Journal of Mathematical Economics 87 (2020) 77–113

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Utilitarianism with and without expected utility

David McCarthy[a,*], Kalle Mikkola[b], Teruji Thomas[c]

[a] Department of Philosophy, University of Hong Kong, Hong Kong
[b] Department of Mathematics and Systems Analysis, Aalto University, Finland
[c] Global Priorities Institute, Faculty of Philosophy, University of Oxford, United Kingdom

ARTICLE INFO

Article history:
Received 25 May 2017
Received in revised form 28 December 2019
Accepted 2 January 2020
Available online 17 January 2020

Keywords:
Harsanyi
Utilitarianism
Expected and non-expected utility
Incompleteness and discontinuity
Uncertainty
Variable populations

ABSTRACT

We give two social aggregation theorems under conditions of risk, one for constant population cases, the other an extension to variable populations. Intra and interpersonal welfare comparisons are encoded in a single 'individual preorder'. The theorems give axioms that uniquely determine a social preorder in terms of this individual preorder. The social preorders described by these theorems have features that may be considered characteristic of Harsanyi-style utilitarianism, such as indifference to *ex ante* and *ex post* equality. However, the theorems are also consistent with the rejection of all of the expected utility axioms, completeness, continuity, and independence, at both the individual and social levels. In that sense, expected utility is inessential to Harsanyi-style utilitarianism. In fact, the variable population theorem imposes only a mild constraint on the individual preorder, while the constant population theorem imposes no constraint at all. We then derive further results under the assumption of our basic axioms. First, the individual preorder satisfies the main expected utility axiom of strong independence if and only if the social preorder has a vector-valued expected total utility representation, covering Harsanyi's utilitarian theorem as a special case. Second, stronger utilitarian-friendly assumptions, like Pareto or strong separability, are essentially equivalent to strong independence. Third, if the individual preorder satisfies a 'local expected utility' condition popular in non-expected utility theory, then the social preorder has a 'local expected total utility' representation. Fourth, a wide range of non-expected utility theories nevertheless lead to social preorders of outcomes that have been seen as canonically egalitarian, such as rank-dependent social preorders. Although our aggregation theorems are stated under conditions of risk, they are valid in more general frameworks for representing uncertainty or ambiguity.

1. Introduction

The subject of this paper is how to evaluate different assignments of welfare to members of society in the presence of risk. We thus consider *distributions* of welfare among individuals, and *lotteries*, probability measures over distributions. Each lottery determines for each relevant individual a *prospect*, a probability measure over welfare states. We assume that the value of prospects for the individuals facing them is represented by a preorder of prospects that we call the *individual preorder*, which thus encodes intra and interpersonal welfare comparisons under risk. We assume that the value of lotteries, from an impartial point of view, is represented by a preorder of lotteries that we call the *social preorder*.[1] We will say more about how to interpret welfare states and the individual and social preorders in Sections 2.1 and 3.2.

How should the individual and social preorders be related? In what we will refer to as his *utilitarian theorem*, Harsanyi (1955) proved (in a slightly different framework) that if the individual preorder satisfies expected utility theory, then it determines a unique social preorder satisfying expected utility theory, the strong Pareto principle, and a suitable condition of impartiality. This social preorder ranks lotteries by their expected total utility.[2]

Our main result is naturally seen as a generalization of Harsanyi's utilitarian theorem. It says that *any* individual preorder determines a unique social preorder satisfying three axioms related to Pareto and impartiality. These axioms are much weaker than Harsanyi's, and in particular we do not require either the

* Corresponding author.
E-mail addresses: mccarthy@hku.hk (D. McCarthy), kalle.mikkola@iki.fi (K. Mikkola), teru.thomas@oxon.org (T. Thomas).

1 A preorder is a reflexive, transitive binary relation.

2 We state expected utility and Pareto axioms in Sections 4.1 and 4.3 respectively. But it is already useful to state the most central expected utility axiom, to which we will often refer: a preorder $\succsim_X$ on a convex set X satisfies *strong independence* if for $p, p', q \in X$ and $\alpha \in (0, 1)$, $p \succsim_X p'$ if and only if $\alpha p + (1-\alpha)q \succsim_X \alpha p' + (1-\alpha)q$.

https://doi.org/10.1016/j.jmateco.2020.01.001
0304-4068/

individual or the social preorder to satisfy expected utility theory. Our first version of this result, Theorem 2.2, considers the constant population case, meaning that the same people exist in every social outcome; Theorem 3.5 extends the result to variable populations. In the variable case, our axioms do entail a mild constraint on the individual preorder related to the possibility of nonexistence that we call Omega Independence (Section 3.2). But it remains true that neither the individual nor the social preorder has to satisfy any axiom of expected utility.[3]

As we will explain, there are good reasons to think of the social preorders described by our theorems as being, like Harsanyi's, utilitarian in flavor; we will ultimately dub them *quasi utilitarian*.[4] Moreover, our weakening of Harsanyi's premises should surely be welcomed by utilitarians, as it provides our framework with considerable flexibility in the kinds of welfare comparisons it can accommodate. Recall that the three main expected utility axioms are completeness, continuity, and independence. In not requiring completeness, we allow for all kinds of incomparabilities between welfare states (and between welfare states and non-existence); in not requiring continuity, we allow some welfare states to be infinitely more valuable than others; in dropping independence, we allow for all sorts of views about welfare comparisons under risk.[5] Indeed, while independence has come to be seen as integral to Harsanyi-style utilitarianism, it is not so clear why it should be seen as a basic utilitarian commitment.[6]

Even so, one may wish to impose further constraints on the individual and social preorders. It turns out that many natural conditions bring quasi utilitarian theories closer to Harsanyi's utilitarianism. We explore this point systematically in sections 4 and 5.

We begin with expected utility in Section 4. If the individual preorder satisfies any one of the main axioms of expected utility theory, then so does the corresponding social preorder (Proposition 4.2). Moreover, if the individual preorder has an expected utility representation, then the social preorder is represented by total expected utility, or equivalently, expected total utility (Theorem 4.4). This is the conclusion of Harsanyi's utilitarian theorem, but extended to variable populations, and resting on much weaker premises (see Section 6.4). For example, strong Pareto and expected utility axioms for the social preorder are derived rather than assumed.

As we explain in Section 4.2, this result holds even if we allow the expected utility representation of the individual preorder to be vector valued; more precisely, we allow it to have values in a *preordered vector space*. Such 'Vector EU' representations generalize standard real-valued expected utility representations, the real numbers being the best known example of a preordered vector space; they also include the so-called multi expected utility representations that have been extensively discussed in the literature. The motivation for considering Vector EU representations is that their existence is equivalent to the central expected utility axiom, strong independence (Lemma 4.3). Thus while Vector EU representations are in many ways similar to ordinary real-valued expected utility representations, they allow for the denial of continuity and completeness. Overall, the upshot of Theorem 4.4 is that strong independence for the individual preorder is sufficient for a quasi utilitarian social preorder to be represented by vector-valued expected total utility, in both the constant and variable population cases. This provides a remarkably general version of Harsanyi's utilitarian theorem.

Still, as we mentioned in note 6, it is curious that any independence axiom should be seen as a fundamental premise of utilitarianism. The results of Section 4.3 are striking in this light. Proposition 4.8 shows that, for quasi utilitarian preorders, independence conditions on the individual preorder are essentially equivalent to corresponding Pareto conditions on the social preorder, and also to corresponding separability conditions. The conditions that correspond to strong independence are what we call Full Pareto, an apparently novel but natural extension of strong Pareto to cases involving incompleteness, and a version of strong separability. This means that, with one qualification, we can derive an expected total utility representation of the social preorder using one of those conditions instead of strong independence; the conceptual advantage is that strong separability and especially Full Pareto seem more central to traditional utilitarian concerns than strong independence. This leads to a particularly economical Harsanyi-like result in Theorem 4.10: Full Pareto plus just one of our aggregation axioms, Two-Stage Anonymity, implies an expected total utility representation of the social preorder. The qualification is that it is only a 'rational coefficients' version of strong independence that is exactly equivalent, given our axioms for aggregation, to Full Pareto. This means that the expected total utility representation just mentioned may be slightly less well behaved than the one that arises from strong independence.

Our results on non-expected utility theory in Section 5 paint a similar picture: we do not need to add much to our basic aggregation axioms to get close to Harsanyi-style utilitarianism. For example, monotonicity, or respect for stochastic dominance, is widely assumed in non-expected utility theory. But Theorems 5.2 and 5.3 show that, given some common background assumptions, monotonicity for the social preorder entails that the social preorder is once more represented by expected total utility. Even if we deny monotonicity, Theorems 5.5 and 5.7 show that when the individual preorder has a 'local expected utility' representation in the style of Machina (1982), the social preorder has a 'local expected total utility' representation.

Given these results, it is natural to ask just how close quasi utilitarian social preorders are to Harsanyi-style utilitarian preorders, and how close our axioms for aggregation are to

3 For notational convenience, we take the expected utility axioms to apply to preorders, so that reflexivity and transitivity do not count as expected utility axioms.

