Aggregate Risk and the Pareto Principle*

Nabil I. Al-Najjar† Luciano Pomatto‡

May 2020

Abstract

In the evaluation of public policies, a crucial distinction is between plans that involve purely idiosyncratic risk and those that generate correlated, or aggregate, 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, it is apt for studying large populations, and is characterized by two simple axioms of preferences aggregation.

Keywords: Harsanyi's theorem; aggregate risk; preference aggregation.

JEL codes: D71, D81.

1 Introduction

Correlated risks, such those associated with catastrophic pandemics or systemic failures of a financial system, have been a major concern for policy makers and academics in recent years. Public policy questions invariably involve a mix of idiosyncratic and correlated risks. However, the policy responses and the tenor of the public debates surrounding the two types of risk are very different because correlation leads to aggregate risk that does not disappear even in a large population.

The collapse of major financial institutions in 2008 provides a vivid illustration of the impact of correlated "systemic risks" that can compromise an entire financial system. This crisis raised the concern that financial institutions may fail to adequately incorporate systemic risks in their decisions. See Lorenzoni (2008), Adrian and Brunnermeier (2016), Acharya, Pedersen, Philippon, and Richardson (2017) for recent models and literature

*We are grateful to the co-editor, Pietro Ortoleva, and to three anonymous referees for their helpful comments. We thank Larry Blume, Amrita Dhillon, Federico Echenique, Atsushi Kajii, Edi Karni, and Kota Saito.

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

‡ *corresponding author,* Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, 91106. Link to Research Page.

reviews. As Adrian and Brunnermeier (2016) argue, individual institutions tend to ignore the impact of correlation, leading to a spiral of precautionary liquidity hoarding. Liquidity accumulation by shedding assets at fire-sale prices may well be prudent for individual institutions, yet can be disastrous for the financial system as a whole. The massive policy interventions triggered by the crisis was predicated on the inability, or unwillingness, of banks and other financial institutions to internalize economy-wide correlated risks. In addition to immediate injections of funds into the system, legislative measures were introduced to constraint financial institutions in apparent recognition of these institutions' inability to handle correlated risk.[1]

The distinctive role played by correlated risks colors public debates on policy responses well beyond the context of financial markets. A key example is the question of how to properly quantify the uncertainty generated by public hazards. It is by now well understood that the expected number of fatalities is an inappropriate measure of risk. In his seminal study, Keeney (1980) proposed alternative criteria for assessing social risks, focusing on the interplay between equity and catastrophe aversion in the context of large-scale accidents (see Bernard, Treich, and Rheinberger (2017) for a recent contribution). Aggregate risk also plays an important role in many other areas, including the Precautionary Principle for dealing with catastrophic risk (Sunstein 2005) and treatment choice under uncertainty (Manski and Tetenov 2007).

To model how correlated and idiosyncratic risks might differ in impact on collective policy decisions, we start by revisiting Harsanyi's (1955) celebrated utilitarian welfare criterion. Under this criterion, social prospects are evaluated according to the expectation of the additive aggregator of the form:

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

where s is a profile of social outcomes; its ith coordinate, s_i, is individual i's outcome; and u_i is this individual's utility function. As common in economic models, utilities are functions of the individual outcomes only. For example, in the context of systemic risk in a financial market, each institution is directly concerned about its profits and losses, rather than the health of the financial system as a whole.

It is clear that additive aggregators like (1) cannot depend on the degree of correlation of the variables $(s_i)_{i \in I}$ and thus fail to distinguish between idiosyncratic and correlated risks. Our goal is to extend this aggregator in a way that can capture sensitivity to correlation.

Consider, for concreteness, two risky policies, labelled A and B. Assume each option will affect a large population and, for each individual, will result in either a good or bad

[1] The Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010. See also the U.S. Congress Committee on Science and Technology hearings on "The Risks of Financial Modeling: VAR and the Economic Meltdown," September 10, 2009.

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 independently 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).

As the earlier discussion suggests, a policy maker may draw a distinction among the two based on equity concerns, social considerations, or economic motives. For example, one may argue that A is more equitable since it leads to a perfectly equal distribution of utility ex post, while B will split society in two subgroups enjoying very different outcomes. Taking a different perspective, A exposes society to the (possibly catastrophic) risk of a uniformly bad outcome while B is very unlikely to lead to such catastrophic outcome. In fact, under policy B, society, as a whole, will not be exposed to any risk represented by an aggregate statistic that is a function of the distribution of outcomes in the population.

Many authors noted the fact that public opinion reacts differently to catastrophic risks compared to more familiar ones like car accidents and house fires.[2] This concern about correlation can reflect different motives that can be difficult to disentangle. A person might decide to vote against a policy that involves aggregate risk partly because of the material long-run implications of a catastrophic outcome, and partly because she sees society as being more than the well-being of its members. She might consider living in a prosperous society an intrinsically valuable goal, above and beyond the consumption or wealth of its members. In this paper we take a general approach that allows us to abstract away from specific normative considerations that favor or oppose aggregate risk. Our goal, instead, is to put forward an approach for incorporating sensitivity to correlation based on few transparent principles.

