

Aggregate Risk and the Pareto Principle*

Nabil I. Al-Najjar[†] Luciano Pomatto[‡]

April 2018

Abstract

In the evaluation of public policies, a crucial distinction is between plans that involve purely idiosyncratic risk and policies that generate aggregate, correlated risk. While natural, this distinction is not captured by standard utilitarian aggregators.

In this paper we revisit Harsanyi's (1955) celebrated theory of preferences aggregation and develop a parsimonious generalization of utilitarianism. The theory we propose can capture sensitivity to aggregate risk in large populations and can be characterized by two simple axioms of preferences aggregation.

*We thank Larry Blume, Amrita Dhillon, Federico Echenique, Atsushi Kajii, Edi Karni, and Kota Saito for their valuable comments.

[†] Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. [Link to Research Page](#).

[‡] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, 91106. [Link to Research Page](#).

1 Introduction

Most public policies involve a mix of idiosyncratic risk (e.g. uncorrelated health accidents) and correlated, aggregate risk (e.g. the risk of pandemics). This elementary, but important, difference has been studied across several fields and has been highlighted in recent institutional debates. Examples include the debate on the Precautionary Principle as an appropriate response to catastrophic risk (Sunstein (2005)), the study of the ethical implications of public risk (Keeney (1984)) and public hazards (Fishburn (1984) and Bernard, Treich, and Rheinberger (2017)), the recent discussion on systemic risk in financial markets (see Acharya, Pedersen, Philippon, and Richardson (2017) and Adrian and Brunnermeier (2016)), and the analysis of the effect of evolution on shaping humans attitude over risk (Robson (1996)), among many others.

In this paper we seek to understand how correlated and idiosyncratic risks should be evaluated by a policy maker. It is common practice, in economics, to evaluate social prospects according to the expectation of an additive utilitarian aggregator of the form:

$$U(s) = \sum_{i \in I} u_i(s_i) \quad (1)$$

where s is a vector of outcome, s_i represents individual i 's outcome and u_i her utility function. The starting point of our analysis is the well known observation that utilitarian aggregators cannot distinguish between idiosyncratic and correlated risk, since the expectation of (1) does not depend on the degree of correlation of the variables $(s_i)_{i \in I}$.¹

Consider, for concreteness, two risky policies, labeled A and B. Assume each option will affect a large population and, for each individual, will result in either a good or bad outcome. Assume under A risk is perfectly correlated: with probability 1/2 all agents will either obtain a high or low level utility. Under B, each individual has a probability 1/2 of receiving one of the two outcomes, but these odds are independent across agents.

From the perspective of the private interests of a single individual, the two policies can be viewed as equivalent. The same conclusion is reached if the two policies are evaluated according to the expectation of the aggregator (1). However, it is not obvious that the policy maker should treat the two options in the same way. A policy maker may draw a distinction among the two based on equity concerns, social considerations, or economic motives.

One may argue that A is more equitable than B. The former will result in a perfectly equal distribution of utility, while the second will split society in two subgroups enjoying very different outcomes. Taking a different perspective, a policy maker may be concerned

¹An important caveat is in order. Throughout the paper we adopt a somewhat narrow interpretation of utilitarianism. We describe functionals as (1) as “utilitarian,” but we do not claim that this captures the whole spectrum of ideas associated with the term. For instance, we maintain the assumption that each individual utility u_i is a function only of i 's outcome. This assumption, while very common in economic models, might be inappropriate in certain contexts.

with the fact that A exposes society to the (possibly catastrophic) risk of a uniformly bad outcome. Such a scenario is guaranteed to happen with only very small probability under the alternative B, which is therefore safer. Another reason to weight the two policies differently is the fact that under policy B, society, as a whole, will not be exposed to any risk. With respect to any aggregate statistic that is a function of the distribution of outcomes in the population, the final effect of policy B is, unlike A, known in advance, up to a vanishing degree of error.

We provide a parsimonious generalization of the utilitarian criterion that can capture sensitivity to aggregate risk. We adopt an axiomatic approach and propose a theory of preference aggregation that revisits Harsanyi's (1955) foundation of utilitarianism.

We consider a standard economic environment given by a large population I of agents and, for each agent i , a set of possible outcomes X_i . Policies are identified with lotteries over profiles $s \in \prod_{i \in I} X_i$. Each individual is endowed with a preference relation \succsim_i over lotteries. These preferences must be aggregated in a social preference relation \succsim , as a guide for a policy maker. Both the individual and the social preference relations are consistent with expected utility.

In Harsanyi's theory, social and individual choices are related by a Pareto condition: if all agents prefer a lottery P to a lottery Q , then society too should rank P as more desirable than Q . While seemingly uncontroversial, Harsanyi's axiom rules out concerns for aggregate risk. In the environment we study in this paper, the axiom forces a policy maker to deem as equally desirable any two policies that induce the same individual risks, regardless of their degrees of correlation. Therefore, in this paper, we depart from Harsanyi's approach.

We relate the social preference and the individual preferences through two simple axioms. The first condition, *Restricted Pareto*, is a weakening of Harsanyi's Pareto axiom. Call *independent* a lottery P under which the individual outcomes $(s_i)_{i \in I}$ are independent random variables. So, an independent lottery describes idiosyncratic risk. The Restricted Pareto axiom requires society to prefer a lottery P to a lottery Q whenever all agents rank the first as preferable to second and both lotteries are independent.

The second axiom, *Anonymity*, posits a limit to the degree by which society can discriminate or favor different groups. Consider a group $a \subset I$ and an allocation s^a that assigns to each agent who belongs to a her *most* favorite outcome, and to each agent who does not belong to a her *worst* favorite outcome. We call such an allocation an *extreme* allocation, since it unambiguously favors a certain group to the rest of the populations. The Anonymity axiom postulates that if two groups a and b represent the same fraction of the population, then the policy maker should be indifferent between the extreme allocations s^a and s^b .

The two axioms, together with a strict version of the Restricted Pareto axiom, are satisfied if and only if the preference relation \succsim ranks social prospects according to the

expectation of the aggregator

$$U(s) = \varphi \left(\int_I u_i(s_i) d\lambda(i) \right) \quad (2)$$

where (I, λ) is a non-atomic space of agents (a primitive of the model), u_i is individual i 's von-Neumann Morgenstern utility, normalized to take values between 0 and 1, and φ is a strictly increasing transformation.

When ranking deterministic allocations, the transformation φ plays no role in the representation (2). So, in the absence of risk, a preference \succsim that satisfies the axioms is consistent with a standard utilitarian aggregator.

The ranking over lotteries, on the other hand, depends crucially on whether risk is idiosyncratic or correlated. Given a lottery P that is independent, a law of large numbers argument implies that the expectation of (2) takes the form

$$E_P[U] = \varphi \left(\int_I E_P[u_i] d\lambda(i) \right)$$

That is, policy P is evaluated by averaging the individual *expected* utilities and then applying the transformation φ . In particular, when ranking two lotteries P and Q that are both independent, the comparison between the expected social utilities $E_P[U]$ and $E_Q[U]$ does not depend on the transformation φ and, is again, consistent with the ranking of a standard utilitarian aggregator. However, unless the function φ is linear, a social preference consistent with the axioms above will display sensitivity to correlated risk.

We illustrate these features of the representation by means of our initial example. Suppose that for every agent a good outcome provides utility 1 and a bad outcome utility 0. Then the policy maker will evaluate policy A as $\frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0)$. Under policy B, almost surely, half of the agents will obtain utility 1 and half of the agents will obtain utility 0, resulting in average realized utility $\int_I u_i(s_i) d\lambda(i)$ equal to 1/2. Hence, A is preferred to B by the policy maker if and only if $\frac{1}{2}\varphi(1) + \frac{1}{2}\varphi(0) \geq \varphi(\frac{1}{2})$ holds. More generally, concavity of φ captures aversion to aggregate risk.

The aggregator we propose in this paper is formally close to utilitarianism and straightforward to apply. While the representation (2) we propose is simple, our main characterization theorem requires us to develop some new techniques. Because we focus on independent lotteries, which form a nonconvex set, we cannot apply some of the standard arguments in the literature on preference aggregation (see Border (1985) for a concise proof of Harsanyi theorem). The proof of our characterization theorem is based instead on a novel probabilistic argument.

1.1 Related Literature

An important reason to study correlation among individual risks a concern about inequality. Two generalizations of utilitarianism that capture inequality aversion are the

Generalized Utilitarian criterion, where $U(s_i) = \int_I \phi(u_i(s_i))d\lambda(i)$ (see, for instance, Adler and Sanchirico (2006)), and the Expected Equally-Distributed Equivalent-Utility representation $U(s_i) = \phi^{-1}(\int_I \phi(u_i(s_i))d\lambda(i))$, introduced by Fleurbaey (2010) and further characterized in Grant, Kajii, Polak, and Safra (2012). Both criteria extend utilitarianism by allowing for a (possibly) nonlinear transformation ϕ . In Section 5.1 we discuss in detail the relation between our work and these alternative classes of social preferences.