4 In speaking of 'utilitarianism', we do not need to take sides in the familiar debate about how Harsanyi's view relates to classical utilitarianism (see also note 6). What is clear is that his view has many formal properties that are by now closely associated with utilitarianism: besides the additive form, we especially have in mind Pareto, separability, and indifference to equality. Thus when discussing the utilitarian features of our social preorders, we only have in mind the extent to which they share such properties.

5 In normative settings, the status of the expected utility axioms has been heavily debated. For entries into a vast critical literature, see e.g. Pivato (2013) and Dubra et al. (2004) (completeness as deniable); Richter (1971) and Fishburn (1988) (continuity as a technical assumption); Luce and Raiffa (1957) and Kreps (1988) (continuity as deniable); and Allais (1979) and Buchak (2013) (independence as deniable). In addition, there are well-documented empirical violations of completeness and independence, and continuity has often been regarded as difficult to test for, raising a doubt about including it with other axioms in positive theories. See, for example, Starmer (2000), Schmidt (2004) and Wakker (2010). Here we mainly avoid positive topics, but see the beginning of Section 5 for a brief discussion.

6 While we stay neutral about the objection to a utilitarian interpretation of Harsanyi's results raised by Sen (1976, 1977) and Weymark (1991), their worry can nevertheless be seen as providing a further motivation for our project. Sen (1986, p. 1123) claims that classical utilitarianism starts with an "*independent* concept of individual utilities of which social welfare is shown to be the sum". In our terminology, the objection pressed by Sen and Weymark is essentially that *given* that the individual preorder satisfies expected utility theory, no reason has been offered for thinking that the cardinalization of welfare provided by expected utility theory coincides with the cardinalization assumed by classical utilitarianism (see Greaves, 2017 for discussion). But there is a perhaps more basic issue. If Sen and Weymark are right, then it does not seem that classical utilitarians are conceptually committed to expected utility theory in the first place. Thus even if one regards expected utility theory, and independence in particular, as normatively plausible, it is far from obvious that they should be seen as fundamental axioms of utilitarianism itself.

Harsanyi's own. We explore this question in Section 6 and draw connections to the literature. While our axioms are much weaker than Harsanyi's, they retain the indifference to *ex ante* and *ex post* equality that is integral to Harsanyi's approach (Section 6.1). More generally, Harsanyi's axioms are often said to combine *ex ante* and *ex post* requirements, and our axioms do this as well, albeit in a weakened sense (Section 6.2). On the other hand, quasi utilitarian social preorders may violate Pareto and separability principles, and failures of separability are sometimes associated with egalitarianism. In relation to this, Proposition 6.1 shows that *any* social preorder on distributions (rather than lotteries), even if apparently egalitarian, is compatible with quasi utilitarianism in the constant population case. For instance, Example 2.9 shows that imposing rank-dependent utility theory on the individual preorder leads to separability-violating rank-dependent social preorders that have been seen as canonically egalitarian. Overall, the sense in which all quasi utilitarian social preorders should be seen as utilitarian is somewhat equivocal. We recommend reserving 'utilitarian' for quasi utilitarian preorders that satisfy strong independence, in large part because they have Harsanyi-like expected total utility representations and are essentially the only ones that satisfy Full Pareto or strong separability (Section 6.3).

We briefly revisit Harsanyi's utilitarian theorem in Section 6.4, but Harsanyi had another famous contribution to the theory of social aggregation: the veil of ignorance construction of Harsanyi (1953), leading to his impartial spectator theorem, another constant population result with the same utilitarian conclusion. However, it is not clear how to justify the use of the veil, and when applied to variable population problems, it is far from obvious even how to interpret it. Nevertheless, our aggregation theorems can be seen as vindicating a specific version of the veil in both the constant and variable population cases (Section 6.5).

We briefly sketch an alternative strategy for generalizing Harsanyi's utilitarian theorem (Section 6.6), and end with a discussion of related literature (Section 6.7). For now we emphasize Pivato (2013) for a generalization of Harsanyi's utilitarian theorem that is closest to ours, and Mongin and Pivato (2015) for also deriving independence from Pareto in a Harsanyi-like framework, though one somewhat different from ours.

Let us make three final comments about our framework. First, our use of a single individual preorder to represent the value of prospects for individuals may appear controversial. For example, individuals may disagree about how to make welfare comparisons, or have different attitudes to risk. However, in Section 2.1 we provide several possible rationales, and also explain why it is compatible with standard presentations of Harsanyi's utilitarian theorem, including Harsanyi's own.

Second, it is nonetheless true that Harsanyi's most general result, Theorem V of Harsanyi (1955), does not rely on a single individual preorder, nor does it explicitly require the kind of interpersonal comparisons implicit in its use.[7] Given an expected utility function for each individual, its conclusion is that the social preorder can be represented by a weighted sum of these functions.[8] We refer to this result, not assuming interpersonal comparisons, as Harsanyi's *social aggregation theorem*; as already noted, we take Harsanyi's utilitarian theorem to be the result whose conclusion is that the social preorder can be represented by an *unweighted* sum of utility functions that have been normalized to reflect interpersonal comparisons. Harsanyi presents the utilitarian theorem as a mere corollary of the social aggregation theorem, and the literature tends to use the same terminology to refer to both results. But we find it helpful to distinguish them, since they indicate a two pronged strategy for generalizing Harsanyi's work: in McCarthy et al. (2019), without assuming interpersonal comparisons, we generalize Harsanyi's social aggregation theorem by dropping continuity and completeness while allowing the population to be infinite; in the present paper, we generalize Harsanyi's utilitarian theorem, assuming interpersonal comparisons at the outset in the form of a single individual preorder, but dropping continuity, completeness, *and* independence. Both approaches have their advantages. Harsanyi's social aggregation theorem, and the generalization just mentioned, may be applied to problems to which the present approach is inapplicable. For example, given a society with heterogeneous values or conflicting views about welfare, the social planner may wish to know how much can be concluded about social aggregation without resorting to controversial interpersonal comparisons. On the other hand, if, like Harsanyi, one is going to introduce interpersonal comparisons eventually, then one can obtain powerful results by assuming them from the outset and examining what other assumptions, like independence, can be modified or dropped.

Third, throughout this paper we work, for simplicity, in the setting of risk, where the uncertainty involved in each option is represented by a single probability measure, which one might think of as objective or universally agreed. But the principles underlying our aggregation theorems are much more general than this, and in Section 2.7 we briefly outline how they work for a variety of other representations of uncertainty, such as convex sets of probability measures and Anscombe–Aumann acts. This enables us to illustrate the relevance of our aggregation theorems to views according to which social evaluation should be based in part on a social consensus about uncertainty. Even if disagreement between individuals means that this social consensus cannot be represented by a single probability measure, our aggregation theorems can still cope.

Most proofs are in Appendix A, while Appendix B contains an index of common notation.

2. A constant population aggregation theorem

In Sections 2.1 to 2.4 we present the basic framework, axioms, and theorem. Section 2.5 introduces some useful terminology, Section 2.6 gives examples, and Section 2.7 explains how our framework and theorem could be adapted to other ways of representing uncertainty.

2.1. The individual and social preorders

Before introducing our formal axioms, we elaborate on our background assumptions, emphasizing our treatment of welfare comparisons and their role in our approach to aggregation. The questions of how to understand welfare and how to make welfare comparisons are much disputed, and our policy will to be remain as neutral as possible. Instead, we lay out our minimal commitments.

We assume that how well off an individual is in a particular social outcome can be adequately represented by specifying that individual's 'welfare state'. A welfare state takes into account all the information that is relevant to an assessment of an individual's overall welfare, with the implication that if two individuals are in identical welfare states, they are equally well off. A welfare state could be a single numerical indicator (see e.g. d'Aspremont and Gevers, 2002, p. 464); it could be a vector

7 Whether this result in some sense *implicitly* assumes interpersonal comparisons is a matter of controversy; see e.g. Harsanyi (1979, p. 294), Harsanyi (1978, p. 227), Harsanyi (1977b, pp. 81–2), Broome (1991, p. 219), Mongin (1994, pp. 348–50), and Mongin and d'Aspremont (1998, p. 432).

8 For elaboration, see e.g. Domotor (1979), Border (1981), Coulhon and Mongin (1989), Weymark (1993, 1995), and De Meyer and Mongin (1995).

with many components corresponding to, say, levels of happiness, pleasure, preference satisfaction, achievement, functioning, capabilities, resources, and so on (Sen, 1981, 1985); it could be a pairing of objective circumstances and subjective tastes (see note 11 and Pivato, 2013).[9] As we explain shortly, it is also possible to think of welfare states as being characterized abstractly in terms of intra and interpersonal comparisons. In any case, we assume that each social outcome can be adequately represented by specifying the welfare state of each member of the population. Thus our approach is welfarist in the sense that we are only concerned with the distribution of welfare under risk, rather than non-welfare factors.[10]

With this in mind, we adopt the following terminology, already sketched in the introduction. A 'distribution' is an assignment of welfare states to individuals. A 'lottery' is a probability measure over distributions. A 'prospect' is a probability measure over welfare states. Each lottery determines a prospect for each individual. The 'social preorder' expresses a view about how good lotteries are from an impartial perspective, while the 'individual preorder' expresses a view about how good prospects are for individuals, allowing interpersonal comparisons. The central question for us is how the social preorder should depend upon the individual preorder.

