Individual and Collective Welfare in Risk Sharing with Many States

Federico Echenique and Farzad Pourbabaee

January 14, 2024

Abstract

We provide a quantitative assessment of welfare in the classical model of risk-sharing and exchange under uncertainty. We prove three kinds of results. First, that in an equilibrium allocation, the scope for improving individual welfare by a given margin (an ε-improvement) vanishes as the number of states increases. Second, that the scope for a change in aggregate resources that may be distributed to enhance individual welfare by a given margin also vanishes. Equivalently: in an inefficient allocation, for a given level of resource sub-optimality (as measured by the coefficient of resource under-utilization), the possibilities for enhancing welfare by perturbing aggregate resources decrease exponentially to zero with the number of states. Finally, we consider efficient risk-sharing in standard models of uncertainty aversion with multiple priors, and show that, in an inefficient allocation, certain sets of priors shrink with the size of the state space.

Echenique (fede@econ.berkeley.edu) is at the Department of Economics, UC Berkeley. Pourbabaee (far@caltech.edu) is at the Division of the Humanities and Social Sciences, Caltech.

Contents

1 Introduction

We provide a *quantitative assessment* of individual and collective welfare in models of risk-sharing, when the state space is large. Our model is completely standard: We consider an exchange economy, with uncertainty, populated by a collection of risk-averse agents who trade to share risk. Uncertainty is captured with a finite, but large, state space. The welfare theorems apply, and Walrasian equilibrium allocations are Pareto optimal; but we are interested in a quantitative assessment of welfare in these economies as the number of states grows large.

Our main results are threefold.

First we consider individual welfare in a Walrasian equilibrium allocation. We quantify welfare by the probability that a random perturbation of an agent's equilibrium consumption would provide a strict utility improvement: specifically, an improvement of "at least ε" over equilibrium consumption; meaning that the perturbation would remain an improvement if we were to shave off a fraction ε from consumption. Our first result is actually more general, and quantifies the perturbations that could make at least one agent better off by at least ε, but we focus on a simple special case in this introduction. Now, it is obvious that an agent's consumption can be improved whenever their preferences are monotonic, and the probability of an improvement by perturbation is always strictly positive. *Our first result says that the probability that a perturbation is an individual ε-improvement converges (exponentially) to zero as the number of states grows.*

To understand our first result, consider the situation of an individual agent i in equilibrium. The usual consumption-maximization diagram is reproduced in Figure 1 on the left. In the figure, there are two states of the world ($d = 2$), and an agent's budget set is depicted in gray. The budget is defined by an equilibrium price p, and the agent's endowment ω_i. The picture shows the agent's equilibrium consumption f_i. The agent's indifference curve, tangent to the budget line at f_i, is depicted as a dashed curve. The *upper contour set* is the region to the northeast of the indifference curve that contains all consumptions that are preferred to f_i.

We wish to assess the size of the upper contour set at f_i: what is the scope for improving i's welfare at the equilibrium allocation f? By virtue of the monotonicity of i's preference, the upper contour set is, of course, infinite. It contains all the consumption vectors that are larger, state by state, than f. So we consider its size relative to a ball that is centered at f_i. Specifically, we measure the size of the dotted region on the right panel of Figure 1. Think of

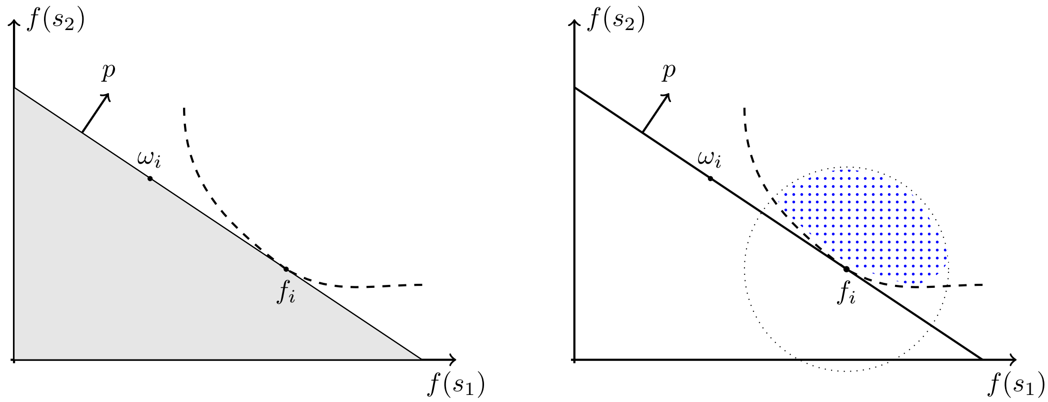

Figure 1: Individual welfare.

the ball as a set of perturbations of f_i.[1] We wish to quantify how often such perturbations would lead to a utility improvement for i. Key to our result is that the improvement should be at least by ε; for simplicity the figure only depicts the strictly better bundles (the case of $\varepsilon = 0$).

The consumer would be better off with more consumption in every state, but they would also be willing to accept a tradeoff of more consumption in one state than in another. Of course, the tradeoff would have to be at terms of trade that are more favorable than equilibrium prices. But the figure shows that there is a significant probability that a perturbation would leave the agent strictly better off, and the upper contour contains much more than the bundles that are larger state-by-state. So does the dotted region. Our first theorem says that, when the number of states is d, the probability of the dotted region, calculated for ε improvements, is bounded above by $\mathrm{e}^{-\varepsilon^2 d/8}$. In particular, it converges rapidly to zero as d grows large.

To put this bound in perspective, note that the fraction of bundles that are larger, state-by-state, than f_i also converges exponentially to zero. So our result says that, despite the potential for utility improvements by trading off more consumption in some states for less consumption in other, the scope for such potential improvement shrinks when the number of states is large. In a sense, *the curvature of utility does not matter in high-dimensional trade*. Our bound of $\mathrm{e}^{-\varepsilon^2 d/8}$ holds for any monotonic and quasi-concave utility function.

Our second main result concerns collective welfare in a Pareto optimal allocation. The exercise is similar to our first result, but relies on redistributing aggregate consumption so

[1]The ball drawn in the figure includes negative consumption, but its radius may be taken to be arbitrarily small so as to avoid negative consumption, if so desired. We may also restrict attention to consumption that is unaffordable at equilibrium prices (the region to the northeast of the budget line) without changing the essential message of our result.

as to make all agents better off. For our second result, we consider an economy without aggregate uncertainty, and fix an allocation that is Pareto optimal. This time we consider a perturbation of aggregate consumption that can be redistributed to make all agents better off by at least ε — meaning again a utility improvement that remains, even after shaving off a fraction ε of consumption. We show that *a perturbation to aggregate consumption that can be redistributed to improve all agents by at least* ε *has vanishingly small probability.*

Pareto optimality refers to efficiency with a fixed aggregate endowment, or aggregate consumption. One cannot make all agents better off without changing the aggregate consumption. Our result says, when the state space is large, the scope for collective welfare improvements by a given margin ε is limited, even by changing aggregate consumption. Again it is useful to compare our bound with the fraction of aggregate allocations that are larger than the endowment, state-by-state. The fraction larger allocations shrinks exponentially to zero, as does (according to our results), the allocation that afford a collective improvement.

Our second result can be interpreted in light of Debreu's coefficient of resource utilization (Debreu, 1951). If we consider an allocation that is *not* Pareto optimal, then Debreu's coefficient of resource utilization measures the degree of inefficiency inherent in the allocation. We can again study the probability of an improving perturbation to aggregate consumption that makes agents better off (not by ε this time, simply strictly better off). Our second main result implies that this probability shrinks to zero exponentially as the number of states grows large. If Debreu's coefficient of resource utilization is CRU, and Pareto inefficiency implies that CRU < 1, then the probability is bounded above by $\mathrm{e}^{-(1-\mathrm{CRU})^2 d/8}$.[2]

Our third main result is not explicitly about welfare, but instead about the attitudes towards uncertainty by agents who engage in efficient risk sharing. We consider the same setting of economic exchange with no aggregate uncertainty, as in our second result; but we strengthen the assumptions on preferences to focus on utilities with multiple priors. Think, to fix ideas, on the max-min expected utility preferences of Gilboa and Schmeidler (1989). Now we prove that, if the risk sharing agreement between the two agents can be improved in the Pareto sense, then at least one of the two must have a small set of prior beliefs. The stronger is the level of Pareto improvement, the smaller will be the size of the prior beliefs set. This means that at least one of the two agents must, in some sense, be close to being ambiguity neutral.

[2] Debreu (1951) proposes CRU $\in (0,1]$ as a measure of inefficiency of an allocation, and shows that CRU $= 1$ if and only if the allocation is Pareto optimal.

Our results rely on the study of concentration of measures, and specifically on the implication of an *Isoperimetric inequality* to the separation of convex sets. Key is the role of ε, or of CRU, in our discussion above. These magnitudes provide a certain "padding," or quantitative bound, on the degree of separation of the relevant convex sets. The phenomenon of concentration of measure then has strong consequences for the volume of one of the sets being separated, and in our economic models we can determine which of the sets must have small volume (except for our third application, which is more subtle). Given its role in our results, we discuss the magnitude of ε in Section 4.2.

2 The Model

2.1 Notations and Conventions

Before presenting our model and main results, we lay down some of the basic definitions we shall make use of, as well as a few notational conventions.

Let A be a subset of a finite-dimensional normed vector space $(\mathbb{R}^m, \|\cdot\|)$. The distance of a point $x \in \mathbb{R}^m$ from A is defined by $\mathsf{dist}(x, A) := \inf_{a\in A}\|x-a\|$. By default, the norms used in the paper are the Euclidean ℓ_2 norm. When a particular p-norm is used, we refer to the distance function by dist_p and the norm by $\|\cdot\|_p$. For the ℓ_2 norm, we omit the subscript 2.