1.1 Outline of the Model

We consider a standard economic environment given by a large population, modeled as a non-atomic space of agents (I, λ). Policies are identified with lotteries over profiles $s \in \prod_{i \in I} X_i$, where X_i is the set of possible outcomes of agent i. Each individual is endowed with a preference relation $\succsim_i$ over lotteries on X_i. These preferences must be aggregated into a social preference relation $\succsim$. 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

[2]See, for example, Sunstein (2005). Robson (1996) provides an evolutionary justification for why Nature might have designed individuals with differing attitudes towards aggregate versus idiosyncratic risks.

for aggregate risk. In the environments we study, the axiom forces a policy maker to deem as equally desirable two policies that induce the same individual risks, regardless of the correlation they induce. For this reason, 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 a lottery P *independent* if the individual outcomes $(s_i)_{i \in I}$ are independent random variables under P. An independent lottery therefore 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 the 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. 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 allocations s^a and s^b.

We show that 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

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

where u_i is individual i's von-Neumann Morgentstern 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 of 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.

To illustrate these features of the representation, 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 correlated risks.

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.

1.2 Related Literature

An important reason to study correlation among individual risks is 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 $U(s_i) = \phi^{-1}\left(\int_I \phi(u_i(s_i))d\lambda(i)\right)$, 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.

While the aggregator we study in this paper can capture some forms of inequity aversion, it remains close to utilitarianism, and hence there are forms of inequity aversion that are outside the scope of our model. For instance, our approach cannot capture ex-post inequity over riskless allocations.

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. In this paper we show that in the context of *large* populations (a natural setup for studying public risk) it is possible to capture sensitivity to correlated risk while retaining most of the features of utilitarian aggregators.

Representation similar to the one we study in this paper have appeared in the literature on time preferences. Consider a decision maker ranking lotteries over consumption streams $(c_t)_{t=1}^{\infty}$. As shown by Dillenberger, Gottlieb, and Ortoleva (2018), many common functional forms are subsumed by a representation in the spirit of Kihlstrom and Mirman (1974), where lotteries over consumption streams are evaluated by the expectation of a cardinal utility of the form $\phi(\sum_t \beta^t u(c_t))$. The function ϕ allows to capture correlation aversion and seeking (e.g. Epstein and Zin 1989, Bansal and Yaron 2004, Andersen, Harrison, Lau, and Rutström 2018), among other phenomena.

Dillenberger, Gottlieb, and Ortoleva (2018) provide a foundation for this representation, by combining the von Neumann-Morgenstern axioms over lotteries with the axioms of discounted utility over consumption streams. No analogous two-stage axiomatization is readily available in our framework. This is because Harsanyi's axiomatization, on which the utilitarian representation $\int_I u_i(s_i)d\lambda(i)$ is founded, leaves no room for disentangling the social preference over riskless allocations from the social preference over lotteries. One of our contributions is to show that for large populations Harsanyi's axiomatization is unduly restrictive, and that it is possible to provide a foundation for utilitarianism that allows for correlation aversion or seeking.

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 countable 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 set X_i could represent consumption bundles or health consequences, but we do not impose any formal structure over X_i.

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. An *individual lottery* for agent i is a (σ-additive) probability measure P_i defined on (X_i, Σ_i). We denote by $\Delta(X_i)$ the set of all such lotteries. We denote by $\Sigma^I = \otimes_{i\in I}\Sigma_i$ the product σ-algebra, i.e. the σ-algebra generated by the projection mappings $\pi_i \colon S \to X_i$ defined for each $i \in I$ as $\pi_i(s) = s_i$. 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 $P \in \Pi(S)$ the 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 often 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 $E \in \Sigma_i$. For every profile of probability measures

$(P_i)_{i\in I} \in \prod_{i\in I} \Delta(X_i)$, there exists a product measure $Q \in \Pi(S)$ with marginals $Q_i = P_i$ for all i (Halmos 1974, §38, Theorem B). 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

Each agent i is endowed with a binary preference relation $\succsim_i$ over the set $\Delta(X_i)$ of individual lotteries that admits the expected utility representation

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

where the von Neumann-Morgenstern utility function $u_i \colon X_i \to \mathbb{R}$ 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 their outcome. Each 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 $\overline{x}_i$ and $\underline{x}_i$ in X_i such that $u_i(\overline{x}_i) = 1$ and $u_i(\underline{x}_i) = 0$.

We consider a social preference relation $\succsim$ over lotteries represented as

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

where $U \colon S \to \mathbb{R}$ 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 Remarks on Individual and Social Preferences

Individual and social preferences are defined over different domains. The preferences of each agent i are defined over the corresponding set $\Delta(X_i)$ of individual lotteries. We interpret the relation $\succsim_i$ as adequate for describing individual decision problems such as choices over consumption bundles, assets, or insurance policies.

In contrast, the social preference relation $\succsim$ is defined over the domain $\Delta(S)$ of policies. The relation represents the preferences of an anonymous individual, when taking an impartial perspective and evaluating policies according to their consequences on the whole population. The relation $\succsim$ can therefore be seen as guiding those choices that are characterized by a prominent impersonal component, such as casting a vote in a referendum. It may be difficult for an agent to formulate a well defined ranking over policies, and from this perspective one can interpret axioms relating personal and social preferences as normative principles, aimed at guiding a decision maker in formulating their social preferences.