To elaborate, our basic assumption is that there is a single preorder, which we call the 'individual preorder', such that for any individuals i and j and prospects P and Q, P is at least as good for i as Q is for j if and only if P is ranked at least as highly as Q by the individual preorder.

As already noted, this assumption may seem implausible to anyone who understands welfare comparisons in terms of preferences, insofar as individuals may have different preferences over welfare states, or different attitudes to risk. However, there are several widely held views about welfare comparisons on which our assumption may be maintained.

First, one may claim that welfare comparisons should not be based on actual preferences, but rather on idealized preferences, formed under conditions of full information and rationality, and argue that at that level, individual preferences do agree.[11] Second, welfare comparisons may be taken to reflect objective evaluative facts; note that on such views, welfare states can still take into account levels of preference satisfaction. Third, welfare comparisons may be interpreted as reflecting the preferences or evaluative judgments of a single impartial person, the 'social planner'. Finally, welfare comparisons may be understood to reflect a (not necessarily unanimous) social consensus, perhaps formed by pooling individual preferences. Each of these positions is well established in the literature,[12] and we have no need to choose between them. We claim that our axioms for aggregation are plausible under each of these views, and we allow for any view that makes them plausible.

In fact, the possibility of using a single individual preorder to encode welfare comparisons is implied by the assumptions of Harsanyi's utilitarian theorem.[13] Harsanyi (1955) starts with a measurable space Y of social outcomes, and a set $\mathcal{P}(Y)$ of probability measures on Y. Each individual i is assumed to have a von Neumann–Morgenstern utility function u_i on Y. These utility functions are assumed to reflect intrapersonal and interpersonal comparisons of welfare (whether understood in terms of preferences or otherwise): for any $p, q \in \mathcal{P}(Y)$, p is at least as good for i as q is for j if and only if $\int_Y u_i \, dp \geq \int_Y u_j \, dq$. The utility functions allow us to define interpersonally comparable utility levels: i's utility level in outcome y is $u_i(y)$. Since each social outcome determines a utility level for each individual, these utility levels can be interpreted as welfare states. With this identification, each element of $\mathcal{P}(Y)$ determines a prospect for each individual, that is, a probability distribution over welfare states understood as utility levels. We can redescribe intra and interpersonal comparisons in terms of these prospects: for any $p, q \in \mathcal{P}(Y)$, p is at least as good for i as q is for j if and only if the prospect p assigns to i has at least as great an expected utility level as the prospect q assigns to j. So, just as we assume, such comparisons are encoded in a single preorder of prospects, namely, the preorder of prospects by expected utility. However, this use of utility levels to define welfare states is not as general as we would like, as it rests on the premise of expected utility theory, which we wish to avoid. Fortunately, an equivalent but more neutral understanding of welfare states is available in Harsanyi's framework: welfare states correspond exactly to equivalence classes in $Y \times \mathbb{I}$, where $\mathbb{I}$ is the set of individuals, and (y, i) is equivalent to (z, j) if and only if y is as good for i as z is for j. This abstract characterization of welfare states as equivalence classes makes sense even without the assumption of expected utility theory.

In short, many understandings of welfare states and justifications for using a single individual preorder are available. For our purposes, we will simply take welfare states and the individual preorder as primitives, and we will impose no formal requirements on the individual preorder beyond preordering, allowing for the flexibility in welfare comparisons outlined in Section 1, including the possibility of violating all of the expected utility axioms.

In summary, we adopt a two-stage approach to social aggregation. At the first stage, welfare comparisons at the individual level are expressed by a single, but possibly highly incomplete, preorder that we call the individual preorder. At the second stage, axioms are introduced to show how the social preorder is determined by the individual preorder. Our focus is on the second stage, and it is not our purpose to defend the two-stage approach as a whole. However, we have noted that it is implicit, and sometimes fully explicit, in presentations of Harsanyi's utilitarian theorem. It also provides one well known response to impossibility theorems concerning social aggregation in the face of individual disagreement about uncertainty (see Section 6.7.4). Aside from Harsanyi-like frameworks specifically involving risk or other forms of uncertainty, the two-stage approach also corresponds to the pioneering work of Sen (1970) in which the social planner first forms a view about interpersonal and intrapersonal comparisons before addressing questions about the social preorder (for discussion, see e.g. d'Aspremont and Gevers, 2002, sec. 1).

[9] See Adler and Fleurbaey (2016) for many essays on these possibilities.

[10] Given that we allow for many views about what constitutes welfare, and about the comparability of different welfare states, this makes our framework akin to 'formal welfarism' in social choice theory; see Mongin and d'Aspremont (1998), d'Aspremont and Gevers (2002), and Fleurbaey (2003).

[11] As a historically important example, Harsanyi (1977b, Ch. 4) assumes that each individual i has an 'extended preference relation' on gambles over 'extended alternatives'; an extended alternative specifies both an objective situation faced by an individual and the individual's subjective attitudes, like tastes. It would then be natural to identify welfare states with extended alternatives. Harsanyi argues in his Section 4.4 that fully rational individuals will have identical extended preferences even though their 'personal' preferences over objective situations may differ. Thus on Harsanyi's view, there is a unique rational extended preference relation, with which we could identify the individual preorder. However, Harsanyi's argument has been heavily criticized (Broome, 1993; Mongin, 2001a), and we are not committed to it. We shortly describe a different way of connecting our framework to Harsanyi's utilitarian theorem.

[12] See Adler and Fleurbaey (2016).

[13] There are several presentational variations of this result. For example, Blackorby et al. (1998) and Fleurbaey (2009) omit one of the steps we describe below and interpret social outcomes directly as profiles of interpersonally comparable von Neumann–Morgenstern utilities. And, closest to our approach, the extensions of Harsanyi's utilitarian theorem in Pivato (2013, 2014) explicitly use a single individual preorder. For more general discussion of interpersonal comparisons in this context, see Hammond (1991).

2.2. Framework

Formally, our basic framework starts with a set $\mathbb{W}$ of welfare states, and a finite, nonempty set $\mathbb{I}$ of individuals. We model social outcomes as what we call *distributions*, elements of $\mathbb{W}^{\mathbb{I}}$, the product of copies of $\mathbb{W}$ indexed by $\mathbb{I}$. We write $\mathcal{W}_i(d)$ for the ith component of distribution d, i.e. the welfare state that individual i has in that outcome. We focus on any set $\mathbb{D} \subset \mathbb{W}^{\mathbb{I}}$ of distributions that satisfies certain conditions shortly to be announced. Besides welfare states and distributions *per se*, we consider probability measures over them. Thus we assume that $\mathbb{W}$ and $\mathbb{D}$ are measurable spaces. We call probability measures over $\mathbb{W}$ *prospects*, and those over $\mathbb{D}$ *lotteries*. Notationally, if P is (say) a prospect and A is a measurable subset of $\mathbb{W}$, then we write $P(A)$ for the probability that P assigns to A. Instead of just considering all prospects and all lotteries, we will, for generality, focus on arbitrary non-empty convex sets $\mathbb{P}$ and $\mathbb{L}$ of prospects and lotteries respectively.

We make the following domain assumptions concerning the finite set $\mathbb{I}$, the measurable spaces $\mathbb{W}$ and $\mathbb{D} \subset \mathbb{W}^{\mathbb{I}}$, and the convex sets of probability measures $\mathbb{P}$ and $\mathbb{L}$.

(A) We assume that for each individual $i \in \mathbb{I}$ the projection $\mathcal{W}_i : \mathbb{D} \to \mathbb{W}$ is a measurable function. This allows us to define a prospect $\mathcal{P}_i(L)$ for each lottery L. Explicitly, if A is a measurable subset of $\mathbb{W}$, then

$$\mathcal{P}_i(L)(A) = L(\mathcal{W}_i^{-1}(A)).$$

We assume that $\mathcal{P}_i(\mathbb{L}) \subset \mathbb{P}$.

(B) For each $w \in \mathbb{W}$, we assume that $\mathbb{D}$ contains the distribution $\mathcal{D}(w)$ in which every individual $i \in \mathbb{I}$ has welfare w. Thus

$$\mathcal{W}_i(\mathcal{D}(w)) = w.$$

We assume that the function $\mathcal{D} : \mathbb{W} \to \mathbb{D}$ is measurable. This allows us to define a lottery $\mathcal{L}(P)$ for each prospect P. Explicitly, if B is a measurable subset of $\mathbb{D}$, then

$$\mathcal{L}(P)(B) = P(\mathcal{D}^{-1}(B)).$$

In $\mathcal{L}(P)$, every individual $i \in \mathbb{I}$ faces prospect P (that is, $\mathcal{P}_i(\mathcal{L}(P)) = P$), and it is certain that all individuals will have the same welfare.[14] We assume that $\mathcal{L}(\mathbb{P}) \subset \mathbb{L}$.

