$, such that the last component will correspond to the allocation or price of money. That is, we will write $\mathbf{x}_i = (\mathbf{x}_i^G, \mathbf{x}_i^M)$ for allocations and $\mathbf{p} = (\mathbf{p}^G, \mathbf{p}^M)$ for prices. Hence, the superscript ^{G} will be used to denote bundles and prices of goods in $\mathbb{R}^m$, while the superscript ^{M} will denote allocations and the price of money.

Realizability over near-feasible core allocations. We start by proving that a random allocation of goods $\tilde{\mathbf{X}}^G$ in the competitive core is realizable over deterministic and near-feasible core allocations. By definition, there exists a limit sequence of δ_k-CEREs $\mathbf{X}^{\delta_k}$ with associated price vectors $\mathbf{p}^{\delta_k}$ that converge to an asymptotic CERE $\tilde{\mathbf{X}}$, where the marginal over bundles is $\tilde{\mathbf{X}}^G$ with associated limit price vector $\mathbf{p}^G$.

Let $\mathbf{x}_i^G$ be a deterministic bundle, which is allocated to agent i with positive probability in the random allocation $\tilde{\mathbf{X}}^G$. We can find a sufficiently small $\delta > 0$, such that all such bundles of goods $\mathbf{x}_i^G$ are realized as a part δ-CERE $\tilde{\mathbf{X}}^\delta = (\tilde{\mathbf{X}}^{\delta,G}, \tilde{\mathbf{X}}^{\delta,M})$ with corresponding price vector $\mathbf{p}^\delta = (\mathbf{p}^{\delta,G}, \mathbf{p}^{\delta,M})$.

For every agent i, let $\mathbf{x}_i^\delta$ denote a realization of the δ-CERE, such that the allocated

bundle is $\mathbf{x}_i^{\delta,G} = \mathbf{x}_i^G$. Denote this deterministic allocation for all agents by $\mathbf{x}^\delta$.

As every realization of a δ-CERE corresponds to a specific realization of the random budget relaxation, let $(\mathbf{b}_1, ..., \mathbf{b}_n)$ denote such a vector that lead to $\mathbf{x}_i^\delta$. For agent i with allocation $\mathbf{x}_i^\delta$ in the δ-CERE, $\mathbf{x}_i^{\delta,G}$ is the optimal affordable bundle:

$$\mathbf{p}^{\delta,G} \cdot \mathbf{x}_i^{\delta,G} \leq \mathbf{p}^\delta \mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}^\delta \mathbf{c}. \tag{28}$$

and any bundle of goods, which agent i prefers to $\mathbf{x}_i^{\delta,G}$ is more expensive.

Let $\mathbf{c}^*$ be the aggregate consumption of the deterministic allocation $\mathbf{x}^\delta$. We next show that the allocation $\mathbf{x}^\delta$ is in the deterministic core with respect to the aggregate consumption $\mathbf{c}^*$. Suppose, that $\mathbf{x}^\delta$ is not in the deterministic core with respect to $\mathbf{c}^*$. Then, there exists a blocking coalition $S \subseteq N$ and a partial allocation $(\mathbf{y}_s)_{s \in S}$, such that for all $s \in S$, it holds that $\mathbf{y}_s \succeq_s \mathbf{x}_s^\delta$, there is at least one agent s', such that $\mathbf{y}_{s'} \succ_{s'} \mathbf{x}_{s'}^\delta$, and the partial allocation is feasible:

$$\sum_{s \in S} \mathbf{y}_s \leq \sum_{s \in S} \mathbf{e}_s. \tag{29}$$

In our δ-CERE, agents point to their favourite affordable and expenditure-minimizing bundle. Hence, it must hold that

$$\mathbf{p}^\delta \mathbf{y}_s \geq \mathbf{p}^\delta \mathbf{e}_s + \mathbf{b}_s \cdot \mathbf{p}^\delta \cdot \mathbf{c}, \tag{30}$$

with strict inequality for at least one agent s'. Adding up these inequalities yields

$$\sum_{s \in S} \mathbf{p}^\delta \mathbf{y}_s > \sum_{s \in S} \left(\mathbf{p}^\delta \mathbf{e}_s + \mathbf{b}_s \cdot \mathbf{p}^\delta \cdot \mathbf{c} \right). \tag{31}$$

This is a contradiction, as Equations (29) and (31) cannot hold at the same time. Hence, the deterministic allocation $\mathbf{x}^\delta$ is in the core with respect to the capacities $\mathbf{c}^*$ from aggregate consumption, as it is feasible with respect to $\mathbf{c}^*$ and cannot be blocked.

Next, we argue that this implies that the deterministic allocation $\mathbf{x}^{\delta,G}$ (without artificial money), is also in the core of the economy without money. Suppose not, this would mean, that there exists a feasible blocking coalition based on bundles of goods alone. However, this could directly be extended to a blocking coalition in the economy with artificial money, where everyone keeps their allocated money, as money does not change the preference order over bundles.