The Restricted Pareto axiom was studied by Keeney (1980), by Bommier and Zuber (2008) and in the context of multi-attribute decision theory (see Keeney and Raiffa (1993)). These papers obtain representations where the aggregator U is multiplicative or, more generally, multilinear with respect to the individual utility functions. In general, multiplicative or multilinear aggregators differ from additively separator aggregators even in the ranking of deterministic allocations. One of the contributions of this papers is to show that in the context of *large* populations (a natural setup for studying public risk) it is possible to capture sensitivity to aggregate risk while retaining most of the features of utilitarian aggregators.

The Anonymity axiom captures a basic principle of impartiality. While anonymity conditions are standard postulates in theories of preference aggregation (see, for instance, May (1952)), the approach we take in this paper is more directly inspired by the work of Karni (1998), Dhillon and Mertens (1999) and Börgers and Choo (2017), and Segal (2000).

In econometrics, Manski and Tetenov (2007) study optimal treatment problems under the social welfare functional (2). Their work provides support for fractional treatments as an optimal way to hedge against risk. The same social welfare functional is also discussed in Al-Najjar and Pomatto (2016) but without providing a foundation based on axioms of preference aggregation.

2 Framework

We consider a society consisting of a set I of agents and a policy maker, or *social planner*. For each agent i we are given a set X_i of *individual outcomes*.

The policy maker must choose among different policies, and each policy induces a different probability distributions over allocations of outcomes. Formally, an *allocation*, or *profile*, of outcomes is a vector $s \in \prod_{i \in I} X_i$ that assigns to each agent i an outcome $s_i \in X_i$. We denote by S the set of all profiles. Each set of outcomes X_i is endowed with a σ -algebra Σ_i containing all singletons. We denote by $\Sigma^I = \otimes_{i \in I} \Sigma_i$ the corresponding product σ -algebra. A *lottery* (or *policy*) is a σ -additive probability measure on (S, Σ^I) and $\Delta(S)$ is the set of all lotteries.

Of particular interest is the subset $\Pi(S) \subseteq \Delta(S)$ of product measures. Under a lottery $P \in \Pi(S)$ individual outcomes $(s_i)_{i \in I}$ are independent (but not necessarily identically distributed) random variables. To simplify the language, we refer to a lottery P in $\Pi(S)$

as an *independent* lottery. Independent lotteries describe idiosyncratic risk. We will identify an allocation s with the degenerate lottery that assigns probability 1 to s . Notice that a degenerate lottery is an independent lottery.

Given $P \in \Delta(S)$ and $i \in I$ let P_i be the corresponding marginal on (X_i, Σ_i) defined as $P_i(E) = P(\{s : s_i \in E\})$ for every $E \in \Sigma_i$. We consider completions of the above σ -algebras. Given a lottery P , let Σ_P^I be the completion of Σ^I with respect to P and denote by Σ the common completion $\Sigma = \bigcap_{P \in \Delta(S)} \Sigma_P^I$.

2.1 Preferences

The policy maker and the agents are expected utility maximizers. Each agent i is endowed with a binary preference relation \succsim_i over the set of lotteries that admits the representation

$$P \succsim_i Q \iff E_{P_i}[u_i] \geq E_{Q_i}[u_i]$$

where the von Neumann-Morgenstern utility function $u_i : X_i \rightarrow \mathbb{R}$ is bounded and Σ_i -measurable. Throughout the paper we maintain the assumption, prevalent in economics, that an individual's utility function depends only on her outcome.

Each utility function u_i is normalized, without loss of generality, to take value in a subset of $[0, 1]$. We assume that for each agent i there are best and worst outcomes \bar{x}_i and \underline{x}_i in X_i such that $u_i(\bar{x}_i) = 1$ and $u_i(\underline{x}_i) = 0$.

The policy maker is endowed with a social preference relation \succsim over lotteries represented as

$$P \succsim Q \iff E_P[U] \geq E_Q[U]$$

where $U : S \rightarrow \mathbb{R}$ is bounded and Σ -measurable. For any two policies P and Q , the ranking $P \succsim Q$ indicates that P is at least as desirable, from a social perspective, as Q .

2.2 The Pareto Axioms and Harsanyi Theorem

We now review the main concepts behind Harsanyi's Theorem. A basic normative tenet for aggregating individual preferences is that society should avoid Pareto dominated allocations:

Axiom (Deterministic Pareto). *For all profiles s and s' , if $s \succsim_i s'$ for every i then $s \succsim s'$.*

Harsanyi's celebrated solution to the problem of preferences aggregation is based on the key idea of extending the Pareto principle from choices among deterministic allocations to choices among lotteries:

Axiom (Extended Pareto). *For all lotteries P and Q , if $P \succsim_i Q$ for every i then $P \succsim Q$.*

In the present framework, Harsanyi Theorem can be stated as follows:²

²Zhou (1997) generalized Harsanyi Theorem to infinite populations. See also Remark 1 below.

Theorem 1 (Harsanyi). *Let I be finite. The preference relation \succsim satisfies the Extended Pareto axiom if and only if there exist $\alpha \in \mathbb{R}$ and weights $(\lambda_i)_{i \in I}$ in \mathbb{R}_+ such that*

$$U(s) = \sum_{i \in I} \lambda_i u_i(s_i) + \alpha \text{ for all } s \in S$$

Hence, a social preference relation that abides by the extended Pareto axiom can be represented by a utilitarian aggregator. In particular, as discussed in the introduction, the Extended Pareto axiom rules out sensitivity to correlated, aggregate, risk. In order to accommodate this basic disposition towards social risk, in the next section we provide a theory of preference aggregation that weakens the Extended Pareto axiom.

3 Axioms and Representation

The contrast between idiosyncratic and correlated risk is more salient in large populations, where individual idiosyncratic risks wash out at the aggregate level. In order to capture this idea we will focus on the case where the population of agents I is large. This approach will also simplify the analysis and facilitate the axiomatic derivation.

We model the population as a nonatomic space (I, λ) , where I is infinite and λ is a nonatomic probability measure defined on a collection of subsets of I . Given a group $a \subseteq I$, $\lambda(a)$ represents the fraction of agents that belong to that group.

In order to avoid the well known measurability issues that arise when dealing with an uncountable family of independent random variables,³ we assume that I is countable and λ is a finitely additive probability measure defined on the collection of all subsets of I . We also assume that the map $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$ is Σ -measurable.⁴

We now turn to describing the axioms. The first axiom restricts Harsanyi's Pareto condition to choices among *independent* lotteries.

Axiom 1 (Restricted Pareto). *For all independent lotteries P and Q , if $P \succsim_i Q$ for every i then $P \succsim Q$.*

The axiom describes a policy maker who abides by the Pareto principle as long as the policy under considerations do not generate aggregate risk, and it reflects our motivation of providing a parsimonious generalization of Harsanyi's extended Pareto axiom.

Under the Extended Pareto axiom, two policies that induce the same marginal distributions are deemed equally desirable by the policy maker. This is true regardless of whether the welfare effect of the two policies is perfectly correlated or completely

³See Judd (1985), and Al-Najjar (2008) among many others.

⁴The latter assumption ensures that expectations with respect to a lottery P that involve integrals with respect to λ are well-defined. Lemma 1 in the Appendix formalizes this claim. The Appendix also contains an existence result. The choice of a countable set of agents and a finitely additive measure allows us to keep the analysis mathematically rigorous, but nothing is lost in terms of intuition by taking I to be the interval $[0, 1]$ and λ the Lebesgue measure.

idiosyncratic, hence regardless of whether the different policies involve risk only at the individual level or also at the societal level. The Restricted Pareto axiom avoids this strong conclusion by applying the Pareto principle only to the ranking of independent lotteries.

The next condition is a strengthening of axiom 1. For every $\alpha \in [0, 1]$, we denote by P^α the independent lottery where each agent i receives her best outcome \bar{x}_i with probability α and the worst outcome \underline{x}_i with probability $1 - \alpha$. The Restricted Pareto axiom ensures that the social preference \succsim is monotone in the odds α of a good outcome, in the sense that $P^\alpha \succsim P^\beta$ whenever $1 \geq \alpha \geq \beta \geq 0$. The next axiom strengthen this property to a strict form of monotonicity.

Axiom 2 (Strict Pareto). *If $\alpha > \beta$ then $P^\alpha \succ P^\beta$.*

The next axiom imposes some limit to the degree by which different groups can be favored or discriminated against by the policy maker.

Given a set $a \subseteq I$ of agents, we denote by s^a the profile defined for every agent i as $s_i^a = \bar{x}_i$ if $i \in a$ and $s_i^a = \underline{x}_i$ if $i \notin a$. Hence, the allocation s^a assigns to every agent her best outcome if she belongs to group a and her worst outcome otherwise. We refer to each s^a as an *extreme* allocation, as in many environments, whether the set of outcomes X_i represents income, consumption bundles or health levels (ranging from “a life-threatening health condition” to “being in perfect health”) an extreme allocation s^a unambiguously favors group a relative to the rest of the population.