The distinction we adopt in this paper between individual and social preferences, as well as their interpretation, is connected to well-known ideas in welfare economics. Arrow (2012, p.18) makes a distinction between "the ordering of social states according to the direct consumption of the individual and the ordering when the individual adds his general standards of equity (or perhaps his standards of pecuniary emulation)." Here, Arrow refers to the first ordering as reflecting the individual *tastes*, and the latter as reflecting their *values*.

A similar distinction plays an important role in Harsanyi (1955). In his framework, the relation $\succsim$ represents the *ethical* preferences of an individual, which must "express what this individual prefers (or, rather, would prefer) on the basis of impersonal social considerations alone" while the *subjective* preferences, represented by $\succsim_i$, must express "what he actually prefers."

In Harsanyi's original framework individual and social preferences are defined over the same domain, which is the collection of all lotteries over an abstract set of outcomes. In contrast, in this paper we focus on a set S of allocations that has a product structure, and adopt different domains for the two preference relations. This modeling choice limits the generality of our framework, but has several advantages.

First, given our focus on modeling idiosyncratic and correlated risk, it is important for our analysis that the set of allocations has a product structure. A second advantage comes from the fact that individual preferences over private outcomes are in principle observable, and are close to the type of primitives routinely studied in economic models. Our approach is therefore in the spirit of the literature on preferences aggregation in "economic environments," following Kalai, Muller, and Satterthwaite (1979), Border (1983), and Bordes and Le Breton (1989). With this body of work we share the motivation of studying aggregation of preferences that satisfy the same structural conditions typically found in economic applications. Finally, adopting different domains for individual and social preferences makes it possible to capture Arrow's and Harsanyi's distinctions between tastes and values, and between ethical and personal preferences.

2.3 Harsanyi's Theorem

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's Theorem can be stated as follows:[3]

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*

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

[3]Zhou (1997) generalized Harsanyi's Theorem to infinite populations. See also Remark 1 below.

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 correlation. In order to accommodate this basic disposition towards social risk, in the next sections we will 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 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 probability space (I, λ). Given a group $a \subseteq I$, $\lambda(a)$ represents the fraction of agents that belong to that group. In Section C in the Appendix we show how our main representation can be applied to finite populations.

To avoid well known measurability issues that arise when dealing with an uncountable family of independent random variables,[4] I is assumed to be countable and λ is a finitely additive probability measure defined on the collection of all subsets of I. We also assume that the map $s \mapsto \lambda(\{i : s_i = \overline{x}_i\})$ is Σ-measurable.[5]

3.1 Restricted Pareto Axiom

It is a familiar fact from everyday life that difficult choices might lead to an internal conflict between a person's taste and their values. This idea has long been recognized in welfare economics. In distinguishing between individual and moral preferences, Harsanyi (1986, p. 51) noted that "the personal and the moral preferences of a given individual may quite dissimilar," to the point that they "may rank two social situations A and B in the opposite way." The same natural difficulty was recognized by Arrow (2012), who wrote "In general, there will, then, be a difference between the ordering of social states according to the direct consumption of the individual and the ordering when the individual adds his general standards of equity." Common to these considerations is the idea that individual preferences are only one of the factor contributing to a person's ranking over policies.

Under the Pareto principle, the conflict behind tastes and values is mediated according to the following doctrine: the ranking of policies must be a function solely of the individual

[4]See Judd (1985), and Al-Najjar (2008) among many others.

[5]The 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.

preferences of the agents. Our starting point is the observation that this reduction to individual tastes is not equally compelling in all decision problems, particularly when choosing among policies that involve risk.

Consider, as an example, the question of whether or not society should prohibit risky behavior such as smoking, drug consumption, or lack of physical exercise. We can imagine a person's decision of whether or not to to engage in risky behavior as a choice between a risky lottery and a safe and riskless outcome. A person is guided in this choice by their preferences over consumption and health prospects, as well as their degree of risk aversion. In this example, the Pareto principle expresses a principle of consumer sovereignty: if individuals are entitled to their tastes, then the optimal policy is to allow each agent to pick their favorite option. Notice that in this example the policy leads to health risks that are independent across agents.

Consider now a different situation, where the question is whether to adopt a policy that guarantees every agent a safe outcome, or a second policy that exposes the population to aggregate risk. As in the previous example, appealing to the Pareto principles implies ranking the two policies solely on the basis of individual tastes. But in this case the argument, while logically coherent, appears less compelling since individual preferences do not encode an intrinsic attitude towards aggregate risk. Any such attitude must be derived from other principles.

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 two policies involve risk only at the individual level or also at the societal level. Our first axiom avoids this strong conclusion by applying the Pareto principle only to the ranking of independent lotteries.

Axiom 1 (Restricted Pareto)**.** *For all independent lotteries P and Q,*

$$\text{if } P \succsim_i Q \text{ for every } i, \text{ then } P \succsim Q.$$

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

For the next axiom, we denote by P^α the independent lottery where each agent i receives her best outcome $\overline{x}_i$ with probability $\alpha \in [0,1]$ 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 probability of the 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 $P^\alpha \succ P^\beta$.*

3.2 Anonimity

The next axiom requires groups that represent equal fractions of the population to be treated similarly. Given a set $a \subseteq I$ of agents, let s^a be the profile defined as $s_i^a = \overline{x}_i$ if $i \in a$ and $s_i^a = \underline{x}_i$ if $i \notin a$. The allocation s^a assigns to every agent their best outcome if they belong to group a and their worst outcome otherwise. We refer to each s^a as an *extreme* allocation.