Finally, we use the following results from Nguyen et al. (2016, Theorem 2.1):

Any feasible random allocation with respect to $\mathbf{c}$ can be realized through

deterministic allocations that are feasible with respect to $\mathbf{c} + (\Delta - 1) \cdot \mathbf{1}$.

As a consequence, we can realize the random allocation $\tilde{\mathbf{X}}^G$ such that each realized aggregate consumption satisfies $\mathbf{c}' \leq \mathbf{c} + (\Delta - 1) \cdot \mathbf{1}$. This realization corresponds to a δ-CERE allocation for sufficiently small δ and budget realizations and is, as we have just proven, in the core with respect to capacities that equal realized aggregate consumption $\mathbf{c}'$. This proves that any random allocation in the competitive core is realizable over $(\Delta - 1)$-near feasible deterministic core allocations.

Justified envy-freeness up to one good. Finally, we show the stronger result, that is, the $(\Delta-1)$-near feasible core allocations can actually be chosen to be JEF1 based on set-inclusion and equilibrium valuations. Given the realizability over near-feasible deterministic core allocations, we write $\tilde{\mathbf{X}}^G = \sum_{l=1}^{L} \pi_l \mathbf{X}_l^G$, where π_i is the probability with which the deterministic near-feasible core allocation $\mathbf{X}_l^G$ is realized. We consider one representative deterministic allocation, and for notational simplicity, we write $\mathbf{X}^G$.

Because this allocation $\mathbf{X}^G$ is realized with positive probability from a competitive core allocation, for sufficiently small $\delta > 0$, it corresponds to the marginal realization of goods from a δ-CERE $\tilde{\mathbf{X}}^\delta$ with corresponding price vector $\mathbf{p}^\delta$, in the augmented economy with artificial money.

As proven above Equation (28), there exist realizations of random budget relaxations $\mathbf{b}_1, ..., \mathbf{b}_n \leq \delta$, such that $\mathbf{X}^G$ is part of the realization $\mathbf{X}^\delta$ of a δ-CERE for a price vector $\mathbf{p}^\delta$. That is, $\mathbf{X}^G = \mathbf{X}^{\delta,G}$. Recall that $\mathbf{p}^\delta = (\mathbf{p}^{\delta,G}, \mathbf{p}^{\delta,M})$.

For agents i and j, we write $\mathbf{x}_i^G$ and $\mathbf{x}_j^G$ for their allocated bundle of goods in that deterministic allocation $\mathbf{X}^G$. Because this allocation $\mathbf{X}^G$ stems from δ-CERE with realized budget relaxations $\mathbf{b}_i$ and $\mathbf{b}_j$, it holds that both agents can afford their allocated bundles:

$$\mathbf{p}^{\delta,G} \cdot \mathbf{x}_i^G \leq \mathbf{p}^\delta \cdot \mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}^\delta \cdot \mathbf{c} \text{ and } \mathbf{p}^{\delta,M} \cdot \mathbf{x}_j^G \leq \mathbf{p}^\delta \cdot \mathbf{e}_j + \mathbf{b}_j \cdot \mathbf{p}^\delta \cdot \mathbf{c}. \tag{32}$$

As a shorthand notation, we call $\mathbf{B}_i = \mathbf{p}^\delta \cdot \mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{p}^\delta \cdot \mathbf{c}$ the *budget* of agent i, and we write $\mathbf{x}_i^G \setminus \{k\}$ to denote the bundle $(\mathbf{x}_i^G - e^k)^+$, where one unit of good k is removed, if there is at least one unit present.

Suppose that envy of agent i towards agent j cannot be resolved by substracting one unit of a good. That is, for all $k \in M$, it holds that $\mathbf{x}_i^G \prec_i \left(\mathbf{x}_j^G \setminus \{k\}\right)$. Given the competitive market foundation of the δ-CERE, this implies that the bundle $\mathbf{x}_j^G \setminus \{k\}$ must not be affordable to agent i:

$$\mathbf{p}^{\delta,G} \cdot \left(\mathbf{x}_j^G \setminus \{k\}\right) > \mathbf{B}_i. \tag{33}$$

That is, because agent i consumed their favourite affordable and expenditure-minimizing bundle within their budget. If they strictly prefer $\mathbf{x}_j^G \setminus \{k\}$ to $\mathbf{x}_i^G$, then the price of the former must have exceeded their budget. Summing Equation (33) over all goods k yields

$$\sum_{k=1}^{m} \mathbf{p}^{\delta,G} \cdot \mathbf{x}_j^G \setminus \{k\} > m \cdot \mathbf{B}_i. \tag{34}$$