(C) We assume that $\mathbb{D}$ is invariant under permutations of individuals. Formally, let Σ be the group of permutations of $\mathbb{I}$. For each $\sigma \in \Sigma$ and $d \in \mathbb{D}$, the assumption is that $\mathbb{D}$ contains the distribution σd such that for all $i \in \mathbb{I}$,

$$\mathcal{W}_i(\sigma d) = \mathcal{W}_{\sigma^{-1} i}(d).$$

We assume that the action of Σ on $\mathbb{D}$ is measurable. That is: if $B \subset \mathbb{D}$ is measurable, then $\sigma^{-1} B$ is measurable, for any $\sigma \in \Sigma$. This allows us to define an action of Σ on lotteries L:

$$(\sigma L)(B) := L(\sigma^{-1} B)$$

for any $\sigma \in \Sigma$, lottery L and measurable $B \subset \mathbb{D}$. We assume that $\mathbb{L}$ is invariant under Σ.

Example 2.1. The various measurability conditions are not very stringent: for example, they are automatically met if $\mathbb{D} \subset \mathbb{W}^{\mathbb{I}}$ has the product sigma algebra, i.e. the *smallest* one for which the functions $\mathcal{W}_i$ are measurable. To check that $\mathcal{D}$ is measurable with respect to that sigma algebra, it suffices to check that $\mathcal{D}^{-1}(\mathcal{W}_i^{-1}(A))$ is measurable whenever A is a measurable subset of $\mathbb{W}$. But, in fact, $\mathcal{D}^{-1}(\mathcal{W}_i^{-1}(A)) = A$. Similarly, if $\mathbb{W}$ is a topological space, and we give $\mathbb{D} \subset \mathbb{W}^{\mathbb{I}}$ the product topology, then the measurability conditions will be met with respect to the Borel sigma algebras (even though the Borel sigma algebra on $\mathbb{D}$ is not necessarily the product one (Dudley, 2002, Prob. 4.1.11)).

2.3. Axioms for aggregation

Now we assume that $\mathbb{P}$ and $\mathbb{L}$ are each preordered. The preorder $\succsim_{\mathbb{P}}$ on $\mathbb{P}$ is the individual preorder; the preorder $\succsim$ on $\mathbb{L}$ is the social preorder. We discussed their significance in Section 2.1. We will use obvious notation, e.g. writing $P \sim_{\mathbb{P}} P'$ to mean the conjunction of $P \succsim_{\mathbb{P}} P'$ and $P' \succsim_{\mathbb{P}} P$, and writing $P \succ_{\mathbb{P}} P'$ to mean that $P \succsim_{\mathbb{P}} P'$ but not $P' \succsim_{\mathbb{P}} P$. Since $\succsim_{\mathbb{P}}$ is allowed to be incomplete, we will also write $P \curlywedge_{\mathbb{P}} P'$ to mean neither $P \succsim_{\mathbb{P}} P'$ nor $P' \succsim_{\mathbb{P}} P$.

We will sometimes informally treat the individual and social preorders as ranking not only prospects and lotteries but also welfare states and distributions respectively. Strictly speaking, this presupposes that we can identify welfare states and distributions with the corresponding delta-measures, for example, writing $w \succsim_{\mathbb{P}} w'$ to mean $1_w \succsim_{\mathbb{P}} 1_{w'}$. (Here, if y is an element of a measurable space Y, then the delta-measure 1_y is the unique probability measure on Y such that for any measurable set A, $1_y(A) = 1$ just in case A contains y.) This does not always make sense in our framework: different welfare states or different distributions may determine the same delta-measure (unless the sigma algebras separate points, an assumption we only take on in Section 6.1); and anyway these delta-measures may not be in the convex sets of probability measures under consideration. But we often ignore this detail in informal discussion.

Our first principle of aggregation says that the social preorder only depends on which prospect each individual faces.

Anteriority. If $\mathcal{P}_i(L) = \mathcal{P}_i(L')$ for every $i \in \mathbb{I}$, then $L \sim L'$.

Second, we need a principle which captures the idea that individual welfare contributes positively towards social welfare.

Reduction to Prospects. For any $P, P' \in \mathbb{P}$, $\mathcal{L}(P) \succsim \mathcal{L}(P')$ if and only if $P \succsim_{\mathbb{P}} P'$.

It says that for lotteries that guarantee perfect equality, social welfare matches individual welfare.

Anteriority can be seen as a very weak form of Pareto indifference,[15] which is obtained by replacing '$\mathcal{P}_i(L) = \mathcal{P}_i(L')$' with '$\mathcal{P}_i(L) \sim_{\mathbb{P}} \mathcal{P}_i(L')$'. In fact, Anteriority and Reduction to Prospects are both restrictions of a natural but apparently novel Pareto principle, which we call Full Pareto (see Section 4.3), that extends strong Pareto to cases involving incompleteness.[16] We further discuss Anteriority and Reduction to Prospects in Section 6.2.1, where we argue that together they express a weak sense in which the social preorder is *ex ante* (hence the term 'Anteriority').

Third, we need a principle of impartiality or permutation-invariance. The simplest such principle is

Anonymity. Given $L \in \mathbb{L}$ and $\sigma \in \Sigma$, we have $L \sim \sigma L$.

---

14 This second statement just means that any measurable subset of $\mathbb{D}$ containing the image of $\mathcal{D}$ has probability 1 according to $\mathcal{L}(P)$. We do not assume that the image of $\mathcal{D}$ is itself measurable. It may not be, even when $\mathbb{D}$ has the product sigma algebra, without modest further assumptions (Dravecký, 1975).

15 The expected utility and Pareto axioms mentioned in this section are formally defined in Sections 4.1 and 4.3.

16 The idea of restricting Pareto principles to lotteries that guarantee equality is familiar; see e.g. Fleurbaey (2010), McCarthy (2015), and Fleurbaey and Zuber (2017). But Reduction to Prospects appears to be novel even in the absence of risk. For example, it implies $w \curlywedge_{\mathbb{P}} w' \Rightarrow \mathcal{D}(w) \curlywedge \mathcal{D}(w')$. This inference is not licensed by any standard Pareto principle we know of, restricted or otherwise.

We will in fact use the following stronger condition.[17]

Two-Stage Anonymity. Given $L, M \in \mathbb{L}$, $\sigma \in \Sigma$, and $\alpha \in [0,1] \cap \mathbb{Q}$,

$$\alpha L + (1-\alpha)M \sim \alpha(\sigma L) + (1-\alpha)M.$$

One motivation for Two-Stage Anonymity is that it follows from the combination of Anonymity and the central axiom of expected utility theory, strong independence, or even the restriction of strong independence to the indifference relation. However, our preferred motivation for Two-Stage Anonymity avoids appealing to any independence axiom.

Define an 'anonymous distribution' to be an element of the quotient $\mathbb{D}/\Sigma$. One natural principle says that L and L' are equally good if they define the same probability measure over anonymous distributions.[18] Here is a convenient reformulation:

Posterior Anonymity. Given $L, L' \in \mathbb{L}$, suppose that $L(B) = L'(B)$ whenever B is a measurable, Σ-invariant subset of $\mathbb{D}$. Then $L \sim L'$.

In Section 6.2.2 we will argue that this principle expresses a weak sense in which the social preorder is *ex post*, hence the term 'Posterior'. Posterior Anonymity is easily seen to logically entail Two-Stage Anonymity, and that is our preferred motivation for accepting the latter as an axiom.[19]

Now strong independence is itself often said to be an *ex post* principle, so one might ask whether Two-Stage Anonymity is genuinely weaker than the conjunction of strong independence and Anonymity. But in Section 6.2.3 we give a precise sense in which Two-Stage Anonymity is much weaker. To anticipate, the following aggregation theorem, our main result, is compatible with rejecting any independence axiom for the individual and social preorders. In fact it is compatible with *any* individual preorder, and therefore with individual preorders that violate even the weakenings of independence axioms that are typical of non-expected utility theory.

2.4. The aggregation theorem

Now we state the main constant population result. We assume given domains $\mathbb{I}$, $\mathbb{W}$, $\mathbb{D}$, $\mathbb{P}$, and $\mathbb{L}$ satisfying the domain conditions **(A)**–**(C)** of Section 2.2.

Theorem 2.2. *Given an arbitrary preorder $\succsim_{\mathbb{P}}$ on $\mathbb{P}$, there is a unique preorder $\succsim$ on $\mathbb{L}$ satisfying Anteriority, Reduction to Prospects, and Two-Stage Anonymity. Namely,*

$$L \succsim L' \iff p_L \succsim_{\mathbb{P}} p_{L'} \tag{1}$$

where p_L (similarly $p_{L'}$) is the prospect

$$p_L = \frac{1}{\#\mathbb{I}} \sum_{i \in \mathbb{I}} \mathcal{P}_i(L).$$

[17] The use of only *rational* numbers α in stating Two-Stage Anonymity is simply a matter of precision: we do not require more. In fact, the obvious generalization to real α will hold for all our social preorders, as a consequence of Theorem 2.2.