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

The assumption can be interpreted as an expression of impartiality. This point of view is analogous to that of Karni (1998). Consider a standard revealed preference approach, according to which the conflict between the desires of different agents should be resolved based solely upon agents' preferences over outcomes. From this perspective, the choice of favoring one agent's most preferred outcome over that of another agent seems difficult to justify. The axiom captures the same principle in the context of a large population. Given two group of agents a and b that represent the same fraction $\lambda(a) = \lambda(b)$ of the populations, the allocations s^a and s^b appear equally compelling given the information available to the social planner. Hence, they should be considered equally desirable.

Axiom 3 admits an additional interpretation as an assumption concerning interpersonal comparisons of preferences. As is well known, comparisons between different agents' utilities are unavoidable in Harsanyi's theory: the weights (λ_i) that appear in the aggregator $\sum_i \lambda_i u_i$ measure the rates by which the social planner is willing to trade off utilities across agents. A contentious aspect of Harsanyi's Theorem is that it imposes no restrictions over the weights, and hence provides no clear indication to the final user of the theory (see, for instance, the discussion in Dhillon and Mertens (1999)). From this perspective, Axiom 3 states an elementary assumption: best and worst outcomes are comparable across agents.

We remark that this assumption, while easy to interpret, is not innocuous. It is a better approximation when the sets of individual outcomes (X_i) are identical across agents, as when each X_i represents income levels or health conditions (ranging from "a life-threatening health condition" to "being in perfect health"). In Section 5.2 we consider a more general framework.

Anonimity conditions in the spirit of Axiom 3 have a long history in the literature on preferences aggregation. An early example is May's (1952) characterization of majority rule. 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), Segal (2000), and Piacquadio (2017), among others.

3.3 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 $\varphi\colon [0,1] \to \mathbb{R}$ such that*

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

When confined to the ranking of deterministic profiles, the ordinal ranking described by U is unaffected by the strictly increasing transformation φ. Hence, in the absence or 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 Section 4.

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 Harsanyi's Extended Pareto axiom.

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

4.1 Conditionally i.i.d. Lotteries

We 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).[6] The following corollary of Theorem 3 gives a convenient formula for evaluating the aggregative utility of conditionally i.i.d. lotteries in homogeneous populations:

[6] By de Finetti's 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.

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

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

4.2 Comparative Attitude to Aggregate Risk

We now illustrate how concavity, or convexity, of the transformation φ characterize aversion to, or preference for, correlated risk. To this end, we say that a lottery P *displays no aggregate risk* if the average utility

$$\int_I u_i(s_i)\, d\lambda(i)$$

is constant P-almost surely. By Theorem 3, this class contains all lotteries that are independent.

Given two social preferences $\succsim_1$ and $\succsim_2$, we say that $\succsim_1$ is *more averse than* $\succsim_2$ *to aggregate risk* if for every lottery P and Q such that Q displays no aggregate risk,

$$P \succsim_1 Q \implies P \succsim_2 Q.$$

That is, whenever the first relation ranks a lottery P above a second lottery Q that displays no aggregate risk, so does the second relation $\succsim_2$. This comparative notion adapts Yaari's (1969) definition of comparative risk aversion to our framework.

Attitudes towards social risk admit a straightforward characterization in terms of concavity or convexity of the transformation φ.

Remark 2. *Let $\succsim_1$ and $\succsim_2$ be two social preference relations that admit the representation* (3) *with respect to the continuous transformations φ_1 and φ_2, respectively. Then, $\succsim_1$ is more averse to aggregate risk than $\succsim_2$ if and only if $\varphi_1 = f \circ \varphi_2$, where f is strictly increasing and concave.*

It is natural to deem *neutral* towards aggregate risk a social preference that ranks equally any two lotteries that share the same marginal distributions for every agent. Under this choice of benchmark, Proposition 2 shows that a social preference $\succsim$ is *averse* to aggregate risk if and only if the corresponding transformation φ is concave.

5 Discussion and Extensions

5.1 Comparison with other Generalizations of Utilitarianism

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

$$U(s) = \int_I \phi(u_i(s_i)) d\lambda(i) \tag{5}$$

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

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

The representation displays 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

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

where $\theta^\mu \in \Delta(X)$ is the marginal induced by P^μ. So, under the EEDEU criterion, a social preference can exhibit aversion to correlation 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 correlation. Decreasing the impact of aggregate uncertainty may require a "diversification" across individuals that exacerbates ex-post welfare differences.

The two classes of social preferences describe separate generalizations of Harsanyi's theorem: a social preference relation that satisfies axioms 1-3 admits an EEDEU representation if and only if it corresponds to a utilitarian aggregator. We establish this fact in Proposition 1 in the Appendix.[7]

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's theorem, in contrast, the weight $\lambda(a)$ is derived from the social

[7] There are other important differences between the two papers. The main result in Fleurbaey (2010) characterizes an abstract aggregator which is more general than (6) and that does not necessarily reduce to a utilitarian aggregator in the absence of risk. Axiomatizations of (6) appear in Fleurbaey and Zuber (2013) and McCarthy (2015), under the assumption that the social preference is separable across subpopulations, a condition that is conceptually distinct from our axioms 1-3.