We represent the Euclidean open ball centered at c and with radius r by $\mathbb{B}(c, r) := \{x \in \mathbb{R}^d_+ : \|x-c\| < r\}$. When the center is omitted, we take it to be the null vector $\mathbf{0}$. When the radius is omitted, we assume it is $r = 1$. Thus $\mathbb{B}$ denotes the standard (open) unit ball. In the same manner, we denote the *uniform* probability law on $\mathbb{B}(r)$ by P^r.

For two subsets $A, B \subseteq \mathbb{R}^m$, we define $\mathsf{dist}(A, B) := \inf\{\|a-b\| : a \in A, b \in B\}$. We refer to the δ-*extension* of the subset A by $A^\delta = \{x : \mathsf{dist}(x, A) < \delta\}$, that coincides with $A+\delta\,\mathbb{B}$.

Given a measurable subset $A \subseteq \mathbb{R}^m$, we denote its Euclidean volume by $\mathsf{Vol}(A)$ that is equal to the Lebesgue integral of the indicator function of A relative to the affine hull of A. For example, if A is a $m-1$ dimensional surface in $\mathbb{R}^m$, then $\mathsf{Vol}(A)$ refers to the surface area of A, as opposed to its m dimensional volume (which is zero).

If S is a finite set, we denote by $\Delta S = \{\mu : S \mapsto \mathbb{R}_+ | \sum_{s\in S}\mu(s) = 1\}$ the set of all probability measures on S. We embed ΔS as a subset of $\mathbb{R}^d_+$, and sometimes we refer to this probability simplex by Δ_d.

2.2 Preferences and Uncertainty

We consider a setting with uncertainty.[3] Let S be a finite set of possible *states of the world*, and d be the number of states of the world, i.e., $d := |S|$.

Consequences are evaluated based on their monetary payoff in $\mathbb{R}_+$. The set of *acts*, mappings from $S \to \mathbb{R}_+$, is denoted by $\mathbb{R}^d_+$, and individual acts are denoted by f and g. We endow this set with its natural topology.

Let $\succeq$ be a binary relation on $\mathbb{R}^d_+$. As usual, we denote the strict part of $\succeq$ by $\succ$, and the associated indifference relation by $\sim$. We say that $\succeq$ is a (weakly monotone) *preference relation* if it satisfies the following properties:

– (Weak Order): $\succeq$ is complete and transitive.

– (Continuity): The upper and lower contour sets are closed. That is for every $f \in \mathbb{R}^d_+$, the sets $\{g : g \succeq f\}$ and $\{g : f \succeq g\}$ are closed.

– (Monotonicity): For all $f, g \in \mathbb{R}^d_+$ if $f(s) \geqslant g(s)$ for all $s \in S$, then $f \succeq g$. Furthermore, if $f(s) > g(s)$ for all $s \in S$, then $f \succ g$.

The space of such preference relations on $\mathbb{R}^d_+$ is denoted by $\mathcal{P}$.

Given a preference $\succeq$ and an act f, the set $\{g : g \succeq f\}$ is called the *upper contour set* of $\succeq$ at f. We say that a preference $\succeq$ is *convex* if its upper contour sets are convex, for all acts f. Convexity jointly with the weak order property and continuity imply that the *strict* upper contour set, denoted by $\mathcal{U}^{(0)}_{\succeq} := \{g : g \succ f\}$, is also convex. We refer to the space of convex preferences by $\mathcal{C} \subset \mathcal{P}$.

Many well-known models in the theory of choice under uncertainty are special cases of convex preferences. Examples are risk-averse subjective expected utility (SEU), max-min expected utility (MEU: Gilboa and Schmeidler, 1989), Multiplier preferences (Hansen and Sargent, 2001), variational preferences (Maccheroni et al., 2006), and smooth ambiguity aversion (Klibanoff et al., 2005).

An approximate notion of upper contour sets is key to our results.

Definition 1 (ε-upper contour set)**.** The approximate upper contour set of preference $\succeq$ at the act f is defined by

$$\mathcal{U}^{(\varepsilon)}_{\succeq}(f) = \left\{g \in \mathbb{R}^d_+ : (1-\varepsilon)g \succ f\right\} .$$

[3] Using the reinterpretation of commodities in Debreu (1959), our results have implications for many other economic environments. For example, the textbook version of the model assumes that consumption is in units of physically distinct goods.

By convexity and continuity of preferences the approximate upper contour set is a convex and open subset. Its convexity is of particular importance to us, as our arguments often rely on the hyperplane separation theorem.

Approximate optimality is often expressed by means of the ε-maximization of some numerical objective function — a *utility function* u representing a preference $\succeq$. Observe that when $\succeq$ is homothetic (which is the case for many preferences used in applications), then u may be taken to be homogenous of degree one. As a consequence, our notion of approximate optimality translates directly into an approximation of utilities. Indeed, we may then write $\mathcal{U}^{(\varepsilon)}_{\succeq}(f) = \{g \in \mathbb{R}^d_+ : (1-\varepsilon)u(g) > u(f)\}$.

2.3 Exchange Economies

The set of agents is denoted by I and a typical member is referenced by indices i, j, and k. We assume throughout that I is finite.

An *exchange economy* is a mapping $\mathcal{E} : I \to \mathcal{P} \times \mathbb{R}^d_+$, where $\mathcal{E}(i) = (\succeq_i, \omega_i)$. Each agent $i \in I$ is described by a preference relation $\succeq_i$ on $\mathbb{R}^d_+$, as well as an *endowment vector* $\omega_i \in \mathbb{R}^d_+$. In an exchange economy, we use $\mathcal{U}^{(\varepsilon)}_i$ to denote the upper contour set $\mathcal{U}^{(\varepsilon)}_{\succeq_i}$.

Given an exchange economy $\mathcal{E}$, the *aggregate endowment* is $\omega := \sum_{i \in I} \omega_i$.[4] A profile of acts across agents, say $f = \{f_i : i \in I\} \in \mathbb{R}^{d \times I}_+$, is called an *allocation* if $\sum_{i \in I} f_i = \omega$. The space of all allocations is denoted by $\mathcal{F}_\omega$.

Definition 2 (ε-Pareto optimality)**.** An allocation $f \in \mathcal{F}_\omega$ is called ε-Pareto optimal if there is no allocation $g \in \mathcal{F}_\omega$ such that $g_i \in \mathcal{U}^{(\varepsilon)}_i(f_i)$ for all $i \in I$.

In words, an allocation f in $\mathcal{E}$ is ε-Pareto optimal if there is no redistribution of the aggregate endowment $\omega = \sum_i f_i$ that would be strictly better for all agents, and that would remain strictly better for all agents after a fraction ε of consumption is "shaved off" in each state of the world. When $\varepsilon = 0$, the definition coincides with the usual notion of *weak Pareto optimality*. Further properties of the set of ε-Pareto optimal allocations are discussed in Appendix A.

Definition 3 (Walrasian equilibrium)**.** An allocation $f = \{f_i : i \in I\}$ is called a Walrasian equilibrium for the exchange economy $\mathcal{E}$, if there exists a price vector $p \in \mathbb{R}^d$ such that $g_i \succ_i f_i$ implies that $p \cdot g_i > p \cdot \omega_i$, and $p \cdot f_i = p \cdot \omega_i$, for every $i \in I$.

[4] In the textbook Walrasian setting, $\omega(s)$ represents the total available amount of good s in the economy, and in the Arrow-Debreu market it represents the total contingent amount that agents can receive in state s.

An exchange economy is *convex* if each preference relation $\succeq_i$ is convex, i.e., $\succeq_i \in \mathcal{C}$. Convexity is required for the basic theory of general equilibrium: existence of Walrasian equilibrium, as well as the second welfare theorem, relies on convex preferences. All of our results, except Proposition 1 and Theorem 3, will also rely on convexity

When $s \mapsto \sum_{i \in I} \omega_i(s)$ is constant, we say that $\mathcal{E}$ exhibits *no aggregate uncertainty.* The aggregate endowment is then same across all states of the world, i.e., $\omega = (\bar{\omega}, \ldots, \bar{\omega})$.

3 Main Results

We proceed to state our main results, leaving a broader discussion our findings to Section 4, and an overview of the methodology behind our proofs to Section 5.

3.1 Optimality and Equilibrium

Our first result is a generalization of the property that we discussed in the introduction. It entails the following exercise: Fix a Walrasian equilibrium allocation, and draw a random perturbation from a ball of radius r. Consider the event that this random perturbation, when applied to *any* agent's equilibrium consumption, results in contingent consumption that is better by at least ε than their equilibrium consumption. The resulting event, that at least one agent is made better off, has vanishingly small probability.

Theorem 1. *Let $\mathcal{E}$ be a convex exchange economy. Suppose that there exists $\tau > 0$ such that $\omega_i \geqslant \tau \mathbf{1}$ for all $i \in I$, and let f be a Walrasian equilibrium allocation. Fix $r > 0$ and let $z \sim \mathsf{P}^r$. Then for every $\varepsilon > 0$,*

$$\mathsf{P}^r\left((1-\varepsilon)(f_i+z) \succ_i f_i \text{ for some } i \in I\right) \leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2}\,. \tag{3.1}$$

The proof of this theorem is in Section 5, as is the proof of the rest of the results we provide in this section. To unpack the statement of the theorem, imagine an economy with two consumers, $I = \{1, 2\}$. If a random perturbation $z \in \mathbb{B}(r)$ is also in $(\mathcal{U}_1^{(\varepsilon)}(f_1) - f_1\}) \cup (\mathcal{U}_2^{(\varepsilon)}(f_2) - \{f_2\})$, then it means that either $(1-\varepsilon)(z+f_1)$ provides a strict welfare improvement for agent 1 over f_1, or that $(1-\varepsilon)(z+f_2)$ does this for agent 2 over f_2 (or that both things are true). If we take $\tau = r$ and consider a 10% welfare improvement ($\varepsilon = 0.1$), then the probability of making at least one agent better off is at most $e^{-d/800}$. If d is, say, the number of stocks trading on the NASDAQ Exchange, then this bound is about 1 %.