A simple counting argument proves that the left-hand side of this equation is upper bounded by $(m-1)\mathbf{p}^{\delta,G} \cdot \mathbf{x}_j^G$. Moreover, the bundle $\mathbf{x}_j^G$ is affordable to agent j, and hence $\mathbf{B}_j \geq \mathbf{p}^{\delta,G} \cdot \mathbf{x}_j^G$. Combining these bounds yields that $\frac{m-1}{m} \cdot \mathbf{B}_j > \mathbf{B}_i$. Thus, the budget of agent i is at most $\frac{m-1}{m}$ times the budget of agent j. Putting our expression for the budget of an agent back in, we get the following inequality:

$$\frac{m-1}{m} \mathbf{p}^{\delta} \left(\mathbf{e}_j + \mathbf{b}_j \cdot \mathbf{c}\right) > \mathbf{p}^{\delta} \left(\mathbf{e}_i + \mathbf{b}_i \cdot \mathbf{c}\right). \tag{35}$$

For the final contradiction, we now need to consider, whether this envy is justified.

JEF1 based on set-inclusion. To boundjustified envy based on set-inclusion, consider that we have the pointwise bound: $\mathbf{e}_i^G \geq \mathbf{e}_j^G$. This implies that also $\mathbf{e}_i \geq \mathbf{e}_j$, as in the augmented economy, all agents receive the same amount of artificial money. This implies

$$\mathbf{p}^{\delta} \cdot \mathbf{e}_i \geq \mathbf{p}^{\delta} \cdot \mathbf{e}_j. \tag{36}$$

As the budget relaxations $\mathbf{b}_i, \mathbf{b}_j \leq \delta$, it follows that for sufficiently small $\delta > 0$, Equations (35) and (36) cannot hold at the same time. This contradicts the assumption, that envy was not bounded by a single good, proving that the representative deterministic allocation $\mathbf{X}^G$ is JEF1 for set-inclusion.

JEF1 based on equilibrium valuations. To bound justified envy based on equilibrium valuations, consider the equilibrium price vector $\mathbf{p}^G$, and suppose that endowments are such that the equilibrium valuation of $\mathbf{e}_i^G$ is great or equal than the equilibrium valuation of $\mathbf{e}_j^G$. That is, $\mathbf{p}^G \cdot \mathbf{e}_i^G \geq \mathbf{p}^G \cdot \mathbf{e}_j^G$. This implies that $\mathbf{p} \cdot \mathbf{e}_i \geq \mathbf{p} \cdot \mathbf{e}_j$, where $\mathbf{p}$ is the limit price vector including the price of artificial money, because in the augmented economy, all agents receive the same amount of artificial money. Recall that $\mathbf{b}_i, \mathbf{b}_j \leq \delta$. Moreover, $\delta > 0$ can be chosen small enough, such that $\mathbf{p}^{\delta}$ is arbitrarily close to $\mathbf{p}$, and hence also such that $\mathbf{p}^{\delta} \cdot \mathbf{e}_i$ is arbitrarily close to $\mathbf{p} \cdot \mathbf{e}_i$. That is, for all $\epsilon > 0$, we can

choose δ small enough such that

$$\mathbf{p}^{\delta} \cdot \mathbf{e}_i \geq \mathbf{p}^{\delta} \cdot \mathbf{e}_j - \epsilon. \tag{37}$$

Again, for sufficiently small $\delta > 0$, Equations (35) and (37) cannot hold at the same time. This contradicts the assumption, that envy was not bounded by a single good, proving that the representative deterministic allocation $\mathbf{X}$ is JEF1 for equilibrium-valuations. □

A.6 Proof of Theorem 5

Proof. Consider an economy with n agents and fix any cardinal representation $(u_t)_{t \in T^*}$ of weak ordinal preferences. By Theorem 3, the Competitive Core Mechanism, denoted by $\Phi^{CCM}(\cdot)$, satisfies the ex-ante property of equal treatment of equals. Specifically, for any reported type vector t and for any pair of agents i and j with types t_i and t_j such that $\tilde{\mathbf{e}}_i = \tilde{\mathbf{e}}_j$, it holds that

$$\Phi_i^{CCM}(t) \succeq_i^{sd} \Phi_j^{CCM}(t) \tag{38}$$

This implies that for the corresponding Von Neumann-Morgenstern utility functions u_t, we have $\mathbb{E}[u_{t_i}(\Phi_i^{CCM}(t)] \geq \mathbb{E}[u_{t_i}(\Phi_j^{CCM}(t)]$. This condition is precisely the envy-freeness condition for semi-anonymous direct mechanisms as introduced in Appendix C of Azevedo and Budish (2019). In other words, no agent envies another agent within the same group; in our context, a group is defined by equal endowment. Thus, the condition of envy-freeness within a group for a fixed utility function follows directly from our concept of ordinal equal treatment of equals. Therefore, by the extension of Azevedo and Budish (2019, Theorem 1) to semi-anonymous mechanisms in their Appendix C, we conclude that the Competitive Core Mechanism is strategyproof in the large. Specifically, it follows directly from their analysis that the maximum possible gains from misreporting decrease at a rate of $n^{-\frac{1}{2}+\epsilon}$ as the number of agents n increases. □