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's theorem has been the subject of considerable scrutiny.

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. First, we require 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 more substantive: it expresses the idea that 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 $\varphi\colon [0,1] \to \mathbb{R}$ 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.[8]

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 important 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.

[8] Formally, this follows from the fact that the only assumptions imposed on λ are that it is non-atomic and satisfies Σ-measurability of the map $s \mapsto \tilde{\lambda}(\{i : s_i = \overline{x}_i\})$. Both assumptions are satisfied by a measure $\tilde{\lambda}$ obtained through Theorem 4.

5.3 Non-Expected Utility

Our main result can be extended to non-expected utility preferences. Our aggregation theorem continues to hold under very minimal assumption on the social planner's preference over lotteries. We consider a policy maker whose preferences over lotteries are represented by a binary relation $\succsim^*$ over $\Delta(S)$, which needs not satisfy the von-Neumann Morgenstern axioms. As usual, we denote by $\succ^*$ and $\sim^*$ the anti-symmetric and symmetric parts of $\succsim^*$, respectively.

We impose two basic axioms on $\succsim^*$. For the next condition, we denote by $\succsim^*_{|S}$ the restriction of $\succsim^*$ over the set S of deterministic allocations.

Axiom I. *The preference $\succsim^*_{|S}$ is complete, transitive and S contains a countable $\succsim^*_{|S}$-order-dense subset.*[9] *For every $s \in S$, the sets $\{t \in S : t \succsim^* s\}$ and $\{t \in S : s \succsim^* t\}$ are Σ-measurable.*

The axiom is equivalent to the existence of a social utility function $U\colon S \to \mathbb{R}$ 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*

$$P\left(\{s \in S : s \sim^* s'\}\right) = Q\left(\{s \in S : s \sim^* s''\}\right) = 1 \tag{7}$$

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

The two axioms are compatible with several models of decision under risk. 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 5. *Let $\succsim^*$ be a binary preference relation on $\Delta(S)$ that satisfies axioms I and II. Let $U : S \to \mathbb{R}$ 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 $\varphi : [0,1] \to \mathbb{R}$ such that*

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

The result shows that 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.

[9] That is, there exists a countable set $S_0 \subset S$ such that for all $s, s' \in S$, if $s \succ^* s'$ there exists $s'' \in S_0$ such that $s \succsim^* s'' \succsim^* s'$.

Acknowledgements

We are grateful to the co-editor, Pietro Ortoleva, and to three anonymous referees for their helpful comments. In addition, we thank Larry Blume, Amrita Dhillon, Federico Echenique, Atsushi Kajii, Edi Karni, and Kota Saito.

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. *$s \mapsto \lambda(\{i : s_i = \overline{x}_i\})$, $s \in S$, 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.*[10]

The following hold: (1) $\Longrightarrow$ (2) and (1) $\Longleftrightarrow$ (3).

Proof: (1) implies (2). Fix a Borel set $A \subseteq [0,1]$ and, for each i, the function $\phi_i : X_i \to \mathbb{R}$ 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 $\phi : S \to S$ defined as $\phi(s) = (\phi_i(s_i))$ is then $\Sigma^I \backslash \Sigma^I$-measurable. We claim that ϕ is also $\Sigma \backslash \Sigma$-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 $E \in \Sigma^I_Q$, 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 $\phi^{-1}(E) \in \Sigma^I_P$. So, $\phi^{-1}(E) \in \Sigma$. It follows that the composition $s \mapsto \lambda(\{i : \phi(s)_i = \overline{x}_i\})$ 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 $A_1, ..., A_n$ of $[0,1]$ and all $\alpha_1, ..., \alpha_n$ 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 $f_n : [0,1] \to [0,1]$ be a function with finite range such that $|f(t) - t| \leq 1/n$ for every $t \in [0,1]$. Then

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

$s \mapsto \int_I f_n(u_i(s_i))d\lambda(i)$ is Σ-measurable. For every s,

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

Hence $s \mapsto \int_I \left(u_i(s_i)\right) d\lambda(i)$ 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 $f_i : \{0,1\} \to X_i$ be defined as $f_i\left(1\right) = \overline{x}_i$ and $f_i\left(0\right) = \underline{x}_i$. Then f_i is measurable. Let $f : \{0,1\}^I \to S$ be defined as $f\left(\xi\right) = \left(f_i\left(\xi_i\right)\right)_{i\in I}$ for all $\xi \in \{0,1\}^I$. By standard arguments f is $\mathcal{B}(\{0,1\}^I)\backslash\Sigma^I$-measurable. By replicating the argument applied in the first part of the proof we obtain that f is also $\mathcal{B}_{um}(\{0,1\}^I)\backslash\Sigma$-measurable. Let $l : S \to \mathbb{R}$ be defined as $l(s) = \lambda(\{i : s_i = \overline{x}_i\})$. The composition $l \circ f$ is $\mathcal{B}_{um}(\{0,1\}^I)$-measurable. For all $\xi \in \{0,1\}^I$, it satisfies

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

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

By Lemma 1, any λ such that $s \mapsto \lambda(\{i : s_i = \overline{x}_i\})$ 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.[11] The result follows directly from Theorem 538S in Fremlin (2008) and lemma 1.