The next axiom requires groups that represent equal fractions of the population to be treated in the same way, at least in the context of choices between extreme allocations.

Axiom 3 (Anonymity). *If $\lambda(a) = \lambda(b)$ then $s^a \sim s^b$.*

Postulates related to axiom 3 have been discussed, in different contexts, by Karni (1998), Dhillon and Mertens (1999), Segal (2000), and Piacquadio (2017), among others.

3.1 Representation

We can now present the main result of the paper.

Theorem 2. *The preference relation \succsim satisfies axioms 1-3 if and only if there exists a strictly increasing function $\varphi : [0, 1] \rightarrow \mathbb{R}$ such that*

$$U(s) = \varphi \left(\int_I u_i(s_i) d\lambda(i) \right) \quad \text{for all } s \in S. \quad (3)$$

When confined to the ranking of deterministic profiles, the ordinal ranking described by the function U is unaffected by the strictly increasing transformation φ . Hence, in the absence of risk, the aggregator U is indistinguishable from an additively separable aggregator. The ranking of lotteries, on the other hand, crucially hinges on whether risk is independent or correlated. We illustrate this and other properties of the representation in the next section.

4 Idiosyncratic and Correlated Risk

We first consider the case where risk is purely idiosyncratic. Our analysis relies on the following law of large numbers:

Theorem 3. *Let U satisfy the representation (3). Then, for every independent lottery P ,*

$$U(s) = \varphi \left(\int_I u_i(s_i) d\lambda(i) \right) = \varphi \left(\int_I E_{P_i}[u_i] d\lambda(i) \right)$$

for P -almost every profile $s \in S$.

Theorem 3 establishes that for almost all realized allocations s , the weighted average realized utility $\int_I u_i(s_i) d\lambda(i)$ will equal the weighted average expected utility $\int_I E_{P_i}[u_i] d\lambda(i)$. This fact formalizes the idea that from the perspective of the policy maker, randomness vanishes in a large population under idiosyncratic risk.

The result implies that the expected social utility with respect to an independent lottery P is given by the expression:

$$E_P[U] = \varphi \left(\int_I E_{P_i}[u_i] d\lambda(i) \right)$$

In particular, the social ranking of two independent lotteries resembles the ranking of a standard utilitarian criterion:

Corollary 1. *Let U satisfy the representation (3). Given two independent lotteries P and Q ,*

$$E_P[U] \geq E_Q[U] \iff \int_I E_{P_i}[u_i] d\lambda(i) \geq \int_I E_{Q_i}[u_i] d\lambda(i)$$

As in the case of a ranking between deterministic allocations, the curvature of the transformation φ plays no role in the evaluation of idiosyncratic risks.

We now study the case where risk has an aggregate component. A simple instance of this class of risky policies is given by a lottery P^* that is a convex combination

$$P^* = \alpha Q + (1 - \alpha)R$$

between two independent lotteries Q and R . The lottery P^* admits a straightforward interpretation. Consider a policy (for instance, a drug) whose effect on the population is known to be distributed according to one of two possible independent distributions, Q or R . Assume that the final distribution is determined by the realization of a common aggregate shock represented by a binary random variable (e.g. whether or not the drug is effective). Then the policy can be represented by the distribution P^* .

It follows from Theorem 3 that the planner's expected utility with respect to the lottery P^* is given by:

$$E_{P^*}[U] = \alpha \varphi \left(\int_I E_{Q_i}[u_i] d\lambda(i) \right) + (1 - \alpha) \varphi \left(\int_I E_{R_i}[u_i] d\lambda(i) \right)$$

The expression makes clear that the non-linearity of φ plays a key role in the evaluation of correlated lotteries. Indeed, as we record in the next remark, non-linearity implies a violation of the Harsanyi Extended Pareto axiom.

Remark 1. *If \succsim satisfies the Extended Pareto axiom then φ is affine.*

4.1 Conditionally i.i.d. Lotteries

We now consider a canonical class of distributions that display a mix of aggregate and idiosyncratic risks. We assume a common set of individual outcomes X , so that $X_i = X$ for every i . We denote by $\Delta(X)$ the set of probability measures on X . An independent lottery is *i.i.d.* if there is a single probability measure over outcomes $\theta \in \Delta(X)$ that satisfies $P_i = \theta$ for every i . Given θ , we denote by P^θ the corresponding i.i.d. lottery. So, P^θ is an independent lottery with marginal θ common to all agents.

A lottery $P^\mu \in \Delta(S)$ is *conditionally i.i.d. with hyper-parameter μ* if μ is a probability over $\Delta(X)$ with finite support and P^μ is the mixture

$$P^\mu = \sum_{\theta \in \Delta(X)} \mu(\theta) P^\theta$$

Informally, P^μ is a mixture of i.i.d. distribution where the parameter θ is unknown and distributed according to μ .

This class of lotteries is widely used in applications for their tractability and because they provide a clear separation between an aggregate common shock that determines the parameter θ , and purely idiosyncratic individual shocks distributed according to P^θ (conditional on the realized θ).⁵

The following corollary of Theorem 3 gives a convenient formula for evaluating the aggregative utility of conditionally i.i.d. lotteries in homogeneous populations:

Corollary 2. *Assume there is a common utility function $u : X \rightarrow \mathbb{R}$ such that $u = u_i$ for every i . Then, for any conditionally i.i.d. lottery P^μ ,*

$$E_{P^\mu}[U] = \sum_{\theta \in \Delta(X)} \mu(\theta) \varphi(E_\theta[u]) \tag{4}$$

Concavity, or convexity, of the transformation φ characterize society's aversion to, or preference for, correlated risk. Given a conditionally i.i.d. lottery P^μ , define $\theta^\mu \in \Delta(X)$ as the mixture

$$\theta^\mu = \sum_{\theta \in \text{supp}(\mu)} \mu(\theta) \theta$$

The social ranking between the conditionally i.i.d. lottery P^μ and the idiosyncratic i.i.d. lottery P^{θ^μ} is crucial in understanding how the policy maker tradeoffs adherence

⁵By de Finetti Theorem, conditional i.i.d. distributions can be characterized axiomatically as the only lotteries that are invariant with respect to any permutation of the agents.

to the Pareto principle and exposure to correlated risk. The two lotteries have identical marginal distributions. Hence, under expected utility, agents must be indifferent between the two, *regardless* of their von Neumann Morgensten utility function, i.e. we must have $P^\mu \sim_i P^{\theta^\mu}$ for every i . Therefore, a strict ranking of the policy maker's preference between the two lotteries reveals a violation of Harsanyi's Extended Pareto axiom.

A social planner who ranks $P^{\theta^\mu} \succ P^\mu$ violates Harsanyi's Extended Pareto axiom by preferring idiosyncratic to correlated risk. An opposite conclusion applies whenever $P^\mu \succ P^{\theta^\mu}$.

We say that the social preference \succsim is *averse to aggregate risk* if $P^{\theta^\mu} \succsim P^\mu$ for every μ . It is *averse to idiosyncratic risk* if $P^\mu \succsim P^{\theta^\mu}$ for every μ . As we record in the next corollary, the social planner's attitude towards social risk admits a straightforward characterization in terms of the concavity of the transformation φ .

Corollary 3. *The social preference \succsim is averse to aggregate risk if and only if φ is concave. It is averse to idiosyncratic risk if and only if φ is convex.*

5 Discussion and Extensions

5.1 Comparison with other Generalizations of Utilitarianism

Here, we compare our work to other related generalizations of utilitarianism. The well-known *Generalized Utilitarian* criterion (see, for instance, Adler and Sanichirico (2006) and Grant, Kajii, Polak, and Safra (2010))

$$U(s) = \int_I \phi(u_i(s_i)) d\lambda(i) \quad (5)$$

can capture aversion to ex-post inequality by applying a concave transformation ϕ to the individual utilities. Generalized utilitarianism cannot, however, capture sensitivity to correlation, since the expectation of the aggregator U in (5) does not depend on the correlation between the (s_i) 's.

Fleurbaey (2010) introduced the *Expected Equally-Distributed Equivalent-Utility* criterion (henceforth, EEDEU) which, in our setting, takes the form

$$U(s) = \phi^{-1} \left(\int_I \phi(u_i(s_i)) d\lambda(i) \right). \quad (6)$$

The representation displays aversion to inequality aversion if and only if the transformation ϕ is concave. Concavity of ϕ translates, by the resulting convexity of ϕ^{-1} , into a social preference that is *averse* to idiosyncratic risk. To illustrate, consider a homogeneous population where all agents have the same utility function u . If ϕ is concave, then for every conditionally i.i.d. lottery P^μ , Jensen's inequality and Theorem 3 imply

$$E_{P^\mu}[U] = \sum_{\theta \in \Delta(X)} \mu(\theta) \phi^{-1}(E_\theta[\phi(u)]) \geq \phi^{-1} \left(\sum_{\theta \in \Delta(X)} \mu(\theta) E_\theta[\phi(u)] \right) = E_{\theta^\mu}[U]$$

So, under the EEDEU criterion, a social preference can exhibit aversion to correlated risk when ranking lotteries if and only if it favors inequity when ranking deterministic allocations. We do not see this as a shortcoming of the EEDEU, but as an illustration of the tension between two conflicting goals: reducing ex post inequality and mitigating correlated risk. Decreasing the impact of aggregate uncertainty may require a “diversification” across individuals that exacerbates ex post welfare differences.