Remark 1. We can express Theorem 1 in terms of volumes. Specifically, the bound in (3.1) is equivalent to

$$\frac{\mathsf{Vol}\left(\bigcup_{i\in I}\left(\mathcal{U}_i^{(\varepsilon)}(f_i)-\{f_i\}\right)\cap\mathbb{B}(r)\right)}{\mathsf{Vol}\left(\mathbb{B}(r)\right)}\leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2}\,. \tag{3.2}$$

By the monotonicity of preferences, $z > 0$, a perturbation that leaves an agent with greater consumption state-by-state, is always a welfare improvement for any agent. Such a perturbation has probability $1/2^d$, which also decreases exponentially in the number of states. The "curvature" of agents' preferences means, however, that many other perturbations should also result in a welfare improvement — recall the dotted area in Figure 1. So the conclusion in Theorem 1 may be surprising; it says that when the improvement is by a margin ε, the scope for welfare improvements shrinks exponentially in d regardless of the shape of agents' preferences, and even if we only ask for an improvement in one agent's welfare.

As a consequence of Theorem 1, we obtain the statement that we discussed at length in the introduction:

Corollary 1. Fix $r > 0$ and let $z \sim \mathsf{P}^r$. Under the hypotheses of Theorem 1, for every $\varepsilon > 0$, and for each $i \in I$,

$$\mathsf{P}^r\left((1-\varepsilon)(f_i+z)\succ_i f_i\right)\leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2}\,. \tag{3.3}$$

Equivalently, in terms of volumes, one has

$$\frac{\mathsf{Vol}\left(\mathcal{U}_i^{(\varepsilon)}(f_i)\cap\mathbb{B}(f_i,r)\right)}{\mathsf{Vol}(\mathbb{B}(f_i,r))}\leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2}\,,\forall i\in I\,. \tag{3.4}$$

We should emphasize that we interpret volume as a uniform probability measure on the ball of radius r, so that the statement in Corollary 1 corresponds to the informal description we provided in the introduction. The volume of the ball itself shrinks exponentially in d, so it is important to condition on the ball.

Observe that the message of Theorem 1 remains the same if we restrict attention to perturbations $z \in \mathbb{B}(f_i, r)$ for which $p\cdot z > p\cdot\omega_i$. A consumption that is affordable cannot provide an improvement in utility over f_i, so it seems natural to consider only z that costs more than i's income at the equilibrium prices. This subset of $\mathbb{B}(f_i, r)$ contains half its volume, so the message of Theorem 1 and Corollary 1 remains unchanged.

Our second result considers welfare improvements outside of equilibrium. An equilibrium allocation must be Pareto optimal, but the Pareto set is strictly larger. So we ask about

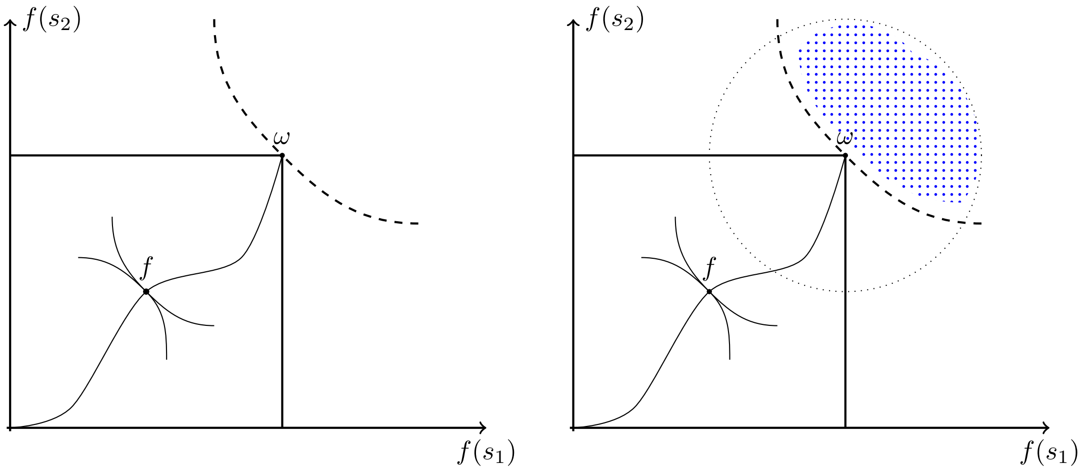

Figure 2: An illustration of Theorem 2

the possibility of a collective improvement after a change in aggregate consumption, starting from an allocation that is a weak Pareto optimum.

In our next theorem, when we assume that there is no aggregate uncertainty, we normalize the aggregate endowment to be $\omega = (1, \dots, 1) = \mathbf{1}$.

Theorem 2. *Let $\mathcal{E}$ be a convex exchange economy, with no aggregate uncertainty, and an aggregate endowment of $\omega = \mathbf{1}$. Assume f is a weakly Pareto optimal allocation. For any given $\varepsilon > 0$, let $\mathcal{V}^{(\varepsilon)} := \sum_{i \in I} \mathcal{U}_i^{(\varepsilon)}(f_i)$ be the Minkowski sum of the ε-approximate upper contour sets. Fix $r > 0$ and let $z \sim \mathsf{P}^r$. Then*

$$\mathsf{P}^r \left(\sum_{i \in I} f_i + z \in \mathcal{V}^{(\varepsilon)} \right) \leqslant \mathrm{e}^{-\varepsilon^2 d / 8 r^2} . \tag{3.5}$$

Remark 2. Similar to the Theorem 1, we can offer a version of the above result in the volume terms. Specifically, the bound in (3.5) is equivalent to

$$\frac{\mathsf{Vol}\left(\mathcal{V}^{(\varepsilon)} \cap \mathbb{B}(\omega, r)\right)}{\mathsf{Vol}(\mathbb{B}(\omega, r))} \leqslant \mathrm{e}^{-\varepsilon^2 d / 8 r^2} . \tag{3.6}$$

Figure 2 illustrates the concepts in the statement of Theorem 2. The figure depicts a case with two states, s_1 and s_2, and two individual traders. For pictorial simplicity we take $\varepsilon = 0$ in the graphs. Aggregate consumption is ω, and entails no aggregate uncertainty. The square defined by the origin and ω is an Edgeworth box that describes all the allocations between the two agents. The Pareto optimal allocations are on the solid curve that connects the origin and ω, while $f = (f_1, f_2)$ is a particular Pareto optimal allocation (indicated by the tangency of the two agents' indifference curves at the point f, with the usual Edgeworth-box convention that one agent's coordinate system has been rotated 180° and its origin coincides with ω). For the allocation f of the initial endowment ω, we may consider all the aggregate consumptions that can be decentralized among the two agents so as to obtain higher utility than in f. This set of aggregate bundles, the Scitovsky contour $\mathcal{V}^{(0)}$, consists of all the bundles to the north-east of the dashed curve passing through ω (a curve that, at ω, has the same slope as the common slope of each individual agent's utility at f).[5]

We want to study the Scitovsky contour $\mathcal{V}^{(0)}$: understand and quantify how large it is. When agents' utilities are monotone, $\mathcal{V}^{(0)}$ is, of course, an infinite set, and has infinite volume (Lebesgue measure; or area in the two-dimensional case). So we consider its volume *relative to a ball centered at* ω. The exercise can be interpreted as follows: if we were to change aggregate consumption by randomly shifting ω, what is the probability that we would end up in the Scitovsky contour. In other words, if we perturb ω by randomly (and uniformly) choosing a perturbation from the ball depicted in Figure 2, how likely is it that the resulting aggregate consumption would make both traders better off.

The exercise amounts to calculating the volume of the dotted region shown on the right of Figure 2. In the figure, the volume represents a significant fraction of the sphere centered in ω. Less than 1/2, but not by much. So it is quite likely that a random perturbation would make both agents better off. Theorem 2 says, however, that when $\varepsilon > 0$ and the number of states is large, this volume will be negligible. It converges to zero very quickly as the number of states grows.

We shall reinterpret Theorem 2 in terms of Debreu's coefficient of resource utilization (see Debreu (1951)). To this end, consider an allocation f in an exchange economy $\mathcal{E}$ that is *not* weakly Pareto optimal. This means that there is an alternative allocation of the aggregate endowment in $\mathcal{E}$ that makes all agents strictly better off. In terms of the Scitovsky contour, this means that the aggregate endowment $\omega = \sum_i \omega_i$ lies in the set $\mathcal{V}^{(0)}$.

Debreu considers the minimum amount of aggregate resources (call it ω^*) that could be used to provide agents with the same utility as in f, and thinks of the gap between ω and ω^* as the inefficiency inherent in the allocation f. In Debreu's words, these are "nonutilized resources." He proposes to measure this gap by means of a "distance with economic meaning:" $p\cdot(\omega - \omega^*)$, where p is an "intrinsic price vector" associated with ω^*. To obtain a measure

[5]The Scitovsky contour is a key concept in the proof of the second welfare theorem. For further discussion of this concept, see Debreu (1951) and Samuelson (1956).

that is, in a sense, scale independent, he works with the ratio of $p{\cdot}\omega^*$ to $p{\cdot}\omega$. Prices p follow from an argument that is analogous to the second welfare theorem; in particular, they are not uniquely defined.