[18] More precisely, we first use the quotient map $\pi: \mathbb{D} \to \mathbb{D}/\Sigma$ to push forward the sigma algebra on $\mathbb{D}$ to a sigma algebra on $\mathbb{D}/\Sigma$. Thus $C \subset \mathbb{D}/\Sigma$ is defined to be measurable if and only if its preimage $\pi^{-1}(C)$ is a measurable subset of $\mathbb{D}$. Note that π^{-1} defines a bijection between {measurable subsets of $\mathbb{D}/\Sigma$} and {measurable, Σ-invariant subsets of $\mathbb{D}$}. Next we push forward each $L \in \mathbb{L}$ to a probability measure L_Σ on $\mathbb{D}/\Sigma$: by definition, $L_\Sigma(C) = L(\pi^{-1}(C))$. Then the principle is that $L \sim L'$ if $L_\Sigma = L'_\Sigma$, and this is easily shown to be equivalent to Posterior Anonymity.

[19] Posterior Anonymity itself follows from Anonymity and the widely accepted principle of monotonicity, provided the social preorder is upper-measurable, a common domain assumption needed for monotonicity to apply (see Section 5.1). Anonymity is the case of Two-Stage Anonymity where $\alpha = 1$.

Proof. First let us show that if the social preorder satisfies the three conditions, then it is has the form (1). Consider the lottery $L_1 := \frac{1}{\#\Sigma}\sum_{\sigma\in\Sigma}\sigma L$. By repeated application of Two-Stage Anonymity, we have

$$L = \frac{1}{\#\Sigma}\sum_{\sigma\in\Sigma} L \sim \frac{1}{\#\Sigma}\sum_{\sigma\in\Sigma}\sigma L = L_1.$$

On the other hand, for any $i \in \mathbb{I}$,

$$\mathcal{P}_i(L_1) = \frac{1}{\#\Sigma}\sum_{\sigma\in\Sigma}\mathcal{P}_i(\sigma L) = \frac{1}{\#\Sigma}\sum_{\sigma\in\Sigma}\mathcal{P}_{\sigma^{-1}i}(L) = p_L.$$

By Anteriority, we must have $L_1 \sim \mathcal{L}(p_L)$, and so $L \sim \mathcal{L}(p_L)$. Similarly, we will have $L' \sim \mathcal{L}(p_{L'})$. Thus $L \succsim L'$ if and only if $\mathcal{L}(p_L) \succsim \mathcal{L}(p_{L'})$. By Reduction to Prospects, the latter holds if and only if $p_L \succsim p_{L'}$.

Now we must check that, conversely, the social preorder defined by (1) necessarily satisfies the three conditions. For Anteriority, suppose that $\mathcal{P}_i(L) = \mathcal{P}_i(L')$ for every $i \in \mathbb{I}$. Then clearly $p_L = p_{L'}$, so $L \sim L'$ by (1). As for Reduction to Prospects, (1) gives $\mathcal{L}(P) \succsim \mathcal{L}(P')$ if and only if $p_{\mathcal{L}(P)} \succsim_{\mathbb{P}} p_{\mathcal{L}(P')}$. However, this biconditional is equivalent to Reduction to Prospects since $p_{\mathcal{L}(P)} = P$ and $p_{\mathcal{L}(P')} = P'$. Finally, suppose given L, M, σ, α as in the statement of Two-Stage Anonymity. To deduce from (1) that $\alpha L + (1-\alpha)M \sim \alpha(\sigma L) + (1-\alpha)M$, it suffices to show that $p_{\alpha L+(1-\alpha)M} = p_{\alpha(\sigma L)+(1-\alpha)M}$. It is easy to see that $p_L = p_{\sigma L}$, and then we can calculate

$$\begin{aligned} p_{\alpha L+(1-\alpha)M} &= \alpha p_L + (1-\alpha)p_M = \alpha p_{\sigma L} + (1-\alpha)p_M \\ &= p_{\alpha(\sigma L)+(1-\alpha)M}. \quad \square \end{aligned}$$

Definition 2.3. We say that a social preorder $\succsim$ is *generated* by the individual preorder $\succsim_{\mathbb{P}}$ whenever the constant population domain conditions **(A)**–**(C)** hold and $\succsim$ satisfies (1). We call such social preorders *quasi utilitarian*.

We defend the 'quasi utilitarian' terminology in Section 6.3.

The following result shows that our favored principle of Posterior Anonymity, which implies Two-Stage Anonymity, could be used in place of the latter in Theorem 2.2.

Proposition 2.4. *If a social preorder is generated by an individual preorder, then it satisfies Posterior Anonymity.*

2.5. Representations

We now introduce some standard terminology which will be useful in the subsequent examples and results.

Definition 2.5. Given two preordered sets $(X, \succsim_X)$ and $(Y, \succsim_Y)$, a function $f: X \to Y$ *represents* $\succsim_X$ (or is a *representation* of $\succsim_X$) when, for all $x_1, x_2 \in X$, $x_1 \succsim_X x_2 \iff f(x_1) \succsim_Y f(x_2)$.

The mere existence of a representation is trivial; let $X = Y$ and f be the identity mapping. The interesting case is where $(Y, \succsim_Y)$ is better behaved or easier to understand or more fundamental than $(X, \succsim_X)$. For example, Y may be $\mathbb{R}$ with the usual ordering. For another example, the conclusion of Theorem 2.2 can be put by saying that the function $\mathbb{L} \to \mathbb{P}$ given by $L \mapsto p_L$ represents $\succsim$.

We will be much concerned with the case where $(Y, \succsim_Y)$ is a preordered vector space, or a slightly more general space that we call a $\mathbb{Q}$-preordered vector space. These are discussed in Sections 4.2 and 4.3, where we show that they are especially useful for making sense of vector-valued expected total utility representations in the absence of continuity or completeness assumptions.

2.6. Examples

Now let us give some examples of individual preorders and the social preorders they generate. For concreteness and simplicity, we will take $\mathbb{W}$ to be the real line $\mathbb{R}$ and take $\mathbb{P}$ to be the set of all finitely supported probability measures on $\mathbb{W}$.[20]

Example 2.6 (*Expected Utility and Total Utility*). Suppose that $\succsim_{\mathbb{P}}$ orders $\mathbb{P}$ by the expectations of a utility function $u: \mathbb{W} \to \mathbb{R}$. That is, $\succsim_{\mathbb{P}}$ is represented by the function $U: \mathbb{P} \to \mathbb{R}$ defined by $U(P) = \sum_{x\in\mathbb{W}} P(\{x\})u(x)$. (This sum, which has finitely many non-zero terms, can also be written as an integral $\int_{\mathbb{W}} u \, dP$.) The corresponding social preorder is represented by the function $V: \mathbb{L} \to \mathbb{R}$ given by $V = \sum_{i\in\mathbb{I}} U \circ \mathcal{P}_i$. We can identify $V(L)$ as the total expected utility of L, or equivalently as the expected total utility, since $V(L) = \sum_{i\in\mathbb{I}} \sum_{x\in\mathbb{W}} \mathcal{P}_i(L)(\{x\})u(x) = \sum_{d\in\mathbb{D}} L(\{d\}) \sum_{i\in\mathbb{I}} u(\mathcal{W}_i(d))$. For a more general statement and proof, see Theorem 4.4.

As we discuss in Section 4.1, the conceptual content of the assumption that $\succsim_{\mathbb{P}}$ has an ordinary (i.e. real-valued) expected utility representation is given by axioms of continuity, completeness, and independence. The next examples illustrate, for one thing, what can happen if one denies each of these axioms. In particular, the first two examples below illustrate the main lesson of Section 4: as long as the individual preorder satisfies strong independence, the social preorder still has an expected total utility representation. The last example illustrates the denial of strong independence.

Example 2.7 (*Leximin*). In this example the individual preorder satisfies strong independence and completeness, but not the axiom of mixture continuity. Let $\succsim_{\mathbb{P}}$ order $\mathbb{P}$ so that $P \succsim_{\mathbb{P}} P'$ if and only if either $P = P'$ or the smallest $x \in \mathbb{W}$ at which $P(\{x\}) \neq P'(\{x\})$ is such that $P(\{x\}) < P'(\{x\})$. When restricted to distributions, the corresponding social preorder is leximin: $d \succ d'$ if and only if the worst off individual in d is better off than the worst off in d'; if they are tied, turn to the next worst off. Although this seems quite different in flavor from Example 2.6, it becomes structurally very similar once we allow the utility function u to have values in a preordered vector space $\mathbb{V}$ rather than the real numbers. We develop this idea in Section 4.2, but a quick explanation is that, since one can average as well as add up vectors, it still makes sense to speak of the expected utility of a prospect, the total utility of a distribution, and the expected total utility of a lottery.[21] In this example, the vector space can be taken to be the space $\mathbb{V}$ of finitely supported functions $\mathbb{W} \to \mathbb{R}$. The 'lexicographic' ordering $\succsim_{\mathbb{V}}$ on $\mathbb{V}$ is defined by the condition that $f \succsim_{\mathbb{V}} g$ if and only if $f = g$ or the least $x \in \mathbb{W}$ for which $f(x) \neq g(x)$ is such that $f(x) > g(x)$. The utility function $u: \mathbb{W} \to \mathbb{V}$ is given by $u(x) = -\chi_{\{x\}}$, that is, minus the characteristic function of $\{x\}$. The social preorder is then represented by expected total utility just as in Example 2.6.