The two approaches may be unified by a more general criterion $U(s_i) = (\varphi \circ \phi^{-1})(\int_I \phi(u_i(s_i))d\lambda(i))$. This more general class of social preferences would be able to display, simultaneously, both ex-post inequity aversion and aversion to correlated risk. We do not know to what extent the analysis in this paper and in Fleurbaey (2010) can be adapted to provide an axiomatic foundation for this more general class of social preferences.

5.2 Interpersonal Comparison of Utilities

In Theorem 2, the weight $\lambda(a)$ represents the fraction of agent in the population that belong to group a . In Harsanyi theorem, in contrast, the weight $\lambda(a)$ is derived from the social preference relation and is a subjective component of the representation. Thus, different policy makers can satisfy Harsanyi axioms and yet attribute different social weights to the same group of individuals. This particular feature of Harsanyi theorem has been the subject of considerable scrutiny and critiques.

In this section we now show how Harsanyi’s approach can be integrated in our analysis by weakening the Anonymity axiom. We consider a social preference relation that satisfies the following three axioms.

The first condition requires each individual to be negligible.

Axiom a. *Fix $j \in I$. If $s_i = s'_i$ for every i other than j , then $s \sim s'$.*

The next requirement is a continuity assumption.

Axiom b. *For every s , $\{\alpha : P^\alpha \succ s\}$ and $\{\alpha : s \succ P^\alpha\}$ are open subsets of $[0, 1]$.*

The final axiom is a more substantive condition. It expresses the following logic: whenever society is facing a choice between two extreme allocations s^a and s^b , the policy maker should choose by taking into account only those agents who are affected by the decision. The axiom is formally equivalent to de Finetti’s (1931) celebrated notion of qualitative probabilities.

Axiom c. *If $(a \cup b) \cap c = \emptyset$ then $s^a \succsim s^b \iff s^{a \cup c} \succsim s^{b \cup c}$.*

Notice that when choosing between s^a and s^b , or between $s^{a \cup c}$ and $s^{b \cup c}$, in both scenarios the final choice is inconsequential for agents who belong to the disjoint group c .

The axiom demands groups who do not have stakes in a decision over extreme allocations to not play a role in determining what allocation will be implemented.

For the next result, we denote by $\mathcal{P}(I)$ the collection of all subsets of I .

Theorem 4. *The preference relation \succsim satisfies axioms 1,2 and a-c if and only if there exists a strictly increasing function $\varphi : [0, 1] \rightarrow \mathbb{R}$ and a nonatomic finitely additive probability $\tilde{\lambda}$ defined on $\mathcal{P}(I)$ such that*

$$U(s) = \varphi \left(\int_I u_i(s_i) d\tilde{\lambda}(i) \right) \quad \text{for all } s \in S.$$

All the results in the paper (including the analysis of Section 4) continue to hold when the measure $\tilde{\lambda}$, derived from the preference \succsim , is substituted to the original measure λ .⁶

Theorem 4 contributes to the literature on the representation of qualitative probabilities. Our result differs from the existing literature (de Finetti (1931), Savage (1972), Niiniluoto (1972), Wakker (1981) and Gilboa (1985), among others) in two main ways. The additional probabilistic structure available in our framework allows substituting Savage's assumptions of *fine* and *tight* qualitative probability by the simple axioms a-b and to provide a proof that is concise and almost self-contained.

5.3 Finite Populations

The social welfare functional introduced in this paper can be easily applied to finite large populations. For every n , consider a finite population $I_n \subseteq I$ of size n and the social welfare functional defined as

$$U_n(s) = \varphi \left(\frac{1}{n} \sum_{i \in I_n} u_i(s_i) \right)$$

The functional U^n is a discretization of the representation in Theorem 2. It satisfies the Restricted Pareto axiom asymptotically as the size of the population grows to infinity, up to a vanishing degree of error.

Theorem 5. *Let φ be continuously differentiable. There exists a sequence $(\epsilon_n) \downarrow 0$ such that for every pair of independent lotteries P and Q ,*

$$\text{if } E_{P_i}[u_i] \geq E_{Q_i}[u_i] \text{ for every } i \in I_n \text{ then } E_P[U_n] \geq E_Q[U_n] - \epsilon_n$$

Notice that the same error term ϵ_n applies uniformly over all pairs of independent lotteries P and Q . The result follows from a concentration of measure argument. As we show in the Appendix (see lemma 5) the law of large numbers described in theorem 3 continues to hold, asymptotically, for finite populations.

⁶Formally, this follows from the fact that the only assumptions imposed on λ are that it is non-atomic and satisfies a measurability conditions. Both assumptions are satisfied by a measure λ^* obtained through Theorem 4.

5.4 Non-Expected Utility

We now extend our main result to non-expected utility preferences. Our aggregation theorem continue to hold under very minimal assumption on the social planner's preference over lotteries. We now consider a policy maker endowed with a binary preference relation \succsim over lotteries which does not necessarily satisfy the von-Neumann Morgenstern axioms.

We impose two basic axioms on \succsim . For the next condition, we denote by $\succsim|_S$ the restriction of \succsim over the set of all deterministic allocations.

Axiom I. *The preference $\succsim|_S$ is complete, transitive and S contains a countable $\succsim|_S$ -order-dense subset.⁷ For every $s^* \in S$, the sets $\{s : s \succsim s^*\}$ and $\{s : s^* \succsim s\}$ are Σ -measurable.*

The axiom is equivalent to the existence of a social utility function $U : S \rightarrow \mathbb{R}$ that represents $\succsim|_S$ and is Σ -measurable and bounded. The next axiom requires \succsim to satisfy a basic form of stochastic dominance.

Axiom II. *Let $P, Q \in \Delta(S)$. If $s', s'' \in S$ are such that*

$$P(\{s : s \sim s'\}) = Q(\{s : s \sim s''\}) = 1 \quad (7)$$

then $P \succsim Q$ if and only if $s' \succsim s''$.

The two axioms are compatible with several models of decision under risk. Notice that axiom II only has bite over lotteries that satisfy (7). Thus, it does not require the social preference \succsim to be complete, or even transitive, over the whole domain of lotteries. The axiom is compatible with a preference that ranks lotteries according to the expectation and the variance of U , as well as with rank dependent preferences. The next result shows how our main result extends to any social preference relation consistent with axioms I and II.

Theorem 6. *Let \succsim be a binary preference relation on $\Delta(S)$ that satisfies axioms I and II. Let $U : S \rightarrow \mathbb{R}$ be a bounded and Σ -measurable function that represents $\succsim|_S$. Then, \succsim satisfies axioms 1-3 if and only if there exists a strictly increasing function $\varphi : [0, 1] \rightarrow \mathbb{R}$ such that*

$$U(s) = \varphi \left(\int_I u_i(s_i) d\lambda(i) \right) \quad \text{for all } s \in S.$$

The result generalizes Theorem 2. Any preference \succsim that satisfies axioms I and II must, under the Anonymity and the Restricted Pareto axioms, rank deterministic allocations in a utilitarian way. It can be shown that the conclusions of Corollary 1 continue to hold under this more general framework. In particular, the ranking of independent lotteries remains consistent with the expectation of a standard utilitarian aggregator.

⁷That is, there exists a countable set $T \subset S$ such that for all $s, s' \in S$, if $s \succ s'$ there exists $t \in T$ such that $s \succ t \succ s'$.