Debreu's *coefficient of resource utilization* for an allocation $f = (f_1, \ldots, f_n)$ is defined as:

$$\mathrm{CRU}(f) := \max_{\omega^* \in \partial\overline{\mathcal{V}^{(0)}}} \frac{p(\omega^*){\cdot}\omega^*}{p(\omega^*){\cdot}\omega},$$

where $\partial\overline{\mathcal{V}^{(0)}}$ consists of the minimal elements of the closure $\overline{\mathcal{V}^{(0)}}$ of $\mathcal{V}^{(0)}$ (meaning there is no smaller element in $\mathcal{V}^{(0)}$), and $p(\omega^*)$ is a supporting price vector at ω^*, what Debreu calls an intrinsic price vector. Debreu (1951) shows that $\mathrm{CRU}(f)$ is well defined (in particular, that it does not depend on the selection of prices $p(\omega^*)$), that it is a number in $(0,1]$, and that $\mathrm{CRU}(f) < 1$ when f is Pareto dominated.

Then we obtain, as a simple consequence of Theorem 2:

Corollary 2. Fix $r > 0$ and let $z \sim \mathsf{P}^r$. Under the hypotheses of Theorem 2, if f is not weakly Pareto optimal, and $\mathrm{CRU}(f)$ its coefficient of resource utilization, then

$$\mathsf{P}^r\left(\sum_{i \in I} f_i + z \in \mathcal{V}^{(0)}\right) \leqslant \mathrm{e}^{-(1-\mathrm{CRU}(f))^2 d/8r^2}.$$

Corollary 2 quantifies the meaning of the coefficient of resource utilization. Debreu writes that one may think of $\mathrm{CRU}(f)$ as a percentage of national income, or GDP. But in an economy with a large state space, even a seemingly large amount of inefficiency — as measured by the coefficient of resource utilization — may not translate into a wide scope for welfare improvements by changing aggregate consumption. To use the NASDAQ example from before, a seemingly large inefficiency of 50% measured by the $\mathrm{CRU}(f)$, translates into a Corollary 2 bound of e^{-112}, which is essentially zero. In words, despite a large inefficiency of 50%, the chance that a random perturbation could be distributed to make all agents better off (not by some $\varepsilon > 0$, just strictly better off) is essentially zero.

3.2 Prior Beliefs and Welfare-Improving Trade

Our third result concerns agents with multiple priors, and the size of the sets of prior beliefs that they may posses. In Section 3.1, we quantified the upper contour sets of individual agents and their sums. Here we instead follow Yaari (1969) to interpret a vector that supports the upper contour sets, at some contingent consumption, as a prior belief. We depend on a relaxation of the equivalence between the existence of a common prior and Pareto efficiency of an allocation. This equivalence has been studied in a number of previous works (e.g., Billot et al., 2000; Ng, 2003; Rigotti et al., 2008; Gilboa et al., 2014; Ghirardato and Siniscalchi, 2018). We shall follow Rigotti et al. (2008) quite closely here.

We consider an exchange economy $\mathcal{E}$ with no aggregate uncertainty. Importantly, here on we will not require convexity of preferences. The aggregate endowment is the same across all states of the world: $\omega = (\bar{\omega}, \dots, \bar{\omega})$. We quantify the space of all allocations, denoted by $\mathcal{F}_{\bar{\omega}}$, by the magnitude

$$\rho := 2\bar{\omega}^{-1} \max_{f \in \mathcal{F}_{\bar{\omega}}} \sum_{i \in I} \|f_i\| \,. \tag{3.7}$$

Yaari (1969) defines the *subjective belief* as a probability distribution vector that supports the upper contour set. When the upper contour has a kink (due, for example, to the ambiguity in preferences and its induced lack of differentiability), there will be multiple supporting vectors. In the spirit of Rigotti et al. (2008), we define the *subjective belief set* as the set of all supporting vectors. Specifically, let f be an act in $\mathbb{R}^d_+$. The upper contour set of agent i is $\{g : g \succeq_i f\}$, and the subjective belief set at f is defined by

$$\mathcal{B}_i(f) = \{\mu \in \Delta S : \mu \cdot g \geqslant \mu \cdot f \text{ for all } g \succeq_i f\} \,.$$

The set $\mathcal{B}_i(f)$ is a convex and closed (hence compact) subset of the d-dimensional probability simplex ΔS. One may also interpret a vector in $\mathcal{B}_i(f)$ as the set of *normalized prices* that support the consumption of the act f.[6]

Following our notation from before, we define the δ-extension of the subjective belief set by

$$\mathcal{B}_i(f_i)^\delta := \{\nu \in \Delta S : \inf_{\mu \in \mathcal{B}_i(f_i)} \|\nu - \mu\| < \delta\} \,.$$

Proposition 1. *Let $\mathcal{E}$ be an exchange economy with preferences $\succeq_i \in \mathcal{P}$ for all $i \in I$ and no aggregate uncertainty. Set $\delta = \varepsilon/\rho$. If the allocation f is ε-Pareto dominated, then $\bigcap_{i \in I} \mathcal{B}_i(f_i)^\delta = \varnothing$.*

In Proposition 1, we show that if an allocation is not *approximately Pareto optimal*

[6] The price interpretation is suitable for the exercise in Section 3.1, where we could have phrased our results in terms of the Walrasian theory of general equilibrium where there is no uncertainty and S captures the space of available consumption goods. Here we are more tightly following the story of uncertain consumption in financial markets.

(measured by the parameter ε), and thus there is room for welfare-improving trade, then the δ-extension of the subjective belief sets share no common prior.

Remark 3. In Proposition 1, the extension of the subjective belief sets and the definition of ρ in (3.7) are both with respect to the ℓ_2 norm. However, one can readily generalize by choosing an arbitrary p-norm for the belief sets, its conjugate q-norm for the definition of ρ, and the proof follows analogously.

Leveraging this result, in the following theorem we examine the volume of the prior sets as the dimension d grows large. Before we state the result, we need to introduce some further notation. For a subset $J \subseteq I$, denote its complement by J^c, and define $\mathcal{B}_J(f_J) := \bigcap_{j \in J} \mathcal{B}_j(f_j)$. We often drop the allocation f from the argument of subjective belief sets, when it is understood from the context.

Theorem 3. *Let $\mathcal{E}$ be an exchange economy with preferences $\succeq_i \in \mathcal{P}$ for all $i \in I$ and no aggregate uncertainty. If the allocation f is ε-Pareto dominated, then there exists a constant $c > 0$, such that for every proper subset $J \subset I$,*

$$\frac{\min\left(\mathsf{Vol}\left(\mathcal{B}_J\right), \mathsf{Vol}\left(\mathcal{B}_{J^c}\right)\right)}{\mathsf{Vol}\left(\Delta_d\right)} \leqslant \frac{1}{2}\, \mathrm{e}^{-c\varepsilon\sqrt{d}}\,. \tag{3.8}$$

Moreover, the constant c is universal: its value is independent of the primitives of the economy, the dimension d and the parameter ε.

In particular, if in a two-agent economy, an allocation is not approximately Pareto optimal, then the volume of the subjective belief set of at least one agent is exponentially smaller than the volume of the probability simplex. Put differently, if there is a possibility for a strong welfare-improving trade, then the subjective belief set of at least one agent must be "very small."

We interpret the theorem as saying that some degree of ambiguity neutrality is needed for the existence of a welfare-improving trade. That agents have a small set of prior beliefs seems to capture the idea of being ambiguity neutral, but the volume measure does not have an obvious behavioral counterpart. We seek to provide such a behavioral interpretation in Section 4.

4 Discussion

4.1 Behavioral Implications of Theorem 3

In this section we study the behavioral implications of Theorem 3. For the purposes of this discussion, we consider a setting with two agents. The message of our theorem is that, if the current risk-sharing arrangement between the two agents is strongly Pareto dominated (namely if there exists a possibility of a strongly welfare improving ex-ante trade), then the volume of the underlying belief set of at least one agent must be very small. The *smaller* is this volume, the *closer* is that agent to the ambiguity neutrality — a connection we seek to quantify in this section.

It is well-established that ambiguity aversion leads to less trade. Theorem 3 offers a quantitative expression of this result: that the possibility of an ε-Pareto improving trade necessitates small ambiguity aversion in high dimensions.

This interpretation is in line with the comparative notion of ambiguity aversion proposed by Ghirardato et al. (2004). Specifically, they show in the max-min setting of Gilboa and Schmeidler (1989), the ambiguity aversion of a decision maker *decreases* as her multiple prior set *shrinks* with respect to the set inclusion order.

To establish the connection between ambiguity aversion and the volume of the belief set, we appeal to the max-min setting of Gilboa and Schmeidler (1989), in which there is a convex compact set of priors $\Pi \subseteq \Delta S$ and the agent's cardinal evaluation of an act $f \in \mathbb{R}^d_+$ is $u(f) = \min_{\mu \in \Pi} f \cdot \mu$. We further assume that Π has constant width, namely the distance between any distinct parallel supporting hyperplanes of Π (residing on the d dimensional probability simplex) is constant.