Theorem 6. *Under Axiom P there exists a nonatomic finitely additive probability λ defined on $\mathcal{P}(I)$ such that $s \mapsto \lambda(\{i : s_i = \overline{x}_i\})$, $s \in S$, 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

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

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 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 nonatomic population models.

Theorem 7. *Let λ be a nonatomic finitely additive probability defined on $\mathcal{P}(I)$ such that $s \mapsto \lambda(\{i : s_i = \overline{x}_i\})$ is Σ-measurable. Then, for every independent lottery P,*

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

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 (2020). ■

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 5

Since Theorem 2 is a special case of Theorem 5 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 5.3. No confusion should arise by denoting $\succsim^*$ simply by $\succsim$ for the rest of this proof. We fix a bounded, Σ-measurable function U such that $s \succsim s'$ iff $U(s) \geq U(s')$ for all $s, s' \in S$.

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 $\varphi\colon [0,1] \to \mathbb{R}$ such that $U(s^a) = \varphi(\lambda(a))$ for every $a \subseteq I$. We now show φ is strictly increasing. Let $\alpha \in (0,1)$. By Theorem 7

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

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

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

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

By Theorem 7 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 $\int_I u_i(s_i) d\lambda(i) = \lambda(\{i : s_i = \overline{x}_i\})$ 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 any such $\tilde{a} \subseteq I$ 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 7 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 we have $\int_I u_i(s'_i)d\lambda(i) \geq \int_I u_i(s''_i)d\lambda(i)$. 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. 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 $P^\alpha(\{s^a : \lambda(a) = \lambda(c)\}) = 1$. 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

We first prove sufficiency of the axioms. Let $\succsim$ satisfy axioms a-c, 1 and 2. Define $\varphi\colon [0,1] \to \mathbb{R}$ 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 $\nu\colon \mathcal{P}(I) \to [0,1]$ 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 $\{\alpha : s^a \succsim P^\alpha\}$. By the Pareto Axiom $s^I \succ s^a \succ s^\emptyset$. Because $P^1(\{s^I\}) = P^0(\{s^\emptyset\}) = 1$, then $1 \in \{\alpha : P^\alpha \succsim s^a\}$ 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 $s^a \backsim P^\alpha$. By the Strict Pareto axiom, such α is unique. Hence we can define a set function $\nu\colon \mathcal{P}(I) \to [0,1]$ such that $s^a \backsim P^{\nu(a)}$ 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 $s^a \backsim P^{\nu(a)}$ 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 U is P^α-almost surely constant. Hence $P^\alpha(\{s^a : U(s^a) = E_{P^\alpha}[U]\}) = 1$. 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 8. *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 $a_1, ..., a_n$ of I such that $\nu(a_1) = ... = \nu(a_n) = 1/n$.*

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 $[0,1/n), ..., [(n-1)/n, 1)$. 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 $\mathcal{A}_\omega$ satisfies the properties in the statement of the theorem. Fix $A \in \mathcal{A}$. Given i define the random variable $Z_i : \Omega \to X_i$ as

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

For each i we have

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

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 $Z : \Omega \to S$ defined as $Z(\omega) = (Z_i(\omega))_{i \in I}$ for all $\omega \in \Omega$. By standard arguments Z is $\mathcal{F} \backslash \Sigma^I$-measurable and satisfies the change of variable identity

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

Because $\mathcal{F}$ is complete then the same identity extends to all events $E \in \Sigma$.[12] By applying Lemma 3, we then obtain

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

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

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

[12] Let $E \in \Sigma$. Then E belongs to the completion $\Sigma^I_{P^{m(A)}}$. 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, $Z^{-1}(E_1) \subseteq Z^{-1}(E) \subseteq Z^{-1}(E_2)$ and $\mu\left(Z^{-1}(E_2)\right) = \mu\left(Z^{-1}(E_1)\right)$. Since $Z^{-1}(E_2), Z^{-1}(E_1) \in \mathcal{F}$ and $\mathcal{F}$ is complete, then $Z^{-1}(E) \in \mathcal{F}$.

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 $\mathcal{A}_\omega$. Let $\omega^{-1}(A_1)$ and $\omega^{-1}(A_2)$ in $\mathcal{A}_\omega$ 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 $m(A_1 \cap A_2) = 0$. 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

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

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

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 $\nu(c) > \nu(a_1^c)$. Thus $\nu(b \cup c) > \nu(a_1 \cup a_1^c) = \nu(I)$. 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 $\{a_1^n, ..., a_n^n\} \subseteq \mathcal{A}$ be a partition of I into n coalitions that have equal weight under ν. For every $n \geq N$ let

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

Then $a_n \cap a_n' = \emptyset$. In addition, $\nu(a_n) \downarrow \nu(b)$ and $\nu(a_n') \downarrow \nu(c)$ as $n \to \infty$. 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

$$\nu(b) + \nu(c) \geq \nu(b \cup c). \tag{10}$$

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

$$\nu(b) + \nu(c) \leq \nu(b \cup c). \tag{11}$$

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. ■

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 $a \subseteq I$.