References

- ACHARYA, V. V., L. H. PEDERSEN, T. PHILIPPON, AND M. RICHARDSON (2017): “Measuring systemic risk,” *The Review of Financial Studies*, 30(1), 2–47.
- ADLER, M. D., AND C. W. SANCHIRICO (2006): “Inequality and uncertainty: Theory and legal applications,” *University of Pennsylvania Law Review*, pp. 279–377.
- ADRIAN, T., AND M. K. BRUNNERMEIER (2016): “CoVaR,” *The American Economic Review*, 106(7), 1705–1741.
- AL-NAJJAR, N. I. (2008): “Large games and the law of large numbers,” *Games and Economic Behavior*, 64(1), 1–34.
- AL-NAJJAR, N. I., AND L. POMATTO (2016): “Choice under aggregate uncertainty,” *Theory and Decision*, 80(2), 187–209.
- (2017): “An abstract law of large numbers,” *mimeo*.
- BERNARD, C., N. TREICH, AND C. RHEINBERGER (2017): “Catastrophe aversion and risk equity under dependent risks,” *Management Science*, forthcoming.
- BOMMIER, A., AND S. ZUBER (2008): “Can preferences for catastrophe avoidance reconcile social discounting with intergenerational equity?,” *Social Choice and Welfare*, 31(3), 415–434.
- BORDER, K. C. (1985): “More on Harsanyi’s utilitarian cardinal welfare theorem,” *Social Choice and Welfare*, 1(4), 279–281.
- BÖRGERS, T., AND Y.-M. CHOO (2017): “Revealed Relative Utilitarianism,” *University of Michigan*.
- DE FINETTI, B. (1931): “Sul significato soggettivo della probabilità,” *Fundamenta Mathematicae*, 17, 298–329.
- DHILLON, A., AND J.-F. MERTENS (1999): “Relative utilitarianism,” *Econometrica*, 67(3), 471–498.
- FISHBURN, P. C. (1984): “Equity axioms for public risks,” *Operations Research*, 32(4), 901–908.
- FLEURBAEY, M. (2010): “Assessing risky social situations,” *Journal of Political Economy*, 118(4), 649–680.
- FREMLIN, D. H. (2008): *Measure theory. Volume 5, Set-theoretic measure theory*. Torres Fremlin.

- GILBOA, I. (1985): “Subjective distortions of probabilities and non-additive probabilities,” *Foerder Institute for Economic Research, Tel Aviv University*.
- GRANT, S., A. KAJII, B. POLAK, AND Z. SAFRA (2010): “Generalized utilitarianism and Harsanyi’s impartial observer theorem,” *Econometrica*, 78(6), 1939–1971.
- (2012): “Equally-distributed equivalent utility, ex post egalitarianism and utilitarianism,” *Journal of Economic Theory*, 147(4), 1545–1571.
- HARSANYI, J. C. (1955): “Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility,” *Journal of political economy*, 63(4), 309–321.
- JUDD, K. L. (1985): “The law of large numbers with a continuum of iid random variables,” *Journal of Economic theory*, 35(1), 19–25.
- KARNI, E. (1998): “Impartiality: definition and representation,” *Econometrica*, 66(6), 1405–1415.
- KEENEY, R. L. (1980): “Equity and public risk,” *Operations Research*, 28(3-part-i), 527–534.
- (1984): “Ethics, decision analysis, and public risk,” *Risk Analysis*, 4(2), 117–129.
- KEENEY, R. L., AND H. RAIFFA (1993): *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.
- MANSKI, C. F., AND A. TETENOV (2007): “Admissible treatment rules for a risk-averse planner with experimental data on an innovation,” *Journal of Statistical Planning and Inference*, 137(6), 1998–2010.
- MAY, K. O. (1952): “A set of independent necessary and sufficient conditions for simple majority decision,” *Econometrica: Journal of the Econometric Society*, pp. 680–684.
- MCDIARMID, C. (1998): “Concentration,” in *Probabilistic methods for algorithmic discrete mathematics*, pp. 195–248. Springer.
- NIINILUOTO, I. (1972): “A note on fine and tight qualitative probabilities,” *The Annals of Mathematical Statistics*, pp. 1581–1591.
- PIACQUADIO, P. G. (2017): “A fairness justification of utilitarianism,” *Econometrica*, 85(4), 1261–1276.
- ROBSON, A. J. (1996): “A biological basis for expected and non-expected utility,” *Journal of economic theory*, 68(2), 397–424.
- SAVAGE, L. J. (1972): *The foundations of statistics*. Dover.

- SEGAL, U. (2000): “Let’s agree that all dictatorships are equally bad,” *Journal of Political Economy*, 108(3), 569–589.
- SUNSTEIN, C. R. (2005): *Laws of fear: Beyond the precautionary principle*, vol. 6. Cambridge University Press.
- WAKKER, P. (1981): “Agreeing probability measures for comparative probability structures,” *The Annals of Statistics*, 9(3), 658–662.
- ZHOU, L. (1997): “Harsanyi’s utilitarianism theorems: general societies,” *Journal of Economic Theory*, 72(1), 198–207.

A Appendix

A.1 Technical Preliminaries

As is well known, nonatomic population models lead to some measure theoretic subtleties. In this paper, a first difficulty consists in making sure that the expectation of the average $\int_I u_i(s_i)d\lambda(i)$ is well-defined with respect to any lottery P . An additional difficulty is establishing the law of large numbers property described in Theorem 3.

It is common in applications to model large populations as the interval $[0, 1]$ endowed with the standard Lebesgue measure; to assume, as a useful heuristic, that results such as Theorem 3 hold; and to omit the measurability issues that arise with a continuum of random variables. It is also natural in many problems to restrict the attention to conditionally i.i.d. social lotteries, for which payoffs can be directly computed using Corollary 2, without any reference to λ .

In this section we provide technical results that allows us to address the aforementioned measurability issues while keeping the analysis rigorous. Recall that $\mathcal{P}(I)$ denotes the collection of all subsets of I .

Lemma 1. *Let λ be a nonatomic finitely additive probability defined on $\mathcal{P}(I)$. Consider the following properties:*

1. $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$, $s \in S$, is Σ -measurable;
2. $s \mapsto \int_I u_i(s_i)d\lambda(i)$, $s \in S$, is Σ -measurable;
3. $\xi \mapsto \lambda(\{i : \xi_i = 1\})$, $\xi \in \{0, 1\}^I$, is universally measurable.⁸

The following hold: (1) \implies (2) and (1) \iff (3).

Proof: (1) implies (2). Fix a Borel set $A \subseteq [0, 1]$ and, for each i , the function $\phi_i : X_i \rightarrow \mathbb{R}$ defined as $\phi_i(x_i) = 1_A(u_i(x_i))$, where 1_A is the indicator function of A . Each ϕ_i is Σ_i -measurable. The function $\phi : S \rightarrow S$ defined as $\phi(s) = (\phi_i(s_i))$ is then $\Sigma^I \setminus \Sigma^I$ -measurable. We claim that ϕ is also $\Sigma \setminus \Sigma$ -measurable. To see this, let $E \in \Sigma$ and $P \in \Delta(S)$. Define $Q \in \Delta(S)$ as the pushforward measure $Q = P\phi^{-1}$. Then $E \in \Sigma_Q^I$, hence there exist $E_1, E_2 \in \Sigma^I$ such that $E_1 \subseteq E \subseteq E_2$ and $Q(E_2) = Q(E_1)$. Hence $\phi^{-1}(E_1) \subseteq \phi^{-1}(E) \subseteq \phi^{-1}(E_2)$ and $P(\phi^{-1}(E_1)) = P(\phi^{-1}(E_2))$. Hence $\phi^{-1}(E) \in \Sigma_P^I$. So, $\phi^{-1}(E) \in \Sigma$. It follows that the composition $s \mapsto \lambda(\{i : \phi(s)_i = \bar{x}_i\})$ is Σ -measurable. Equivalently, $s \mapsto \int_I 1_A(u_i(s_i))d\lambda(i)$ is Σ -measurable. The linearity of the integral with respect to λ implies that for every partition A_1, \dots, A_n of $[0, 1]$ and all $\alpha_1, \dots, \alpha_n$ in $[0, 1]$, the function $s \mapsto \int_I \sum_{k=1}^n \alpha_k 1_{A_k}(u_i(s_i))d\lambda(i)$ is Σ -measurable. For every n , let $f_n : [0, 1] \rightarrow [0, 1]$ be a function with finite range such that $|f(t) - t| \leq 1/n$ for every $t \in [0, 1]$. Then $s \mapsto \int_I f_n(u_i(s_i))d\lambda(i)$ is Σ -measurable. For every s ,

$$\left| \int_I f_n(u_i(s_i))d\lambda(i) - \int_I u_i(s_i)d\lambda(i) \right| \leq \int_I |f_n(u_i(s_i)) - u_i(s_i)|d\lambda(i) \leq 1/n$$

⁸The space $\{0, 1\}^I$ is endowed with the product topology.

Hence $s \mapsto \int_I (u_i(s_i)) d\lambda(i)$ is the limit of a sequence of Σ -measurable functions. Hence it is Σ -measurable.