One may define the level of ambiguity aversion by the difference between the maximum and minimum expected utility of a *normalized* act f (i.e., $\|f\|_2 = 1$) over Π, namely

$$\theta(f) := \max_{\mu \in \Pi} f \cdot \mu - \min_{\mu \in \Pi} f \cdot \mu \,.$$

When Π has constant width, $\theta(f)$ becomes constant in f, and we write $\theta(f) \equiv \theta$. Thus we can take θ as a measure of ambiguity aversion. In the next proposition we show how the upper bound on the relative volume in Theorem 3 means that θ vanishes as $d \to \infty$.

Proposition 2. *Under the conditions of Theorem 3, let Π coincide with the set of priors with the smaller volume, and suppose that it has constant width θ. Then there exists a universal* *constant $c > 0$ such that*

$$\theta \leqslant 4\,\mathrm{e}^{-c\varepsilon/\sqrt{d}}(d!)^{-1/2d}\,. \tag{4.1}$$

4.2 On the Magnitude of ε

The interpretation of our results depends on the magnitude of two numbers: the number of states d, and the fraction ε accounted for in a utility improvement. When d is large while ε remains constant, the volumes discussed in Theorems 1 and 2 shrink to zero. So one could question our interpretation by arguing that we should use a small value of ε when the number of states d is large. In particular, one might argue that we should impose that $\varepsilon = O(1/\sqrt{d})$. We would disagree, however, essentially because we think of ε as a dimension-free fraction.

First, ε is expressed as a fraction of physical units of stage-contingent consumption. If we think that $f(s)$, for an act f, is a monetary payment, then ε is a percentage of a monetary payment. It seems odd to impose a smaller percentage in monetary terms when the number of states is large than when it is small. For example, if we identify states with the number of assets in a market: is the meaning of a 5% return different in a market with many assets than in a market with few assets? If we instead consider $\|\omega\|$ to be a measure of the "size" of the economy, then we may want to impose a value of ε that represents a constant fraction of $\|\omega\|$. For example, with the assumption that $\omega = \mathbf{1}$ we have $\|\omega\| = \sqrt{d}$. Of course, the resulting ε would then grow with d, and only strengthen our results.

Second, taking $\varepsilon = O(1/\sqrt{d})$ is problematic because it seems very hard to reconcile with the common practice of using a numerical objective function in calculating approximately optimal outcomes. As we discussed above, many applications make use of a homothetic preference and a resulting utility function that is homogenous. Examples include the max-min representation in choice under uncertainty, or the Cobb-Douglas utility in consumer choice. In this case, the ε equals that tolerance level assumed in the agents' maximization problem: the number ε is then measured in "utils," the same unit of account as used for the utility function. Utils are, however, dimension free. In a model with many states of the world, we would hardly be allowing for any relaxation in our notion of approximate optimality.

Finally, we should emphasize the reinterpretation of our results using the coefficient of resource utilization (see Corollary 2). The CRU is usually thought of as a fraction of national income, also a dimension free measure, and one that one might expect is constant, or even grow, with the size of the economy.

5 Methodology and Proofs

At a high level, the primary concept underpinning our proofs is that in high-dimensional metric spaces, and under relatively mild conditions, probability measures tend to concentrate. As a result, for a probability measure μ and a subset A that encompasses at least half of the probability space, namely $\mu(A) \geqslant 1/2$, the metric extension $A^\delta = \{z : \mathsf{dist}(z, A) < \delta\}$ covers a substantial portion and growing of the unit measure. In particular, the complement $1-\mu(A^\delta)$ diminishes rapidly, often exhibiting exponential decay with respect to the dimension d. These types of concentration bounds are commonly referred to by *Isoperimetric inequalities*. Their importance lie on the *independence* of the concentration rate from the set A. The source of our results, that certain welfare improvements have a vanishingly small probability, independently of the shape of the agents' preferences, can be traced to basic results in the theory of Isoperimetric inequalities and concentration of measure.

A consequence of the uniform concentration of measure is that, if two subsets A and B are separated with a positive distance, then as the dimension d grows, the measure of at least one of them must be exponentially small. We apply this basic idea in the proofs of our results.

In the following, we begin by introducing the core inequality in concentration of measure: Brunn-Minkowski. Then, we apply a variant of this inequality to prove Theorems 1 and 2. Followed by that, we present the preliminaries of Isoperimetric inequalities, and employ them to prove Theorem 3.

5.1 Brunn-Minkowski Inequality

For two subsets $A, B \subseteq \mathbb{R}^d$, their Minkowski sum is defined by $A+B = \{a+b : a \in A, b \in B\}$. The Brunn-Minkowski Inequality provides a crucial connection between volumes and Minkowski sum in Euclidean spaces.

Let A and B be two non-empty compact subsets of $\mathbb{R}^d$. The Brunn-Minkowski inequality claims that

$$\mathsf{Vol}(A+B)^{1/d} \geqslant \mathsf{Vol}(A)^{1/d} + \mathsf{Vol}(B)^{1/d}. \tag{5.1}$$

If one makes the additional assumption that A and B are restricted to convex subsets, then the inequality binds if and only if A and B are homothetic (that is one is the translated and scaled version of another). This inequality implies the concavity of the volume operator with respect to the Minkowski sum. There is a *dimension-free* version of this inequality that often proves more useful. In particular, for $\lambda \in [0,1]$, inequality (5.1) implies that

$$\mathsf{Vol}(\lambda A + (1-\lambda)B)^{1/d} \geqslant \lambda \mathsf{Vol}(A)^{1/d} + (1-\lambda)\mathsf{Vol}(B)^{1/d}\,.$$

Applying the arithmetic geometric inequality to the above provides the following dimension-free version:

$$\mathsf{Vol}(\lambda A + (1-\lambda)B) \geqslant \mathsf{Vol}(A)^{\lambda}\,\mathsf{Vol}(B)^{1-\lambda}\,. \tag{5.2}$$

We henceforth refer to this inequality by BM inequality. A useful application of BM is the following lemma: Expressing an upper bound for the minimum volume of two positively distanced subsets (its proof can be found in Artstein-Avidan et al. (2015), but we also state it for completeness).

Lemma 1. *Assume A and B are Borel subsets of $\mathbb{B}(r)$, and $\mathsf{dist}(A,B) \geqslant \delta$. Then,*

$$\frac{\min\{\mathsf{Vol}(A), \mathsf{Vol}(B)\}}{\mathsf{Vol}(\mathbb{B}(r))} \leqslant \mathrm{e}^{-\delta^2 d/8r^2}\,. \tag{5.3}$$

Proof. Since the volume of any Borel set can be approximated arbitrarily close by the inner measure of its closed subsets, we can assume without any loss that A and B are closed and hence compact. By the parallelogram law for the ℓ_2-norm if $a \in A$ and $b \in B$ then

$$\|a+b\|^2 = 2\|a\| + 2\|b\|^2 - \|a-b\|^2 \leqslant 4r^2 - \delta^2\,,$$

where the inequality holds because $a, b \in \mathbb{B}$ and $\|a-b\| \geqslant \delta$. Hence, it follows that

$$\frac{A+B}{2} \subseteq \sqrt{1 - \frac{\delta^2}{4r^2}}\,\mathbb{B}(r)\,,$$

and therefore,

$$\mathsf{Vol}\left(\frac{A+B}{2}\right) \leqslant \left(1 - \frac{\delta^2}{4r^2}\right)^{d/2} \mathsf{Vol}(\mathbb{B}(r)) \leqslant \mathrm{e}^{-\delta^2 d/8r^2}\,\mathsf{Vol}(\mathbb{B}(r))\,.$$

Setting $\lambda = 1/2$ in (5.2) and using the above inequality justify the claim in (5.3). $\square$

Therefore, as the dimension grows, two subsets in the ℓ_2 ball with a bounded radius, will have positive distance from each other only if at least one them has a very small volume. Of course, the larger is the distance, the smaller would be the implied volume.

5.2 Proof of Results in Section 3.1

We proceed with the proof of our first two theorems. In both cases, using the optimality or the equilibrium property of the allocation, we apply a type of convex separation argument. In the first theorem, the separation is provided by the given equilibrium price. While in the second theorem, the separation argument follows the idea used in the proof of the second welfare theorem. Consequently, as we argue below, the approximate versions of upper contour sets stay at a positive distance from their separating counterpart. An application of Lemma 1 on appropriately chosen subsets imply the volume bounds in Theorems 1 and 2.