Example 2.8 (*Incompleteness*). Here the individual preorder satisfies strong independence and mixture continuity, but it is not in general complete. Let $\mathcal{U}$ be a set of real-valued functions on $\mathbb{W}$. Let $\succsim_{\mathbb{P}}$ preorder $\mathbb{P}$ so that $P \succsim_{\mathbb{P}} P'$ if and only if, for all u in $\mathcal{U}$, the expected value of u is at least as great under P as under P'. The corresponding social preorder ranks $L \succsim L'$ if and only if, for each u in $\mathcal{U}$, the expected total value of u is at least as great under L as under L'. In Section 4.2 we explain how this type of 'multi expected utility' representation by *many* real-valued utility functions is equivalent to an expected utility representation by a *single*, vector-valued utility function. With respect to this single utility function, the social preorder again ranks lotteries by their expected total utility.

Example 2.9 (*Risk-Avoidance and Rank-Dependence*). Finally, here is an example in which the individual preorder violates strong independence, even though it is complete and satisfies mixture continuity. This has interesting consequences for the social preorder: it illustrates a connection between strong independence and strong separability that we develop in Section 4.3.

Say that $\succsim_{\mathbb{P}}$ is a 'rank-dependent' individual preorder (RDI) if it has a 'rank-dependent utility' representation.[22] In other words, besides a utility function $u: \mathbb{W} \to \mathbb{R}$, there is an increasing function $r: [0, 1] \to [0, 1]$, with $r(0) = 0$ and $r(1) = 1$, which we will call the 'risk function'; $\succsim_{\mathbb{P}}$ is represented by $U: \mathbb{P} \to \mathbb{R}$ defined by the following sum (which has finitely many non-zero terms):

$$U(P) := \sum_{x\in\mathbb{W}} P_r(x)u(x), \quad \text{where } P_r(x) := r(P[x, \infty)) - r(P(x, \infty)).$$

If in addition r is convex, we will say that $\succsim_{\mathbb{P}}$ is 'risk-avoidant'.[23]

Although $U(1_w) = u(w)$ holds in general, and ordinary expected utility theory is satisfied when $r(x) = x$, $U(P)$ is not in general simply the expected utility of P. To see the deviation from ordinary expected utility, assume for concreteness $r(x) = x^2$ and $u(x) = x$. Consider the following distributions containing four individuals with listed welfare states.

$$d_A = (1, 1, 1, 1), \quad d_B = (5, 0, 1, 1)$$
$$d_C = (1, 1, 0, 0), \quad d_D = (5, 0, 0, 0).$$

Each of these distributions d_X determines a prospect $P_X := p_{1_{d_X}}$ that gives equal chances to each individual's welfare state; for example, P_B gives probability 1/4 to welfare states 5 and 0, and probability 1/2 to welfare state 1. Computing the value of U for each prospect yields $P_A \succ_{\mathbb{P}} P_B$ and $P_D \succ_{\mathbb{P}} P_C$. This has the structure of the Allais paradox, violating strong independence. For the corresponding social preorder, our aggregation theorem then implies that $d_A \succ d_B$ and $d_D \succ d_C$, violating strong separability.

Such violations of strong separability have been seen as expressions of egalitarianism. Thus it might be said that while the perfect equality in d_A outweighs the greater total welfare in d_B, there is not much difference in inequality between d_C and d_D, so the greater total in d_D is decisive (Sen 1973, p. 41, Broome 1989).

Returning to the general case, assume a population of size n. Say that a preorder $\succsim$ on distributions is a rank-dependent social preorder (RDS) if, for some $a_1, \ldots, a_n \geq 0$ with $\sum_k a_k = 1$, $\succsim$ ranks a distribution d with welfare states $w_1 \leq w_2 \leq \cdots \leq w_n$ according to the aggregate score

$$V(d) := a_1 u(w_1) + a_2 u(w_2) + \cdots + a_n u(w_n).$$

[20] Conceptually, a probability measure p on a measurable space Y is finitely supported if it is supported on a finite set (in the sense of note 33). Equivalently, though, it just means that p can be written as a convex combination of delta-measures. The equivalence, which we prove in Lemma A.4, holds regardless of whether the sigma algebra separates the points of Y, e.g. by making singletons measurable. For reasons not to require a separating sigma algebra, see Halpern (2003, §2.3).

[21] For the purpose of these examples, we can understand expectations as weighted sums, as in Example 2.6. In Section 4.2 we explain how to understand vector-valued expectations as integrals.

[22] Rank-dependent utility representations were introduced in Quiggin (1982). For further discussion of this popular non-expected utility theory, see e.g. Quiggin (1993), Wakker (1994), Schmidt (2004, §4.2), and Buchak (2013, Ch. 2). The function r is often required to be continuous and strictly increasing, but the weaker definition will be useful.

[23] This term is from Buchak (2013, p. 66); Yaari (1987) and Chateauneuf and Cohen (1994) use 'pessimistic'.

If in addition $a_1 \geq a_2 \geq \cdots \geq a_n$, we will say that $\succsim$ is 'downwards increasing'.

Downward increasing RDSs are called 'generalized Gini' by Blackorby and Donaldson (1980) and Weymark (1981), who take them to be natural examples of egalitarian preorders. We will say more about the relationship between apparently egalitarian preorders and our aggregation theorems in Section 6.1. But for now, by setting $a_k = r(\frac{n-k+1}{n}) - r(\frac{n-k}{n})$, we see that $\succsim$ is a [downwards increasing] RDS if and only if it is generated by a [risk-avoidant] RDI. Thus what has been taken to be a canonical form of egalitarianism at the social level emerges from what has been characterized as 'pessimism about risk' at the individual level. For example, by setting $r(x) = 1$ if $x = 1$, $r(x) = 0$ otherwise, we obtain the social preorder on distributions given by the Rawlsian maximin rule.

Curiously, though, the empirically best supported RDIs have S-shaped risk functions.[24] Provided the population is large enough, such RDIs lead to RDSs which are apparently inegalitarian at the high end, favoring unit transfers from the relatively well-off (but perhaps absolutely badly off) to the relatively better off. Given the lack of enthusiasm for inegalitarian ideas, this might call into question the sometimes mooted idea that people's attitudes to inequality reflect their attitudes to risk.

The examples illustrate how distributive views which are traditionally seen as very different can be obtained while maintaining our axioms for aggregation simply by varying the form of welfare comparisons.[25] General results corresponding to such possibilities will be given in Section 4.

2.7. Uncertainty

As laid out in Section 2.2, we model uncertainty using probability measures on sets of outcomes (whether welfare states or distributions). But analogues of our aggregation theorems hold for many other ways of modeling uncertainty. All we need is that we can take well-behaved mixtures (even just with rational coefficients) of the appropriate analogues of lotteries and prospects. Thus there is no difficulty in dealing with infinitesimal probabilities, non-additive 'capacities', or many other variations of standard probability theory.[26] Even in Savage's decision theory, in which there is no explicit representation of uncertainty, it is sometimes possible to endow the set of acts with convex structure, as in Ghirardato et al. (2003).

More formally, suppose we have a finite set $\mathbb{I}$ with permutation group Σ; convex sets $\mathbb{P}$ and $\mathbb{L}$ (or, more generally, associative mixture sets[27]); a mixture-preserving[28] map $\mathcal{L}: \mathbb{P} \to \mathbb{L}$ with, for each $i \in \mathbb{I}$, a mixture-preserving left-inverse $\mathcal{P}_i: \mathbb{L} \to \mathbb{P}$; and finally a mixture-preserving action of Σ on $\mathbb{L}$ such that $\mathcal{P}_i(\sigma L) = \mathcal{P}_{\sigma^{-1}i}(L)$ for every $i \in \mathbb{I}, L \in \mathbb{L}, \sigma \in \Sigma$. Then Theorem 2.2 makes sense as stated, and is still valid, with the same proof.[29]

Here are two more detailed illustrations. In both, we assume given the population $\mathbb{I}$ with permutation group Σ, as well as sets $\mathbb{P}$ and $\mathbb{L}$ of probability measures satisfying the domain conditions in Section 2.2; we use these to construct more complicated domains that do not themselves consist of probability measures.