We next show that $s \mapsto \tilde{\lambda}(\{i : s_i = \overline{x}_i\})$ is Σ-measurable. Let $f \colon S \to S$ map each s to s^a where $a = \{i : s_i = \overline{x}_i\}$. It is immediate to verify that f is $\Sigma^I \backslash \Sigma^I$-measurable. A routine argument implies it is $\Sigma\backslash\Sigma$-measurable. For every s we have $\tilde{\lambda}(\{i : s_i = \overline{x}_i\}) = \varphi^{-1}(U(f(s)))$. The proof of sufficiency is concluded by noting that $U \circ f$ is Σ-measurable and φ strictly increasing. The proof of sufficiency is now concluded by applying Theorem 2 after substituting λ with $\tilde{\lambda}$.

We now turn to proving necessity of the axioms. Given the representation of Theorem 4, that $\succsim$ satisfies axiom a follows from the fact that $\tilde{\lambda}$ is non-atomic. Axiom b follows from Theorem 3. It implies that $P^\alpha \succ s$ if and only if $\varphi(\alpha) > \varphi(\int_I u_i(s_i) d\lambda(i))$, which holds if and only if $\alpha > \int_I u_i(s_i) d\lambda(i)$. Hence, the set $\{\alpha : P^\alpha \succ s\}$ is open. The same argument implies $\{\alpha : s \succ P^\alpha\}$ is open as well. Axiom c is a direct implication of the additivity of $\tilde{\lambda}$. That axioms 1 and 2 are satisfied by the representation follows from Theorem 2.

A.4 Other Proofs

Proof of Theorem 3: The result follows immediately from Theorem 7. ■

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 $E_P[U] = \alpha\varphi\left(\int_I u_i(s_i)\, d\lambda(i)\right) + (1-\alpha)\varphi\left(\int_I u_i(s'_i)\, d\lambda(i)\right)$ while Theorem 7 implies

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

Because α, s and s' are arbitrary, it follows that φ is affine. ■

Proof of Remark 2: The proof is standard. Let $\varphi_1 = f \circ \varphi_2$. The function f is strictly increasing. Assume first f is concave, and consider two lotteries P and Q, such that $\int_I u_i(s_i) d\lambda(i) = t \in [0,1]$, Q-almost surely. If $P \succsim_1 Q$ then

$$f(\varphi_2(t)) \leq \int (f \circ \varphi_2)\left(\int_I u_i(s_i)\right) dP(s) \leq f\left(\int \varphi_2\left(\int_I u_i(s_i)\right) dP(s)\right)$$

where the last inequality follows from the concavity of f. Since f is strictly increasing, we obtain that $P \succsim_2 Q$. Hence $\succsim_1$ is more averse to aggregate risk than $\succsim_2$.

We now turn to the second part of the proof. Let $\succsim_1$ be more averse to aggregate risk than $\succsim_2$.

As a first step, we show that the map $s \mapsto \int_I u_i(s_i)d\lambda(i)$, $s \in S$, has range equal to $[0,1]$. This follows by Theorem 3: for every $\alpha \in [0,1]$ we have that

$$\int_I u_i(s_i)d\lambda(i) = \int_I E_{P_i^\alpha}[u_i]d\lambda(i) = \alpha$$

for P^α-almost every $s \in S$.

Now let $t, t' \in [0,1]$. Without loss of generality assume $t \neq t'$. Fix $s, s' \in S$ such that $\int_I u_i(s_i)d\lambda(i) = t$ and $\int_I u_i(s_i')d\lambda(i) = t'$. Let $\alpha \in [0,1]$ and define $P = \alpha\delta_s + (1-\alpha)\delta_{s'}$. Since φ_2 is continuous, there exists $t'' \in [0,1]$ such that

$$\alpha\varphi_2(t) + (1-\alpha)\varphi_2(t') = \varphi_2(t''). \tag{12}$$

Let $s'' \in S$ satisfy $\int_I u_i(s_i'')d\lambda(i) = t''$. Since $t \neq t'$ then $\varphi(t'') > \varphi(0)$, and hence

$$P \sim_2 s'' \succ_2 s^\emptyset$$

where, as usual, $s^\emptyset$ is the profile where each agent obtains utility 0. Given $\epsilon \in (0,1)$ define $R = (1-\epsilon)P + \epsilon\delta_{s^0}$. Then $\delta_{s''} \succ_2 R$. Since $\delta_{s''}$ does not display aggregate risk, it must be that $\delta_{s''} \succsim_1 R$ (if not then $R \succ_1 \delta_{s''}$ and then $R \succsim_2 \delta_{s''}$, since $\succsim_1$ is more averse than $\succsim_2$ to aggregate risk). The ranking $\delta_{s''} \succsim_1 R$ can be rewritten as

$$(f \circ \varphi_2)(t'') \geq (1-\epsilon)\left(\alpha(f\circ\varphi_2)(t) + (1-\alpha)(f\circ\varphi_2)(t')\right) + \epsilon(f\circ\varphi_2)(0).$$

Since ϵ is arbitrary, we obtain

$$f(\varphi_2(t'')) \geq \alpha f(\varphi_2(t)) + (1-\alpha) f(\varphi_2(t'))$$

which by (12) it implies f is concave. ■

B Relation with EEDEU

We say that a social preference relation $\succsim$ admits an EEDEU representation if

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

for some bounded and strictly increasing function $\phi : [0,1] \to \mathbb{R}$.