We begin by laying down some terminology. For a vector $p \in \mathbb{R}^d$ and a constant b, we define two half-spaces:

$$H^+(p\,;\,b) := \left\{x \in \mathbb{R}^d : p \cdot x \geqslant b\right\}\,,$$
$$H^-(p\,;\,b) := \left\{x \in \mathbb{R}^d : p \cdot x \leqslant b\right\}\,,$$

that are, respectively, called upper and lower half-spaces. One can readily verify that the ℓ_2 distance between the two half-spaces $H^+(p\,;\,b_2)$ and $H^-(p\,;\,b_1)$, where $b_2 > b_1$, is equal to

$$\mathsf{dist}\left(H^+(p\,;\,b_2), H^-(p\,;\,b_1)\right) = \frac{b_2 - b_1}{\|p\|}\,. \tag{5.4}$$

Proof of Theorem 1. Since $f = \{f_i : i \in I\}$ is a Walrasian equilibrium allocation in $\mathcal{E}$, and preferences are monotone, there exists a price vector $p \in \mathbb{R}^d_+$ such that $p \cdot g_i > p \cdot \omega_i$ for all $i \in I$ and $g_i \in \mathcal{U}_i^{(0)}(f_i)$. Next, observe that if $g \in \mathcal{U}_i^{(\varepsilon)}(f_i)$ then $(1-\varepsilon)g \in \mathcal{U}_i^{(0)}(f_i)$ and therefore $p\cdot((1-\varepsilon)(g-\omega_i)-\varepsilon\omega_i) > 0$. This in turn implies that

$$p \cdot (g - \omega_i) > \frac{\varepsilon p \cdot \omega_i}{1-\varepsilon} > \varepsilon p \cdot \omega_i \geqslant \varepsilon\tau\|p\|_1\,,$$

where the last inequality holds because $p \in \mathbb{R}^d_+$ and $\omega_i \geqslant \tau\mathbf{1}$. Therefore, since $i \in I$ and $g \in \mathcal{U}_i^{(\varepsilon)}(f_i)$ were arbitrary, by the above inequality $\mathcal{U}_i^{(\varepsilon)}(f_i) - \{\omega_i\} \subseteq H^+(p; \varepsilon\tau\|p\|_1)$ for all $i \in I$. Let us define $\mathcal{Q} := \bigcup_{i\in I}\left(\mathcal{U}_i^{(\varepsilon)}(f_i) - \{\omega_i\}\right)$.[7] Then $\mathcal{Q} \subseteq H^+(p; \varepsilon\tau\|p\|_1)$, so for an arbitrary $r > 0$, one has

$$\begin{aligned}\mathsf{dist}\left(\mathcal{Q}\cap\mathbb{B}(r), H^-(p;0)\cap\mathbb{B}(r)\right) &\geqslant \mathsf{dist}\left(H^+(p;\varepsilon\tau\|p\|_1)\cap\mathbb{B}(r), H^-(p;0)\cap\mathbb{B}(r)\right)\\ &\geqslant \mathsf{dist}\left(H^+(p;\varepsilon\tau\|p\|_1), H^-(p;0)\right) = \varepsilon\tau\frac{\|p\|_1}{\|p\|} \geqslant \varepsilon\tau\,.\end{aligned}$$

[7]The set $\mathcal{Q}$ was originally used by Debreu and Scarf (1963) to prove core convergence, and by Barman and Echenique (2023) to characterize approximate Walrasian equilibria.

The equality above follows from (5.4), and the last inequality holds because $\min_{p\neq \mathbf{0}} \|p\|_1/\|p\|_2 = 1$, namely the minimum of $\|p\|_1/\|p\|_2$ is achieved on the standard unit basis vectors, and is equal to 1. Now set $A := \mathcal{Q} \cap \mathbb{B}(r)$ and $B := H^-(p;0) \cap \mathbb{B}(r)$. By the above inequality $\mathsf{dist}(A, B) \geqslant \varepsilon\tau$. Since p is a nonzero vector in $\mathbb{R}^d_+$, the subset B covers at least half of the volume of $\mathbb{B}(r)$. So it must be that $\mathsf{Vol}(A) \leqslant \mathsf{Vol}(B)$. Therefore, Lemma 1 implies that $\mathsf{Vol}(A)/\mathsf{Vol}(\mathbb{B}(r)) \leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2}$, namely:

$$\frac{\mathsf{Vol}\left(\mathcal{Q} \cap \mathbb{B}(r)\right)}{\mathsf{Vol}(\mathbb{B}(r))} \leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2} ,$$

and thereby

$$\frac{\mathsf{Vol}\left(\bigcup_{i\in I} \left(\mathcal{U}_i^{(\varepsilon)}(f_i) - \{\omega_i\}\right) \cap \mathbb{B}(r)\right)}{\mathsf{Vol}\left(\mathbb{B}(r)\right)} \leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2} .$$

Now observe that if $f = \{f_i : i \in I\}$ is a Walrasian equilibrium for the exchange economy $\mathcal{E}$, it is also a Walrasian equilibrium for the exchange economy $\mathcal{E}'$ that is identical to $\mathcal{E}$ except that each agent i's endowment is $\omega'_i = f_i$. Therefore, we can replace ω_i in the above inequality with f_i, and obtain the volume bound in (3.2) that is equivalent to (3.1). □

Proof of Corollary 1. Because of the union bound and (3.1) one has

$$\frac{\mathsf{Vol}\left(\left(\mathcal{U}_i^{(\varepsilon)}(f_i) - \{f_i\}\right) \cap \mathbb{B}(r)\right)}{\mathsf{Vol}\left(\mathbb{B}(r)\right)} \leqslant \mathrm{e}^{-\varepsilon^2\tau^2 d/8r^2} , \forall i \in I .$$

Since the volume is translation invariant, we can shift the subsets in the above inequality by f_i and thus achieve the bound in (3.4) and thereby (3.3). □

Proof of Theorem 2. Since $f = \{f_i : i \in I\}$ is weakly Pareto optimal, then there is no allocation $g \in \mathcal{F}_{\mathbf{1}}$ such that $g_i \succ_i f_i$ for every $i \in I$. That is $\mathcal{F}_{\mathbf{1}} \cap \prod_{i\in I} \mathcal{U}_i^{(0)}(f_i) = \varnothing$. That in turn means the normalized endowment vector $\omega = \mathbf{1}$ is disjoint from $\mathcal{V}^{(\varepsilon)} := \sum_{i\in I} \mathcal{U}_i^{(\varepsilon)}(f_i)$ for all $\varepsilon \geqslant 0$. Since preferences are convex the approximate upper contour sets are convex, so is their sum $\mathcal{V}^{(\varepsilon)}$. Therefore, by the hyperplane separation theorem and monotonicity of preferences there exists a nonzero vector $p \in \mathbb{R}^d_+$ such that $p \cdot v \geqslant p \cdot \mathbf{1} = \|p\|_1$ for all $v \in \mathcal{V}^{(0)}$. That is $\mathbf{1} \in H^-(p; \|p\|_1)$ and $\mathcal{V}^{(0)} \subseteq H^+(p; \|p\|_1)$.

Now suppose $v \in \mathcal{V}^{(\varepsilon)}$. Then, there are $g_i \in \mathcal{U}_i^{(\varepsilon)}$ for all $i \in I$, such that $v = \sum_{i\in I} g_i$ and $(1-\varepsilon)g_i \succ_i f_i$. For each $i \in I$, it holds that $(1-\varepsilon)g_i \in \mathcal{U}_i^{(0)}$ and hence $(1-\varepsilon)\sum_{i\in I} g_i \in \mathcal{V}^{(0)}$. Consequently, one has $p \cdot (1-\varepsilon)v \geqslant \|p\|_1$. That, in turn, implies $v \in H^+\left(p; \|p\|_1/(1-\varepsilon)\right)$, and thereby $\mathcal{V}^{(\varepsilon)} \subseteq H^+(p; \|p\|_1/(1-\varepsilon))$. As a result of this set inclusion, for an arbitrary $r > 0$, one obtains that

$$\begin{aligned}
\mathsf{dist}\left(\mathcal{V}^{(\varepsilon)} \cap \mathbb{B}(\mathbf{1}, r)\,,\, H^-(p\,;\, \|p\|_1) \cap \mathbb{B}(\mathbf{1}, r)\right) &\geqslant \\
\mathsf{dist}\left(H^+\left(p\,;\, \frac{\|p\|_1}{1-\varepsilon}\right) \cap \mathbb{B}(\mathbf{1}, r)\,,\, H^-\left(p\,;\, \|p\|_1\right) \cap \mathbb{B}(\mathbf{1}, r)\right)& \\
= \mathsf{dist}\left(H^+\left(p\,;\, \frac{\|p\|_1}{1-\varepsilon}\right)\,,\, H^-(p\,;\, \|p\|_1)\right)&\,,
\end{aligned}$$

where the equality holds because the two half-spaces are parallel. Their distance by (5.4) is equal to $\varepsilon\|p\|_1/(1-\varepsilon)\|p\|$. Therefore, we arrive at

$$\mathsf{dist}\left(\mathcal{V}^{(\varepsilon)} \cap \mathbb{B}(\mathbf{1}, r)\,,\, H^-(p\,;\, \|p\|_1) \cap \mathbb{B}(\mathbf{1}, r)\right) \geqslant \frac{\varepsilon\|p\|_1}{(1-\varepsilon)\|p\|} \geqslant \varepsilon\,,$$

where the last inequality follows as before, because $\min_{p\neq\mathbf{0}}\|p\|_1/\|p\|_2 = 1$. Now set $A := \mathcal{V}^{(\varepsilon)} \cap \mathbb{B}(\mathbf{1}, r)$ and $B := H^-(p\,;\, \|p\|_1) \cap \mathbb{B}(\mathbf{1}, r)$. By the above inequality one has $\mathsf{dist}(A, B) \geqslant \varepsilon$. Since p is a nonzero vector in $\mathbb{R}^d_+$, the subset B covers at least half of the volume of $\mathbb{B}(\mathbf{1}, r)$. So it must be that $\mathsf{Vol}(A) \leqslant \mathsf{Vol}(B)$. Therefore, Lemma 1 implies that $\mathsf{Vol}(A)/\mathsf{Vol}(\mathbb{B}(\mathbf{1}, r)) \leqslant \mathrm{e}^{-\varepsilon^2 d/8r^2}$, thus proving the volume bound (3.6) that is equivalent to (3.5). $\square$

Proof of Corollary 2. Let $\beta := \mathrm{CRU}(f)$. Since f is not weakly Pareto optimal, then $\beta < 1$. By Debreu (1951), one has $\beta\omega \in \partial\overline{\mathcal{V}^{(0)}}$. Thus, ω is a minimal element of the closure of $\beta^{-1}\mathcal{V}^{(0)}$. Observe that for every i, one has

$$\frac{1}{\beta}\mathcal{U}_i^{(0)}(f_i) = \mathcal{U}_i^{(1-\beta)}(f_i)\,.$$

Therefore $\beta^{-1}\,\mathcal{V}^{(0)} = \mathcal{V}^{(1-\beta)}$, and ω becomes a minimal element of $\mathcal{V}^{(1-\beta)}$ as well. This fact, in turn, means that f is $(1-\beta)$-Pareto optimal. Thus the corollary now follows from Theorem 2. $\square$

5.3 Concentration and Isoperimetric Inequalities

Isoperimetric inequalities provide lower bounds for the surface measure of Borel subsets. Specifically, suppose μ is a given probability measure on $\mathbb{R}^d$, and let $A \subset \mathbb{R}^d$ be a Borel subset, whose δ-extension is denoted by $A^\delta = A + \delta\mathbb{B}$. Then, the *Minkowski content* (denoted by μ^+) of the subset A relative to the measure μ is defined by

$$\mu^+(A) := \liminf_{\delta\to 0} \frac{\mu(A^\delta)-\mu(A)}{\delta}\,. \tag{5.5}$$

Given this definition, we can think of $\mu^+(A)$ as the area measure of the boundary of A.