(1) implies (3). We denote by $\mathcal{B}(\{0, 1\}^I)$ and $\mathcal{B}_{um}(\{0, 1\}^I)$ the collections of, respectively, Borel and universally measurable subsets of $\{0, 1\}^I$. Let $f_i : \{0, 1\} \rightarrow X_i$ be defined as $f_i(1) = \bar{x}_i$ and $f_i(0) = \underline{x}_i$. Then f_i is measurable. Let $f : \{0, 1\}^I \rightarrow S$ be defined as $f(\xi) = (f_i(\xi_i))_{i \in I}$ for all $\xi \in \{0, 1\}^I$. By standard arguments f is $\mathcal{B}(\{0, 1\}^I) \setminus \Sigma^I$ -measurable. By replicating the argument applied in the first part of the proof we obtain that f is also $\mathcal{B}_{um}(\{0, 1\}^I) \setminus \Sigma$ -measurable. Let $l : S \rightarrow \mathbb{R}$ be defined as $l(s) = \lambda(\{i : s_i = \bar{x}_i\})$. The composition $l \circ f$ is $\mathcal{B}_{um}(\{0, 1\}^I)$ -measurable. For all $\xi \in \{0, 1\}^I$, it satisfies

$$l(f(\xi)) = \lambda(\{i : f_i(\xi_i) = \bar{x}_i\}) = \lambda(\{i : \xi_i = 1\}).$$

(3) implies (1). For every i , consider the map $\phi_i : X_i \rightarrow \{0, 1\}$ defined as the indicator function of \bar{x}_i , and define $\phi : S \rightarrow \mathbb{R}$ as $\phi(s) = (\phi_i(s_i))$ for all s . Then ϕ is $\Sigma^I \setminus \mathcal{B}(\{0, 1\}^I)$ -measurable. As before, the same argument applied in the first part of the proof shows that ϕ is $\Sigma \setminus \mathcal{B}_{um}(\{0, 1\}^I)$ -measurable. The map $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$ is the composition of ϕ and $\xi \mapsto \int_I \xi_i d\lambda(i)$ and is therefore Σ -measurable. \blacksquare

By Lemma 1, any λ such that $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$ is Σ -measurable guarantees that the expectation of $\int_I u_i(s_i) d\lambda(i)$ is well defined with respect to any lottery P .

The next theorem, a direct corollary of a result by Fremlin, establishes the existence of a nonatomic probability λ that satisfies the appropriate measurability properties under an additional set theoretic axiom. Let \mathfrak{c} denote the cardinality of the continuum.

Axiom (P) The interval $[0, 1]$ cannot be covered by less than \mathfrak{c} meager sets.

As implied by the Baire category theorem, the interval $[0, 1]$ cannot be covered by countably many meager sets. Axiom P strengthens this conclusion to any collection of meager sets which cardinality is less than the continuum. In particular, it is implied by the Continuum Hypothesis.⁹ The result follows directly from Theorem 538S in Fremlin (2008) and lemma 1.

Theorem 7. *Under Axiom P there exists a nonatomic finitely additive probability λ defined on $\mathcal{P}(I)$ such that $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$, $s \in S$, is Σ -measurable.*

The use of set theoretic assumption may appear peculiar. Substantively, our view is that decision theoretic and economic considerations dictate the choice of a mathematical structure, not the other way around. A decision maker is justified to question whether expected utility or additive separability are appropriate on economic, ethical, or other normative grounds. But modelers and practitioners who accept the stylized nature of

⁹In fact, it is implied by Martin's Axiom, which is weaker than the continuum hypothesis.

abstract models should not, and are unlikely to, take a stand on the status of the axioms of set theory. For example, one may disagree with Savage's theory of decision making under uncertainty on substantive grounds, but usually not because it requires an *infinite* set of states of the world. Axiom P is just what it takes to make the analysis mathematically consistent.

The next result is a law of large numbers for abstract nonatomic population models.

Theorem 8. *Let λ be a nonatomic finitely additive probability defined on $\mathcal{P}(I)$ such that $s \mapsto \lambda(\{i : s_i = \bar{x}_i\})$ is Σ -measurable. Then, for every independent lottery P ,*

$$\int_I u_i(s_i) d\lambda(i) = \int_I E_P[u_i] d\lambda(i) \quad P\text{-a.s.}$$

Proof: By Lemma 1 the map $\xi \mapsto \lambda(\{i : \xi_i = 1\})$, $\xi \in \{0, 1\}^I$, is universally measurable. The result now follows from Theorem 1 in Al-Najjar and Pomatto (2017). ■

Given any independent lottery, the realized average $\int_I u_i(s_i) d\lambda(i)$ is almost surely equal to the average expectation $\int_I E_P[u_i] d\lambda(i)$.

A.2 Proof of Theorems 2 and 6

Since Theorem 2 is a special case of Theorem 6 it is sufficient to prove the latter. To this end, we fix a preference relation \succsim defined on $\Delta(S)$ that satisfies the two axioms introduced in Section 5.4. In particular, we fix a bounded, Σ -measurable function U such that $s \succsim s'$ iff $U(s) \geq U(s')$ for all $s, s' \in S$, and assume that \succsim satisfies axiom II.

We first show the sufficiency of the axioms. By the Anonymity axiom, if $\lambda(a) = \lambda(b)$ then $U(s^a) = U(s^b)$. Hence, there exists a function $\varphi : [0, 1] \rightarrow \mathbb{R}$ such that $U(s^a) = \varphi(\lambda(a))$ for every $a \subseteq I$.

We now show that φ is strictly increasing. Let $\alpha \in (0, 1)$. By Theorem 8

$$\lambda(\{i : s_i = \bar{x}_i\}) = \int_I u_i(s_i) d\lambda(i) = \int_I E_{P^\alpha}[u_i] d\lambda(i) = \alpha$$

for P^α -almost every $s \in S$. We can therefore conclude that for every $\alpha \in [0, 1]$,

$$P^\alpha(\{s^a : \lambda(a) = \alpha\}) = 1 \tag{8}$$

Now let $1 \geq \alpha > \beta \geq 0$. By (8) we can find two subsets $c, d \subseteq I$ such that $\lambda(c) = \alpha$ and $\lambda(d) = \beta$. In addition,

$$P^\alpha(\{s^a : \lambda(a) = \alpha\}) = P^\beta(\{s^a : \lambda(a) = \beta\}) = 1.$$

Hence

$$P^\alpha(\{s^a : \varphi(\lambda(a)) = \varphi(\lambda(c))\}) = P^\beta(\{s^a : \varphi(\lambda(a)) = \varphi(\lambda(d))\}) = 1.$$

So,

$$P^\alpha(\{s^a : s^a \sim s^c\}) = P^\beta(\{s^a : s^a \sim s^d\}) = 1.$$

By the strict Pareto axiom, $P^\alpha \succ P^\beta$. Hence, by axiom II, $s^c \succ s^d$. So, $\varphi(\lambda(c)) = \varphi(\alpha) > \varphi(\beta) = \varphi(\lambda(d))$. It follows that φ is strictly increasing.

Now fix a profile $\tilde{s} \in S$. Let P be the independent lottery defined so that for each i the marginal $P_i \in \Delta(X_i)$ satisfies

$$P_i(\{\bar{x}_i\}) = u_i(\tilde{s}_i) \text{ and } P_i(\{\underline{x}_i\}) = 1 - u_i(\tilde{s}_i)$$

for every $i \in I$. By the Restricted Pareto axiom, $\tilde{s} \sim P$.

By Theorem 8 we have

$$\int_I u_i(s_i) d\lambda(i) = \int_I E_P[u_i] d\lambda(i)$$

for P -almost every profile s . Notice that $\int_I u_i(s_i) d\lambda(i) = \lambda(\{i : s_i = \bar{x}_i\})$ for P -almost every profile s and $\int_I E_P[u_i] d\lambda(i) = \int_I u_i(\tilde{s}_i) d\lambda(i)$ by construction. Since P assigns probability 1 to extreme allocations, we obtain

$$P\left(\left\{s^a : \lambda(a) = \int_I u_i(\tilde{s}_i) d\lambda(i)\right\}\right) = 1.$$

Fix a set $\tilde{a} \subseteq I$ such that $\lambda(\tilde{a}) = \int_I u_i(\tilde{s}_i) d\lambda(i)$. Then

$$P(\{s^a : \varphi(\lambda(a)) = \varphi(\lambda(\tilde{a}))\}) = 1.$$

Equivalently, $P(\{s^a : s^a \sim s^{\tilde{a}}\}) = 1$. Since $P \sim \tilde{s}$, axiom II implies $s^{\tilde{a}} \sim \tilde{s}$.