Proposition 1. *Let $\succsim$ satisfy Axioms 1-3. Then $\succsim$ admits an EEDEU representation if and only if it admits a utilitarian representation $U(s) = \int_I u_i(s_i)d\lambda(i)$.*

Proof: Given two profiles s and s', the relation $\succsim$ satisfies

$$\int_I u_i(s_i)d\lambda(i) \geq \int_I u_i(s_i')d\lambda(i) \iff \int_I \phi(u_i(s_i))d\lambda(i) \geq \int_I \phi(u_i(s_i'))d\lambda(i). \tag{13}$$

Let L be the set of probability distributions on $[0,1]$ with finite support. Since λ is non-atomic, for every $l \in L$ with support $\{t_1, \dots, t_n\} \subseteq [0,1]$ there exists a profile s such that

$$l(t_k) = \lambda(\{i : u_i(s_i) = t_k\}) \quad k = 1, \dots, n.$$

and hence for every bounded function $\psi \colon [0,1] \to \mathbb{R}$,

$$\sum_{t \in [0,1]} l(t)\psi(t) = \sum_{t \in [0,1]} \int_{\{i : u(s_i) = t\}} \psi(t) d\lambda(i) = \int \psi(u_i(s_i)) d\lambda(i)$$

It now follows from (13) that for every $l, m \in L$

$$\sum_{t \in [0,1]} l(t)t \geq \sum_{t \in [0,1]} m(t)t \iff \sum_{t \in [0,1]} l(t)\phi(t) \geq \sum_{t \in [0,1]} m(t)\phi(t).$$

Therefore, by the uniqueness properties of the expected utility representation, the function ϕ must be affine. In turn, this implies U is a utilitarian aggregator. ■

C 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 9. *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 below, the law of large numbers described in theorem 3 continues to hold, asymptotically, for finite populations.

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. ■

Proof of Theorem 9: Fix a sequence $\delta_n \downarrow 0$ and set $M = \max \varphi$, $L(n) = 1 - 2e^{-2n(\frac{\delta_n}{K})^2}$ and $A(n, P) = \phi\left(\frac{1}{n}\sum_{i\in I_n} E_{P_i}[u_i]\right)$. 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 $\epsilon_n = 2\left((1 - L(n))M + L(n)\delta_n\right)$. 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

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

■

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.

——— (2020): "An Abstract Law of Large Numbers," *Sankhya A*, 82(1), 1–12.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2018): "Multiattribute utility theory, intertemporal utility, and correlation aversion," *International Economic Review*, 59(2), 537–555.

Arrow, K. J. (2012): *Social choice and individual values*, vol. 12. Yale university press.

Bansal, R., and A. Yaron (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The journal of Finance*, 59(4), 1481–1509.

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. (1983): "Social welfare functions for economic environments with and without the Pareto principle," *Journal of Economic Theory*, 29(2), 205–216.

——— (1985): "More on Harsanyi's utilitarian cardinal welfare theorem," *Social Choice and Welfare*, 1(4), 279–281.

Bordes, G., and M. Le Breton (1989): "Arrovian theorems with private alternatives domains and selfish individuals," *Journal of Economic Theory*, 47(2), 257–281.

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.

Dillenberger, D., D. Gottlieb, and P. Ortoleva (2018): "Stochastic Impatience and the separation of Time and Risk Preferences," .

Epstein, L. G., and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, pp. 937–969.

Fleurbaey, M. (2010): "Assessing risky social situations," *Journal of Political Economy*, 118(4), 649–680.

Fleurbaey, M., and S. Zuber (2013): "Inequality aversion and separability in social risk evaluation," *Economic Theory*, 54(3), 675–692.

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.

Halmos, P. R. (1974): *Measure theory*. Springer.

Harsanyi, J. C. (1955): "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of political economy*, 63(4), 309–321.

——— (1986): *Rational behaviour and bargaining equilibrium in games and social situations*. CUP Archive.

Judd, K. L. (1985): "The law of large numbers with a continuum of iid random variables," *Journal of Economic theory*, 35(1), 19–25.

Kalai, E., E. Muller, and M. A. Satterthwaite (1979): "Social welfare functions when preferences are convex, strictly monotonic, and continuous," *Public Choice*, 34(1), 87–97.

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.

Keeney, R. L., and H. Raiffa (1993): *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.

Kihlstrom, R. E., and L. J. Mirman (1974): "Risk aversion with many commodities," *Journal of Economic Theory*, 8(3), 361–388.

Lorenzoni, G. (2008): "Inefficient credit booms," *The Review of Economic Studies*, 75(3), 809–833.

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.

McCarthy, D. (2015): "Distributive equality," *Mind*, 124(496), 1045–1109.

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.

Yaari, M. E. (1969): "Some remarks on measures of risk aversion and on their uses," *Journal of Economic Theory*, 1(3), 315–329.

Zhou, L. (1997): "Harsanyi's utilitarianism theorems: general societies," *Journal of Economic Theory*, 72(1), 198–207.