Among subsets with measures in a certain range, the *Isoperimetric function* $\mathcal{I}_\mu : [0,1/2) \to \mathbb{R}_+$ returns the Minkowski content of the subset with the smallest boundary area. Formally, it is defined by

$$\mathcal{I}_\mu(a) := \inf_{1/2<\mu(A)\leqslant 1-a} \mu^+(A)\,. \tag{5.6}$$

In many environments, where the probability measure satisfy some mild regularity conditions, there exist universal lower bounds for the Isoperimetric function. One particular case that is of interest to us is the following lemma.

Lemma 2 (Barthe and Wolff (2009))**.** *Let u be the uniform measure on the probability simplex Δ_d. That is $u(A) = \mathsf{Area}(A)/\mathsf{Area}(\Delta_d)$, for every $A \subseteq \Delta_d$. Then, there exists a universal constant $c > 0$ such that for $a \in [0,1/2)$:*

$$\mathcal{I}_u(a) \geqslant c\,a\,d\,. \tag{5.7}$$

In the following, we use $u(\cdot)$ to refer to the uniform measure on Δ_d. As a corollary of the previous lemma we show that if a subset covers at least half of the measure on Δ_d, then the measure of its δ-extension is very close to 1.

Corollary 3. Assume $u(A) \geqslant 1/2$, then

$$u(A^\delta) \geqslant 1-\frac{1}{2}\,\mathrm{e}^{-c\delta d}\,. \tag{5.8}$$

Proof. Because of the definition of the Minkowski content in (5.5) and the Isoperimetric function in (5.6) — both based on the limit inferior — one obtains

$$\begin{aligned}
u(A^\delta) &\geqslant u(A)+\int_0^\delta u(A^t)\,\mathrm{d}t \\
&\geqslant u(A)+\int_0^\delta \mathcal{I}_u(1-u(A^t))\,\mathrm{d}t \\
&\geqslant u(A)+cd\int_0^\delta \big(1-u(A^t)\big)\,\mathrm{d}t\,,
\end{aligned}$$

where the third inequality follows from (5.7). Define $z(0) := u(A)$, and let $z : [0, \delta] \to \mathbb{R}$ be the solution to the following integral equation:

$$z(\delta) = z(0) + cd \int_0^\delta (1 - z(t))\, \mathrm{d}t\,.$$

Grönwall's inequality implies that $u(A^\delta) \geqslant z(\delta)$. One can simply verify that $z(\delta) = 1 - ((1 - z(0))\, \mathrm{e}^{-c\delta d}$, and this establishes the claim in (5.8). □

An important consequence of this result, that lies at the core of the proof of Theorem 3, is that if two subsets in Δ_d have positive distance from each other, then the area of at least one of them must be exponentially smaller than $\mathsf{Area}(\Delta_d)$. Intuitively, this resembles the separation argument in Lemma 1, although its proof is not a direct consequence of the BN inequality, and follows from the more elaborate construct of the aforementioned Isoperimetric lower bound.

Lemma 3. *Assume A and B are two Borel subsets of Δ_d, where $A^\delta \cap B^\delta = \varnothing$. Then, there exists a universal constant $c > 0$ such that:*

$$\min\{u(A), u(B)\} \leqslant \frac{1}{2}\, \mathrm{e}^{-c\delta d}\,. \tag{5.9}$$

Proof. Without any loss we shall assume that $u(A^\delta) \leqslant u(B^\delta)$. Since $A^\delta \cap B^\delta = \varnothing$, then $u(A^\delta) + u(B^\delta) \leqslant 1$. Hence the measure of the complement of A^δ is greater than or equal to $1/2$, i.e., $u(\Delta_d \backslash A^\delta) \geqslant 1/2$. We claim that $A \subseteq \Delta_d \setminus [\Delta_d \setminus A^\delta]^\delta$, which is equivalent to $\Delta d \setminus A \supseteq [\Delta_d \setminus A^\delta]^\delta$. To show the latter, pick any point $x \in [\Delta_d \setminus A^\delta]^\delta$. Recall that we defined the δ-extension with strict inequality, hence, $\mathsf{dist}(x, \Delta_d \setminus A^\delta) < \delta$. Since A^δ is an open subset, then $\Delta_d \setminus A^\delta$ is compact, and thus there exists $y \in \Delta_d \setminus A^\delta$ such that

$$\|x - y\| = \mathsf{dist}(x, \Delta_d \setminus A^\delta) < \delta\,.$$

On the other hand, $y \in \Delta_d \setminus A^\delta$ implies that $\mathsf{dist}(y, A) \geqslant \delta$. The previous two inequalities imply that $x \notin A$, and hence our claim is verified. Therefore, we have $u(A) \leqslant 1 - u\big([\Delta_d \setminus A^\delta]^\delta\big)$. Since $u\big(\Delta_d \setminus A^\delta\big) \geqslant 1/2$, Corollary 3 implies that

$$u\big([\Delta_d \setminus A^\delta]^\delta\big) \geqslant 1 - \frac{1}{2}\, \mathrm{e}^{-c\delta d}\,,$$

thereby verifying the inequality in (5.9). □

5.4 Proof of Results in Section 3.2

Proposition 1 implies that if an allocation is ε-Pareto dominated, then the extension of the subjective belief sets have empty intersection. Hence, any arbitrary split of the agents' index set I into two groups results into two subsets whose extensions have empty intersection as well. Thus, we can employ Lemma 3 to conclude that at least one of these subsets should have a small volume.

Proof of Proposition 1. Assume towards a contradiction that $\bigcap_{i\in I}\mathcal{B}_i(f_i)^\delta \neq \varnothing$, while f is ε-Pareto dominated. Choose $\eta \in \bigcap_{i\in I}\mathcal{B}_i(f_i)^\delta$. Since f is ε-Pareto dominated, there is $g \in \mathcal{F}_{\bar{\omega}}$ such that $(1-\varepsilon)g_i \succ_i f_i$ for all i. By definition of the subjective belief set, and continuity one obtains that $\mu_i\cdot\big[(1-\varepsilon)g_i - f_i\big] > 0$ for all $\mu_i \in \mathcal{B}_i(f_i)$. Choose $\tilde{\mu}_i \in \arg\min\{\|\eta-\mu_i\| : \mu_i \in \mathcal{B}_i(f_i)\}$. Observe that $\|\eta-\tilde{\mu}_i\| < \delta$, hence

$$\big|\eta\cdot\big[(1-\varepsilon)g_i-f_i\big]-\tilde{\mu}_i\cdot\big[(1-\varepsilon)g_i-f_i\big]\big| \leqslant \|\eta-\tilde{\mu}_i\|\|g_i(1-\varepsilon)-f_i\| < \delta\|(1-\varepsilon)g_i-f_i\|\,.$$

Therefore, it holds that

$$\eta\cdot\big[(1-\varepsilon)g_i-f_i\big] > \tilde{\mu}_i\cdot\big[(1-\varepsilon)g_i-f_i\big]-\delta\|(1-\varepsilon)g_i-f_i\| > -\delta\|(1-\varepsilon)g_i-f_i\|\,,$$

where the last inequality follows because $\tilde{\mu}_i\cdot\big[g_i(1-\varepsilon)-f_i\big] > 0$. Summing over all i's leads to

$$\eta\cdot\sum_{i=1}^{n}\big[(1-\varepsilon)g_i-f_i\big] > -\delta\sum_{i=1}^{n}\|(1-\varepsilon)g_i-f_i\| \geqslant -\delta\bar{\omega}\rho = -\varepsilon\bar{\omega}\,.$$

Since both f and g belong to $\mathcal{F}_{\bar{\omega}}$, the left most side above is equal to $-\varepsilon\bar{\omega}$, thus leading to a contradiction. □

Proof of Theorem 3. By Proposition 1, if f is ε-Pareto dominated, then $\bigcap_{i\in I}\mathcal{B}_i(f_i)^\delta = \varnothing$. For an arbitrary $J \subset I$, one can readily verify that $\mathcal{B}_J^\delta \subseteq \bigcap_{j\in J}\mathcal{B}_j(f_j)^\delta$. Therefore, $\mathcal{B}_J^\delta \cap \mathcal{B}_{J^c}^\delta = \varnothing$. Applying Lemma 3 implies the conclusion in (3.8). □

A Remaining Proofs

Continuity of the Approximate Pareto Correspondence

Fix an exchange economy $\mathcal{E}$ and denote the set of all ε-Pareto optimal allocations by $\boldsymbol{P}^{(\varepsilon)}$. The subset $\boldsymbol{P}^{(\varepsilon)}$ is evidently increasing with respect to the set inclusion order: $\boldsymbol{P}^{(\varepsilon')} \subseteq \boldsymbol{P}^{(\varepsilon'')}$ for $\varepsilon' \leqslant \varepsilon''$. Using continuity and monotonicity of preferences, we show that the correspondence $\boldsymbol{P} : [0,1] \twoheadrightarrow \mathbb{R}_+^{d\times I}$ is upper-hemicontinuous on $[0,1]$ and continuous at $\varepsilon = 0$. This means as $\varepsilon \to 0$ the subset $\boldsymbol{P}^{(\varepsilon)}$ "*approximates*" the space of weakly Pareto optimal allocations.