Therefore,

$$U(\tilde{s}) = U(s^{\tilde{a}}) = \varphi(\lambda(\tilde{a})) = \varphi\left(\int_I u_i(\tilde{s}_i) d\lambda(i)\right).$$

Because \tilde{s} is arbitrary, this concludes the proof of sufficiency. We now turn to the proof of necessity. We first show that the Restricted Pareto axiom is implied by the representation. Consider two independent lotteries P and Q such that $P \succsim_i Q$ for every i . We can apply Theorem 8 to conclude

$$P\left(\left\{s : \int_I u_i(s_i) d\lambda(i) = \int_I E_{P_i}[u_i] d\lambda(i)\right\}\right) = 1$$

$$Q\left(\left\{s : \int_I u_i(s_i) d\lambda(i) = \int_I E_{Q_i}[u_i] d\lambda(i)\right\}\right) = 1$$

Fix two profiles s' and s'' such that $\int_I u_i(s'_i) d\lambda(i) = \int_I E_{P_i}[u_i] d\lambda(i)$ and $\int_I u_i(s''_i) d\lambda(i) = \int_I E_{Q_i}[u_i] d\lambda(i)$. Then

$$1 = P\left(\left\{s : \varphi\left(\int_I u_i(s_i) d\lambda(i)\right) = \varphi\left(\int_I u_i(s'_i) d\lambda(i)\right)\right\}\right)$$

$$1 = Q\left(\left\{s : \varphi\left(\int_I u_i(s_i) d\lambda(i)\right) = \varphi\left(\int_I u_i(s''_i) d\lambda(i)\right)\right\}\right)$$

Hence, $P(\{s : s \sim s'\}) = Q(\{s : s \sim s''\}) = 1$. By assumption, $\int_I u_i(s'_i)d\lambda(i) \geq \int_I u_i(s''_i)d\lambda(i)$. Hence, by the Restricted Pareto axiom applied to $\delta_{s'}$ and $\delta_{s''}$, $U(s') \geq U(s'')$, so $s' \succsim s''$. Axiom II therefore implies $P \succsim Q$.

We now verify that the strict Pareto axiom holds. To this end, let $\alpha > \beta$ and fix two subsets $c, d \subseteq I$ such that $\lambda(c) = \alpha$ and $\lambda(d) = \beta$. Then $\varphi(\lambda(c)) > \varphi(\lambda(d))$ by the strict monotonicity of φ , hence $s^c \succ s^d$.

As in the proof of sufficiency, we have $P^\alpha(\{s^a : \lambda(a) = \alpha\}) = 1$ and $P^\beta(\{s^a : \lambda(a) = \beta\}) = 1$. Hence $P^\alpha(\{s^a : \lambda(a) = \alpha\}) = 1$. That is, $P^\alpha(\{s^a : s^a \sim s^c\}) = 1$. Similarly, $P^\beta(\{s^a : s^a \sim s^d\}) = 1$. Axiom II, together with the fact that $s^c \succ s^d$, implies $P^\alpha \succ P^\beta$.

Finally, the Anonymity axiom follows immediately from the representation.

A.3 Proof of Theorem 4

Define $\varphi : [0, 1] \rightarrow \mathbb{R}$ as $\varphi(\alpha) = E_{P^\alpha}[U]$ for each $\alpha \in [0, 1]$. By the Strict Pareto axiom φ is strictly increasing.

Lemma 2. *There exists a capacity $\nu : \mathcal{P}(I) \rightarrow [0, 1]$ such that $\nu(I) = 1$ and $U(s^a) = \varphi(\nu(a))$ for every $a \subseteq I$.*

Proof: Given $a \subseteq I$, consider the sets $\{\alpha : P^\alpha \succsim s^a\}$ and $\{\alpha : s^a \succsim P^\alpha\}$. Because $P^1(\{s^I\}) = P^0(\{s^\emptyset\}) = 1$, then the Pareto axiom implies $1 \in \{\alpha : P^\alpha \succsim s^a\}$ and $0 \in \{\alpha : s^a \succsim P^\alpha\}$. By the Continuity axiom the two sets are closed and their union is $[0, 1]$. Hence there exists $\alpha \in [0, 1]$ such that $s^a \sim P^\alpha$. By the Strict Pareto axiom, such α is unique. Hence we can define a set function $\nu : \mathcal{P}(I) \rightarrow [0, 1]$ such that $s^a \sim P^{\nu(a)}$ for every $a \subseteq I$. Whenever $a \subseteq b$ the Pareto axiom implies $P^{\nu(b)} \sim s^b \succsim s^a \sim P^{\nu(a)}$ so $\nu(b) \geq \nu(a)$. In addition, because $s^\emptyset \sim P^0$ then $\nu(\emptyset) = 0$. Thus ν is a capacity. To conclude, notice that $s^a \sim P^{\nu(a)}$ implies $U(s^a) = E_{P^{\nu(a)}}[U] = \varphi(\nu(a))$. ■

Lemma 3. *For every $\alpha \in [0, 1]$, ν satisfies $P^\alpha(\{s^a : \nu(a) = \alpha\}) = 1$.*

Proof: By axiom a the function U is unaffected by changing the outcome of any finite set of agents. Kolmogorov's 0-1 law implies $P^\alpha(\{s^a : U(s^a) = E_{P^\alpha}[U]\}) = 1$. Hence, by the definition of φ and Lemma 2 we obtain $P^\alpha(\{s^a : \varphi(\nu(a)) = \varphi(\alpha)\}) = 1$. Because, φ is strictly increasing, then $\nu(a) = \alpha$ for P^α -almost every profile s^a . ■

The next result constructs an algebra $\mathcal{A} \subseteq \mathcal{P}(I)$ such that ν , when restricted to \mathcal{A} , is additive and strongly non-atomic. Axiom c is not needed for the result.

Theorem 9. *There exists an algebra $\mathcal{A} \subseteq \mathcal{P}(I)$ such that ν when restricted to \mathcal{A} is a finitely additive probability. For every $n \geq 1$, \mathcal{A} contains a partition a_1, \dots, a_n of I such that $\nu(a_1) = \dots = \nu(a_n) = 1/n$.*

Proof: In the proof we will use an auxiliary probability space $(\Omega, \mathcal{F}, \mu)$, where $\Omega = [0, 1]^I$, μ is the product $\mu = \otimes_{i \in I} m$ where m is the Lebesgue measure on $[0, 1)$, and \mathcal{F} the completion with respect to μ of the σ -algebra of Borel subsets of Ω . Given $\omega \in \Omega$, denote by $\omega_i \in [0, 1]$ its i -th coordinate. For each $n \geq 1$ let \mathcal{A}^n be the algebra on $[0, 1)$ generated by the partition $[0, 1/n), \dots, [(n-1)/n, 1)$. Consider the algebra $\mathcal{A} = \bigcup_n \mathcal{A}^n$. Given $A \in \mathcal{A}$ and $\omega \in \Omega$ let $\omega^{-1}(A) = \{i : \omega_i \in A\}$. For each ω , the collection

$$\mathcal{A}_\omega = \left\{ \omega^{-1}(A) : A \in \mathcal{A} \right\}$$

is an algebra of subsets of I . We now show that for μ -almost every ω the realized algebra \mathcal{A}_ω satisfies the properties in the statement of the theorem. Fix $A \in \mathcal{A}$. Given i define the random variable $Z_i : \Omega \rightarrow X_i$ as

$$Z_i(\omega) = \begin{cases} \bar{x}_i & \text{if } \omega_i \in A \\ \underline{x}_i & \text{otherwise} \end{cases}$$

For each i we have

$$\mu(\{\omega : Z_i(\omega) = \bar{x}_i\}) = m(A)$$

By construction, the random variables (Z_i) are independent. They form an i.i.d. process whose distribution is the lottery P^α where $\alpha = m(A)$. Formally, consider the map $Z : \Omega \rightarrow S$ defined as $Z(\omega) = (Z_i(\omega))_{i \in I}$ for all $\omega \in \Omega$. By standard arguments Z is $\mathcal{F} \setminus \Sigma^I$ -measurable and satisfies the change of variable identity

$$P^{m(A)}(E) = \mu(Z^{-1}(E)) \text{ for all } E \in \Sigma^I.$$

Because \mathcal{F} is complete then the same identity extends to all events $E \in \Sigma$.¹⁰ By applying Lemma 3, we then obtain

$$\begin{aligned} 1 &= P^{m(A)}\{s^a : \nu(a) = m(A)\} = \mu\left(Z^{-1}(\{s^a : \nu(a) = m(A)\})\right) \\ &= \mu(\{\omega : \nu(\{i : Z_i(\omega) = \bar{x}_i\}) = m(A)\}) = \mu(\{\omega : \nu(\{i : \omega_i \in A\}) = m(A)\}). \end{aligned}$$

That is, $\mu(\{\omega : \nu(\omega^{-1}(A)) = m(A)\}) = 1$. Therefore,

$$\Omega^* = \bigcap_{n=1}^{\infty} \bigcap_{A \in \mathcal{A}^n} \left\{ \omega : \nu(\omega^{-1}(A)) = m(A) \right\} \quad (9)$$

is a countable intersection of sets that have probability 1 under μ . Hence $\mu(\Omega^*) = 1$.