Example 2.10 (*Convex Sets of Measures*)**.** In some choice frameworks one uses a convex set of probability measures, instead of a single one, to model uncertainty, as in, for example, the maxmin expected utility decision theory of Gilboa and Schmeidler (1989) and the Knightian decision theory of Bewley (2002). In any case, let $\mathrm{con}(\mathbb{P})$ and $\mathrm{con}(\mathbb{L})$ be the sets of nonempty convex subsets of $\mathbb{P}$ and $\mathbb{L}$. We can apply our aggregation theorem to relate an individual preorder $\succsim_{\mathrm{con}(\mathbb{P})}$ on $\mathrm{con}(\mathbb{P})$ to a social preorder $\succsim_{\mathrm{con}(\mathbb{L})}$ on $\mathrm{con}(\mathbb{L})$. To do this we first have to define suitable mixing operations on $\mathrm{con}(\mathbb{P})$ and $\mathrm{con}(\mathbb{L})$: for any $\alpha \in [0, 1]$ and $\mathbf{P}, \mathbf{Q} \in \mathrm{con}(\mathbb{P})$, set $\alpha\mathbf{P}+(1-\alpha)\mathbf{Q} = \{\alpha P+(1-\alpha)Q : P \in \mathbf{P}, Q \in \mathbf{Q}\}$, and similarly for $\mathrm{con}(\mathbb{L})$. Second, we need suitable maps $\mathcal{L}^{\mathrm{con}}: \mathrm{con}(\mathbb{P}) \to \mathrm{con}(\mathbb{L})$, $\mathcal{P}_i^{\mathrm{con}}: \mathrm{con}(\mathbb{L}) \to \mathrm{con}(\mathbb{P})$, and an action of Σ on $\mathrm{con}(\mathbb{L})$. Define $\mathcal{L}^{\mathrm{con}}(\mathbf{P}) = \{\mathcal{L}(P) : P \in \mathbf{P}\}$, $\mathcal{P}_i^{\mathrm{con}}(\mathbf{L}) = \{\mathcal{P}_i(L) : L \in \mathbf{L}\}$, and $\sigma\mathbf{L} = \{\sigma L : L \in \mathbf{L}\}$.

Example 2.11 (*Anscombe–Aumann*)**.** In the Anscombe and Aumann (1963) framework, perhaps the most popular decision-theoretic treatment of uncertainty, the objects of choice are probability-measure-valued functions on a set S of states of nature. In our setting, consider the function spaces $\mathbb{P}^S$ and $\mathbb{L}^S$. Our aggregation theorem can be used to relate an individual preorder $\succsim_{\mathbb{P}^S}$ on $\mathbb{P}^S$ to a social preorder $\succsim_S$ on $\mathbb{L}^S$. First we can define mixtures in $\mathbb{P}^S$, and similarly $\mathbb{L}^S$: for any $\mathbf{P}, \mathbf{Q} \in \mathbb{P}^S$ and $\alpha \in [0, 1]$, $(\alpha\mathbf{P}+(1-\alpha)\mathbf{Q})(s) = \alpha(\mathbf{P}(s))+(1-\alpha)(\mathbf{Q}(s))$. Then we need suitable maps $\mathcal{L}^S: \mathbb{P}^S \to \mathbb{L}^S$, $\mathcal{P}_i^S: \mathbb{L}^S \to \mathbb{P}^S$, and an action of Σ on $\mathbb{L}^S$. For this we can define $\mathcal{L}^S(\mathbf{P})(s) = \mathcal{L}(\mathbf{P}(s))$, $\mathcal{P}_i^S(\mathbf{L})(s) = \mathcal{P}_i(\mathbf{L}(s))$, and $(\sigma\mathbf{L})(s) = \sigma(\mathbf{L}(s))$.

The formalism explained in these examples can be interpreted in different ways, along the lines laid out in Section 2.1. But they especially illustrate the relevance of our aggregation theorems to the much discussed problem of social aggregation in the face of individual disagreement about uncertainty.[30] Assume, for example, a social choice perspective in which it is seen as desirable for social evaluation to reflect the social consensus about uncertainty. One model might assume that each individual is equipped with a subjective probability measure (or a convex set of measures), and regard the social consensus about uncertainty as represented by their convex hull (Brunnermeier et al., 2014; Alon and Gayer, 2016). Another could take each individual to be equipped with a preorder of Anscombe–Aumann acts, and see the social consensus about welfare comparisons under uncertainty as given by their intersection, or some extension thereof. As the examples illustrate, variations on our aggregation theorems still apply, even when, as in these cases, the social consensus about uncertainty can fall a long way short of being representable by a single real-valued probability measure.[31] We plan to discuss these and other examples in future work, but to focus on other problems, we stick here with our simpler framework in which uncertainty is always represented by a single probability measure.

---

24 See Schmidt (2004, §4.2.2) for references.

25 This reflects a theme in social choice theory, where, for example, classical utilitarianism and leximin can be derived from common axioms except for different assumptions about the measurability of welfare; see d'Aspremont and Gevers (1977).

26 One possibility particularly relevant to the themes of this paper is the use of probability-like measures with values in preordered vector spaces. In McCarthy et al. (2019) we explain how such 'vector measures' generalize other ways of representing likelihood, give a representation theorem, and apply this idea to the problem of opinion pooling.

27 For a definition of 'mixture set', see axioms (A1)–(A3) of Mongin (2001b), following Herstein and Milnor (1953); we call 'associative' a mixture set that also satisfies Mongin's axiom (A4). Example 2.10 will require this level of generality.

28 A map $f: X_1 \to X_2$ between convex sets, or more generally between mixture sets, is *mixture-preserving* (by some authors called 'affine' or 'convex-linear') if $f(\alpha x + (1-\alpha)y) = \alpha f(x) + (1-\alpha)f(y)$ for all $x, y \in X_1, \alpha \in [0, 1]$. Below, the action of Σ on $\mathbb{L}$ is assumed mixture-preserving in the sense that $L \mapsto \sigma L$ is mixture-preserving for each $\sigma \in \Sigma$.

29 An analogous remark applies to the variable population analogue Theorem 3.5, but we omit the details.

30 Section 6.7.4 gives references to discussions of this problem in the framework of Harsanyi's social aggregation theorem.

31 Of course, the literature on pooling uncertain opinion is vast, and we are not arguing for any particular approach. See Dietrich and List (2016) for an entry to the literature.

3. A variable population aggregation theorem

In this section we present a version of the aggregation theorem in which the population is allowed to vary from one distribution to another. In Sections 3.1–3.3 we present the basic framework, axioms, and theorem. In Section 3.4 we show that any constant population individual preorder can be extended, in many different ways, to a variable population one that generates a social preorder. In Sections 3.5 and 3.6 we consider some examples.

3.1. Framework

At a basic level, the generalization to variable populations is straightforward: we simply introduce a new element Ω representing nonexistence, and use an expanded set $\mathbb{W}^{\mathrm{v}} := \mathbb{W} \cup \{\Omega\}$ of welfare states. (In general, we will mark variable population objects with a superscript v, to distinguish them from their constant population analogues.) This allows each distribution to represent some individuals as nonexistent and, otherwise, Theorem 2.2 remains unchanged. To be sure, there are some questions of interpretation. For example, we will speak of Ω as a welfare state, but one need not take this literally. We will say more about comparisons involving Ω in Section 3.2.

The shortcoming of the approach just mentioned is there is only a finite set $\mathbb{I}$ of possible individuals. The interesting generalization is to allow the population size to be unbounded. We will, however, insist that any given lottery involves only finitely many individuals. We spell this out as assumption **(D$^{\mathrm{v}}$)** below.[32] In comparing two lotteries, then, only a finite population will be relevant, and we can apply the ideas of Section 2. Only a little more work is required to ensure that these pairwise comparisons combine into a well-defined social preorder. That is what we now explain.

Thus let $\mathbb{I}^{\infty}$ be an infinite set of possible individuals. Assume that $\mathbb{W}^{\mathrm{v}}$ and $\mathbb{D}^{\mathrm{v}} \subset (\mathbb{W}^{\mathrm{v}})^{\mathbb{I}^{\infty}}$ are measurable spaces, with $\Omega \in \mathbb{W}^{\mathrm{v}}$, and that $\mathbb{P}^{\mathrm{v}}$ and $\mathbb{L}^{\mathrm{v}}$ are non-empty convex sets of probability measures on $\mathbb{W}^{\mathrm{v}}$ and $\mathbb{D}^{\mathrm{v}}$ respectively. We make the following domain assumptions, parallel to those of Section 2.2.

(A$^{\mathrm{v}}$) We assume that, for each $i \in \mathbb{I}^{\infty}$, the projection $\mathcal{W}_i^{\mathrm{v}}: \mathbb{D}^{\mathrm{v}} \to \mathbb{W}^{\mathrm{v}}$ is measurable. This again allows us to define a function $\mathcal{P}_i^{\mathrm{v}}$ from lotteries to prospects, so that $\mathcal{P}_i^{\mathrm{v}}(L)(A) = L((\mathcal{W}_i^{\mathrm{v}})^{-1}(A))$ for measurable $A \subset \mathbb{W}^{\mathrm{v}}$. We assume that $\mathcal{P}_i^{\mathrm{v}}(\mathbb{L}^{\mathrm{v}}) \subset \mathbb{P}^{\mathrm{v}}$.