Proposition A.1. *The correspondence $\boldsymbol{P} : [0,1] \twoheadrightarrow \mathbb{R}_+^{d\times I}$ is upper-hemicontinuous.*

Proof. Let $\varepsilon_k \to \varepsilon$, $f^{(k)} \in \boldsymbol{P}^{(\varepsilon_k)}$. Since $\mathcal{F}_w$ is a compact subset of $\mathbb{R}_+^{d\times I}$, then $f^{(k)}$ has a limit point, i.e., $\exists f \in \mathcal{F}_w$ such that the subsequence $f^{(k_m)} \to f$ as $m \to \infty$. To avoid clutter, we denote the indices of this subsequence by k instead of k_m. To justify the upper-hemicontinuity, we need to show that $f \in \boldsymbol{P}^{(\varepsilon)}$. Since $f^{(k)} \to f$, then $\sum_{i\in I}^{n} \|f_i^{(k)} - f_i\| \to 0$, that in turn implies $f_i^{(k)} \to f_i$ uniformly over $i \in I$. Now assume by contradiction that $f \notin \boldsymbol{P}^{(\varepsilon)}$. Then, there exists $g \in \mathcal{F}_w$ such that $(1-\varepsilon)g_i \succ_i f_i$ for all $i \in I$. Since preferences are continuous (A2), there exists a $\bar{k}$ such that $(1-\varepsilon)g_i \succ_i f_i^{(k)}$ for all $k \geqslant \bar{k}$ and for all $i \in I$. Because of continuity again, for each i there exists $\hat{\varepsilon}_i > 0$, such that $(1-\varepsilon-\hat{\varepsilon}_i)g_i \succ_i f_i^{(k)}$ for all $k \geqslant \bar{k}$. Let $\hat{\varepsilon} := \min_{i\in I} \hat{\varepsilon}_i$, then because of monotonicity (A3) one has $(1-\varepsilon-\hat{\varepsilon})g_i \succ_i f_i^{(k)}$ for all $k \geqslant \bar{k}$. Pick $k_0 = \min\{k \geqslant \bar{k} : \varepsilon_k \leqslant \varepsilon + \hat{\varepsilon}\}$. Since $\varepsilon_k \to \varepsilon$, then k_0 is finite. Observe that the previous statement implies that $f^{(k_0)}$ is not ε_{k_0}-Pareto optimal. Thus, the contradiction is reached and hence f must belong to $\boldsymbol{P}^{(\varepsilon)}$. ☐

Corollary A.1 (Continuity at $\varepsilon = 0$)**.** Suppose $\varepsilon_k \to 0$ and $f^{(k)} \in \boldsymbol{P}^{(\varepsilon_k)}$. Then, there exists a subsequence $\{f^{(k_m)}\}$ that converges to a weakly Pareto optimal allocation f. This is resulted from the upper hemicontinuity of $\boldsymbol{P}^{(\cdot)}$ at $\varepsilon = 0$. Conversely, the correspondence is clearly lower-hemicontinuous at $\varepsilon = 0$, because every weakly Pareto optimal allocation is in fact ε-Pareto optimal.

Proof of Proposition 2

Since Π is a $d-1$ dimensional surface with *constant* width, then Schramm (1988) implies that

$$\mathsf{Vol}(\Pi) \geqslant \left(\sqrt{3+\frac{2}{d}}-1\right)^{d-1} \mathsf{Vol}\left(\mathbb{B}^{d-1}(\mathbf{0}, \theta/2)\right) \geqslant \left(\sqrt{3}-1\right)^{d-1} \left(\frac{\theta}{2}\right)^{d-1} \mathsf{Vol}\left(\mathbb{B}^{d-1}\right) .$$

By Theorem 3 one has $\mathsf{Vol}(\Pi) \leqslant \frac{1}{2}\,\mathrm{e}^{-c\varepsilon/\sqrt{d}}\,\mathsf{Vol}(\Delta_d)$, therefore the above inequality leads to

$$\left(\frac{\theta}{2}\right)^{d-1} \leqslant \frac{\mathrm{e}^{-c\varepsilon\sqrt{d}}}{2\left(\sqrt{3}-1\right)^{d-1}}\,\frac{\mathsf{Vol}(\Delta_d)}{\mathsf{Vol}\left(\mathbb{B}^{d-1}\right)}\,.$$

The volume of the $d-1$ dimensional unit ℓ_2 ball and the $d-1$ dimensional probability simplex Δ_d are respectively equal to $\mathsf{Vol}(\mathbb{B}_2^{d-1}) = \pi^{(d-1)/2}/\Gamma\left(\frac{d+1}{2}\right)$ and $\mathsf{Vol}(\Delta_d) = \sqrt{d}/\Gamma(d)$. Therefore,

$$\left(\frac{\theta}{2}\right)^{d-1} \leqslant \frac{\sqrt{d}\,\mathrm{e}^{-c\varepsilon\sqrt{d}}}{2\left(\sqrt{\pi}(\sqrt{3}-1)\right)^{d-1}}\,\frac{\Gamma\left(\frac{d+1}{2}\right)}{\Gamma(d)}\,.$$

Since the Gamma function is log-convex, then $\Gamma\left(\frac{d+1}{2}\right) \leqslant \sqrt{\Gamma(d)} = \sqrt{(d-1)!}$. Hence the above inequality simplifies to

$$\left(\frac{\theta}{2}\right)^{d-1} \leqslant \frac{d\,\mathrm{e}^{-c\varepsilon\sqrt{d}}}{2\left(\sqrt{\pi}(\sqrt{3}-1)\right)^{d-1}}\,\frac{1}{\sqrt{d!}}\,.$$

Let us denote $\sqrt{\pi}\left(\sqrt{3}-1\right)$ by α. Since the width θ is smaller than 2, then $(\theta/2)^d \leqslant (\theta/2)^{d-1}$ and thus

$$\theta \leqslant \frac{2}{\alpha}\left(\frac{\alpha d}{2}\right)^{1/d}\mathrm{e}^{-c\varepsilon/\sqrt{d}}(d!)^{-1/2d}\,.$$

We can readily verify that $\frac{2}{\alpha}\left(\frac{\alpha d}{2}\right)^{1/d} \leqslant 4$ for all integers d, thereby proving the inequality (4.1).

□

References

[1] Shiri Artstein-Avidan, Apostolos Giannopoulos, and Vitali D. Milman (2015). *Asymptotic Geometric Analysis, Part I*, volume 202, American Mathematical Society.

[2] Siddarth Barman and Federico Echenique (2023). "The Edgeworth Conjecture with Small Coalitions and Approximate Equilibria in Large Economies," *Mathematics of Operations Research*, 48(1): 313–331.

[3] Franck Barthe and Paweł Wolff (2009). "Remarks on Non-Interacting Conservative Spin Systems: the Case of Gamma Distributions," *Stochastic Processes and their Applications*, 119(8): 2711–2723.

[4] Antoine Billot, Alain Chateauneuf, Itzhak Gilboa, and Jean-Marc Tallon (2000). "Sharing Beliefs: Between Agreeing and Disagreeing," *Econometrica*, 68(3): 685–694.

[5] Gerard Debreu (1951). "The Coefficient of Resource Utilization," *Econometrica*: 273–292.

[6] ——— (1959). *Theory of Value: An axiomatic Analysis of Economic Equilibrium*, volume 17, Yale University Press.

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

[8] Paolo Ghirardato, Fabio Maccheroni, and Massimo Marinacci (2004). "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118(2): 133–173.

[9] Paolo Ghirardato and Marciano Siniscalchi (2018). "Risk Sharing in the Small and in the Large," *Journal of Economic Theory*, 175: 730–765.

[10] Itzhak Gilboa, Larry Samuelson, and David Schmeidler (2014). "No-Betting-Pareto Dominance," *Econometrica*, 82(4): 1405–1442.

[11] Itzhak Gilboa and David Schmeidler (1989). "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18(2): 141–153.

[12] Lars Peter Hansen and Thomas J. Sargent (2001). "Robust Control and Model Uncertainty," *American Economic Review*, 91(2): 60–66.

[13] Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji (2005). "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73(6): 1849–1892.

[14] Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini (2006). "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74(6): 1447–1498.

[15] Man-Chung Ng (2003). "On the Duality Between Prior Beliefs and Trading Demands," *Journal of Economic Theory*, 109(1): 39–51.

[16] Luca Rigotti, Chris Shannon, and Tomasz Strzalecki (2008). "Subjective Beliefs and Ex ante Trade," *Econometrica*, 76(5): 1167–1190.

[17] Paul A. Samuelson (1956). "Social Indifference Curves," *The Quarterly Journal of Economics*, 70(1): 1–22.

[18] Oded Schramm (1988). "On the Volume of Sets Having Constant Width," *Israel Journal of Mathematics*, 63: 178–182.

[19] Menahem E. Yaari (1969). "Some Remarks on Measures of Risk Aversion and on their Uses," *Journal of Economic Theory*, 1(3): 315–329.