Fix $\omega \in \Omega^*$. We now show that ν is additive on \mathcal{A}_ω . Let $\omega^{-1}(A_1)$ and $\omega^{-1}(A_2)$ in \mathcal{A}_ω be disjoint. Equivalently, $\omega^{-1}(A_1 \cap A_2) = \emptyset$. Suppose $m(A_1 \cap A_2) > 0$. Then (9) implies $\nu(\omega^{-1}(A_1 \cap A_2)) = m(A_1 \cap A_2) > 0$, i.e. $\nu(\emptyset) > 0$. A contradiction. Thus

¹⁰ Let $E \in \Sigma$. Then E belongs to the completion $\Sigma_{P^{m(A)}}^I$. So, there exists $E_1, E_2 \in \Sigma^I$ such that $E_1 \subseteq E \subseteq E_2$ and $P^{m(A)}(E_2) = P^{m(A)}(E_1)$. Therefore, $Z^{-1}(E_1) \subseteq Z^{-1}(E) \subseteq \phi^{-1}(E_2)$ and $\mu(Z^{-1}(E_2)) = \mu(Z^{-1}(E_1))$. Since $Z^{-1}(E_2), Z^{-1}(E_1) \in \mathcal{F}$ and \mathcal{F} is complete, then $Z^{-1}(E) \in \mathcal{F}$.

$m(A_1 \cap A_2) = 0$. So, $m(A_1 \cup A_2) = m(A_1) + m(A_2)$. Hence, $\nu(\omega^{-1}(A_1) \cup \omega^{-1}(A_2))$ is equal to

$$\nu(\omega^{-1}(A_1 \cup A_2)) = m(A_1 \cup A_2) = m(A_1) + m(A_2) = \nu(\omega^{-1}(A_1)) + \nu(\omega^{-1}(A_2)).$$

Hence ν is additive on \mathcal{A}_ω . Finally, given $n \geq 1$, let A_1, \dots, A_n be the atoms of \mathcal{A}^n . Then $\omega^{-1}(A_1), \dots, \omega^{-1}(A_n)$ is a partition of I and $\nu(\omega^{-1}(A_1)) = \dots = \nu(\omega^{-1}(A_n)) = 1/n$. ■

The next result establishes that the capacity ν is in fact additive.

Lemma 4. *ν is a nonatomic finitely additive probability.*

Proof: Throughout the proof we apply the following implication of axiom c. For every $a, b, c, d \subseteq I$ such that $b \cap d = \emptyset$, if $\nu(a) \leq \nu(b)$ and $\nu(c) < \nu(d)$ then $\nu(a \cup c) < \nu(b \cup d)$. See Fishburn (1970) (Lemma C3, p. 195) for a proof.

Let $b \cap c = \emptyset$. We first show that $\nu(b) + \nu(c) \leq 1$. Suppose not. Notice that the range $\nu(\mathcal{A})$ is a dense subset of $[0, 1]$. Hence, we can find a_1 and a_2 in \mathcal{A} such that $\nu(b) > \nu(a_1)$, $\nu(c) > \nu(a_2)$ and $\nu(a_1) + \nu(a_2) > 1$. Because ν is additive on \mathcal{A} , then $\nu(a_2) > 1 - \nu(a_1) = \nu(a_1^c)$. So, $\nu(b) > \nu(a_1)$ and $\nu(c) > \nu(a_1^c)$. A contradiction. Hence, $\nu(b) + \nu(c) \leq 1$.

Consider the case where $\nu(b) + \nu(c) \in [0, 1)$. Let (k_n) and (l_n) be sequences in \mathbb{N} such that $k_n/n \downarrow \nu(b)$ and $l_n/n \downarrow \nu(c)$. Let N be such that $k_n + l_n \leq n$ for every $n \geq N$. Let $\{a_1^n, \dots, a_n^n\} \subseteq \mathcal{A}$ be a partition of I into n coalitions that have equal weight under ν . For every $n \geq N$ let

$$a_n = a_1^n \cup \dots \cup a_{k_n}^n \text{ and } a'_n = a_{k_n+1}^n \cup \dots \cup a_{k_n+l_n}^n.$$

Then $a_n \cap a'_n = \emptyset$. In addition, $\nu(a_n) \downarrow \nu(b)$ and $\nu(a'_n) \downarrow \nu(c)$ as $n \rightarrow \infty$. So, $\nu(a_n \cup a'_n) > \nu(b \cup c)$. Hence $\nu(a_n) + \nu(a'_n) > \nu(b \cup c)$ for every n . Hence

$$\nu(b) + \nu(c) \geq \nu(b \cup c). \quad (10)$$

In particular, $\nu(b) + \nu(c) = \nu(b \cup c)$ if $\nu(b) + \nu(c) = 0$. By approximating $\nu(b)$ and $\nu(c)$ from below, the same argument can be replicated to show that if $\nu(b) + \nu(c) \in (0, 1]$, then

$$\nu(b) + \nu(c) \leq \nu(b \cup c). \quad (11)$$

In particular $\nu(b) + \nu(c) = \nu(b \cup c)$ if $\nu(b) + \nu(c) = 1$. In the case where $\nu(b) + \nu(c) \in (0, 1)$, then (10) and (11) imply $\nu(b) + \nu(c) = \nu(b \cup c)$. Hence ν is additive. ■

The proof of Theorem 4 can now be concluded as follows. By Lemma 4 there exists a finitely additive probability $\tilde{\lambda}$ such that $s^a \succsim s^b \iff \tilde{\lambda}(a) \geq \tilde{\lambda}(b)$. In particular, the preference \succsim satisfies the Anonymity axiom with respect to $\tilde{\lambda}$. In addition, $U(s^a) = \varphi(\tilde{\lambda}(a))$ for every $a \subseteq I$.

Theorem 4 will then follow from Theorem 2 once λ is substituted by $\tilde{\lambda}$. It remains only to show that $s \mapsto \tilde{\lambda}\{i : s_i = \bar{x}_i\}$ is Σ -measurable. To this end let $f : S \rightarrow S$ map each s to s^a where $a = \{i : s_i = \bar{x}_i\}$. It is immediate to verify that f is $\Sigma^I \setminus \Sigma^I$ -measurable. A routine argument implies it is $\Sigma \setminus \Sigma$ -measurable. For every s we have $\tilde{\lambda}\{i : s_i = \bar{x}_i\} = \phi^{-1}(U(f(s)))$. The proof is concluded by noting that $U \circ f$ is Σ -measurable and ϕ strictly increasing.

A.4 Other Proofs

Proof of Theorem 3: The result follows immediately from Theorem 8. \blacksquare

Proof of Remark 1: Assume \succsim satisfies the Extended Pareto axiom. Fix two profiles s, s' and $\alpha \in [0, 1]$. Let $P = \alpha\delta_s + (1 - \alpha)\delta_{s'}$ and let P° be an independent lottery with the same marginals as P . We have $E_P[U] = \alpha\varphi(\int u_i(s_i) d\lambda(i)) + (1 - \alpha)\varphi(\int u_i(s'_i) d\lambda(i))$ while Theorem 8 implies

$$\begin{aligned} E_{P^\circ}[U] &= \varphi\left(\int (\alpha u_i(s_i) + (1 - \alpha)u_i(s'_i)) d\lambda(i)\right) \\ &= \varphi\left(\alpha \int u_i(s_i) d\lambda(i) + (1 - \alpha) \int u_i(s'_i) d\lambda(i)\right) \end{aligned}$$

Because s and s' can be chosen such that $\int u_i(s_i) d\lambda(i)$ and $\int u_i(s'_i) d\lambda(i)$ correspond to any two points in the domain of φ , it follows that φ is affine. \blacksquare

Lemma 5. *Let φ be continuously differentiable. Then there exists $K > 0$ such that for every $\epsilon > 0$, $n \in \mathbb{N}$ and independent lottery P ,*

$$P \left\{ s : \left| \varphi\left(\frac{1}{n} \sum_{i \in I_n} u_i(s_i)\right) - \varphi\left(\frac{1}{n} \sum_{i \in I_n} E_{P_i}[u_i]\right) \right| < \epsilon \right\} > 1 - 2e^{-2n(\frac{\epsilon}{K})^2}.$$

Proof: Since φ is continuously differentiable then it is K -Lipschitz where $K = \max \varphi'$. The result now follows by applying McDiarmid concentration inequality McDiarmid (1998) (Theorem 3.1) applied to the function $\frac{1}{n} \sum_{i \in I_n} u_i$ and the fact that φ is K -Lipschitz. \blacksquare

Proof of Theorem 5: Fix a sequence $\delta_n \downarrow 0$ and set $M = \max \varphi$, $L(n) = 1 - 2e^{-2n(\frac{\delta_n}{K})^2}$ and $A(n, P) = \varphi\left(\frac{1}{n} \sum_{i \in I_n} E_{P_i}[u_i]\right)$. By Lemma 5,

$$|E_P[U_n] - A(n, P)| \leq E_P[|U_n - A(n, P)|] \leq (1 - L(n))M + L(n)\delta_n$$

Let $\epsilon_n = 2((1 - L(n))M + L(n)\delta_n)$. Then $\epsilon_n \downarrow 0$. If $E_{P_i}[U_n] \geq E_{Q_i}[U_n]$ for every $i \in I_n$ then $A(n, P) \geq A(n, Q)$, hence

$$E_P[U_n] \geq A(n, P) - \frac{\epsilon_n}{2} \geq A(n, Q) - \frac{\epsilon_n}{2} \geq E_Q[U_n] - \epsilon_n$$

\blacksquare