(B$^{\mathrm{v}}$) For each $w \in \mathbb{W}^{\mathrm{v}}$ and each finite population $\mathbb{I} \subset \mathbb{I}^{\infty}$, we assume that our set $\mathbb{D}^{\mathrm{v}}$ of distributions contains the distribution $\mathcal{D}_{\mathbb{I}}^{\mathrm{v}}(w)$ such that

$$\mathcal{W}^{\mathrm{v}}{}_i(\mathcal{D}_{\mathbb{I}}^{\mathrm{v}}(w)) = \begin{cases} w & \text{if } i \in \mathbb{I} \\ \Omega & \text{if not.} \end{cases}$$

We further assume that $\mathcal{D}_{\mathbb{I}}^{\mathrm{v}}: \mathbb{W}^{\mathrm{v}} \to \mathbb{D}^{\mathrm{v}}$ is measurable. We can then define a corresponding function $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}$ from prospects to lotteries. Thus if B is a measurable subset of $\mathbb{D}^{\mathrm{v}}$, $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(P)(B) = P((\mathcal{D}_{\mathbb{I}}^{\mathrm{v}})^{-1}(B))$. We assume that $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(\mathbb{P}^{\mathrm{v}}) \subset \mathbb{L}^{\mathrm{v}}$.

(C$^{\mathrm{v}}$) We assume that $\mathbb{D}^{\mathrm{v}}$ is invariant under permutations of $\mathbb{I}^{\infty}$. We write Σ^{∞} for the group of all such permutations. We further assume that the action of Σ^{∞} on $\mathbb{D}^{\mathrm{v}}$ is measurable. This allows us to define the action of Σ^{∞} on lotteries. We assume that $\mathbb{L}^{\mathrm{v}}$ is Σ^{∞}-invariant.

Finally, we will assume in **(D$^{\mathrm{v}}$)** that each distribution in $\mathbb{D}^{\mathrm{v}}$ and each lottery in $\mathbb{L}^{\mathrm{v}}$ involves only finitely many individuals. This requires careful formulation. For a distribution d, the assumption is that $\mathcal{W}_i^{\mathrm{v}}(d) = \Omega$ for all but finitely many $i \in \mathbb{I}^{\infty}$. One might guess that for a lottery L to 'involve only finitely many individuals', it would suffice that $\mathcal{P}_i(L) = 1_{\Omega}$ for all but finitely many $i \in \mathbb{I}^{\infty}$. But this is not conceptually the right criterion, as the following example shows.

Example 3.1. Suppose that $\mathbb{I}^{\infty} = [0, 1]$, and let d_i be the distribution in which only individual i exists, with welfare state w. Let L be the uniform probability measure over these d_i. Then each person i is certain not to exist – each has prospect 1_{Ω} – yet there is a clear sense in which L involves infinitely many individuals, rather than no individuals. Namely, for any finite population $\mathbb{I} \subset \mathbb{I}^{\infty}$, it is certain that someone not in $\mathbb{I}$ exists. One reason that this is problematic is that it would be natural to reject Anteriority in this example. Anteriority would say that L is just as good as no one existing at all, but intuitively it is rather as good as having one person who is certain to exist in welfare state w.

To state a better criterion, given finite $\mathbb{I} \subset \mathbb{I}^{\infty}$, let $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ be the subset of $\mathbb{D}^{\mathrm{v}}$ consisting of distributions d such that $\mathcal{W}_i^{\mathrm{v}}(d) = \Omega$ for all $i \notin \mathbb{I}$. We always consider $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ as a measurable space, with its sigma algebra restricted from the one on $\mathbb{D}^{\mathrm{v}}$. In other words, its measurable sets are those of the form $B \cap \mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$, with B measurable in $\mathbb{D}^{\mathrm{v}}$. The assumption we make is

(D$^{\mathrm{v}}$) We assume that each distribution $d \in \mathbb{D}^{\mathrm{v}}$ is a member of some $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$, and each lottery $L \in \mathbb{L}^{\mathrm{v}}$ is supported on some $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$,[33] with $\mathbb{I} \subset \mathbb{I}^{\infty}$ finite in both cases.

We write $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ for the subset of $\mathbb{L}^{\mathrm{v}}$ consisting of lotteries which are supported on $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$. In this notation, $\mathbb{I} \subset \mathbb{I}^{\infty}$ is always assumed to be finite. Note that, if $\mathbb{I}$ is contained in some larger population $\mathbb{I}'$, then $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}} \subset \mathbb{D}_{\mathbb{I}'}^{\mathrm{v}}$, and any lottery supported on $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ is also a lottery supported on $\mathbb{D}_{\mathbb{I}'}^{\mathrm{v}}$. Because of this, any two lotteries in $\mathbb{L}^{\mathrm{v}}$ are members of some common $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$, with $\mathbb{I} \subset \mathbb{I}^{\infty}$ finite.

Example 3.2. The various measurability assumptions are again guaranteed if $\mathbb{D}^{\mathrm{v}} \subset (\mathbb{W}^{\mathrm{v}})^{\mathbb{I}^{\infty}}$ has the product sigma algebra, or if $\mathbb{W}^{\mathrm{v}}$ is a topological space, $\mathbb{D}^{\mathrm{v}}$ has the product topology, and we use Borel sigma algebras (cf. Example 2.1). However, it would be natural to consider a finer-grained sigma algebra by including the sets $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$. Then $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ would be the subset of $\mathbb{L}^{\mathrm{v}}$ containing lotteries L such that $L(\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}) = 1$. But we do not need this assumption.

The implications of the domain assumptions are illustrated by the following lemma.

Lemma 3.3. *Assume the variable population domain conditions* **(A$^{\mathrm{v}}$)**–**(D$^{\mathrm{v}}$)**.

(i) Given $L \in \mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$, we have $\mathcal{P}_i^{\mathrm{v}}(L) = 1_{\Omega}$ for any $i \in \mathbb{I}^{\infty} \setminus \mathbb{I}$. In particular, $1_{\Omega} \in \mathbb{P}^{\mathrm{v}}$.

[32] Aggregating the welfare of infinitely many individuals raises quite formidable problems which we thus set aside. For example, assuming interpersonal comparisons as in Harsanyi's utilitarian theorem, full Anonymity is inconsistent with strong Pareto; see Bostrom (2011) for an overview of such problems and Pivato (2014) for a careful study of separable aggregation in the infinite setting with applications to the present setting of risk. Without assuming interpersonal comparisons, see Zhou (1997) for an infinite population version of Harsanyi's social aggregation theorem; Danan et al. (2015) for an infinite population version without assuming completeness; and McCarthy et al. (2019) for an infinite population version that dispenses with continuity and completeness.

[33] We say that a probability measure p on a measurable space Y is supported on $A \subset Y$ (not necessarily measurable) if $p(B) = 0$ whenever $B \subset Y$ is measurable and disjoint from A. More conceptually, the condition is that p is the pushforward to Y of a probability measure on A by the inclusion of A in Y, assuming that A is given the sigma algebra restricted from Y. This pushforward is a bijection between probability measures on A and probability measures on Y supported on A, and it is convenient to identify these things informally.

(ii) $\mathbb{D}^{\mathrm{v}}$ *contains the 'empty distribution'* d_Ω *such that* $\mathcal{W}_i^{\mathrm{v}}(d_\Omega) = \Omega$ *for all* $i \in \mathbb{I}^\infty$.
(iii) $\mathbb{L}^{\mathrm{v}}$ *contains the 'empty lottery'* 1_{d_Ω}, *and* $\mathcal{P}_i^{\mathrm{v}}(1_{d_\Omega}) = 1_\Omega$ *for all* $i \in \mathbb{I}^\infty$.
(iv) Suppose $\{\Omega\}$ *is measurable in* $\mathbb{W}^{\mathrm{v}}$. *If* $\mathcal{P}_i^{\mathrm{v}}(L) = 1_\Omega$ *for all* $i \in \mathbb{I}^\infty$, *then* $L = 1_{d_\Omega}$.

3.2. Axioms for aggregation

In parallel with the constant population case, we assume that there is an 'individual preorder' $\succsim_{\mathbb{P}^{\mathrm{v}}}$ on $\mathbb{P}^{\mathrm{v}}$, and a 'social preorder' $\succsim^{\mathrm{v}}$ on $\mathbb{L}^{\mathrm{v}}$. The significance of these preorders was discussed in Section 2.1, but there are some new subtleties having to do with the possibility of individual non-existence, as we explain shortly. The key axioms for aggregation are much as before, replacing constant population objects by variable population ones. The only notable point is that Reduction to Prospects must be formulated relative to every finite, non-empty subset of $\mathbb{I}^\infty$.

Anteriority (Variable Population). If $\mathcal{P}_i^{\mathrm{v}}(L) = \mathcal{P}_i^{\mathrm{v}}(L')$ for every $i \in \mathbb{I}^\infty$, then $L \sim^{\mathrm{v}} L'$.

Reduction to Prospects (Variable Population). For any $P, P' \in \mathbb{P}^{\mathrm{v}}$ and any finite, nonempty $\mathbb{I} \subset \mathbb{I}^\infty$, $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(P) \succsim^{\mathrm{v}} \mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(P')$ if and only if $P \succsim_{\mathbb{P}^{\mathrm{v}}} P'$.

Two-Stage Anonymity (Variable Population). Given $L, M \in \mathbb{L}^{\mathrm{v}}$, $\sigma \in \Sigma^\infty$, and $\alpha \in [0, 1] \cap \mathbb{Q}$,

$$\alpha L + (1-\alpha)M \sim^{\mathrm{v}} \alpha(\sigma L) + (1-\alpha)M.$$

In line with Theorem 2.2, Anteriority, Reduction to Prospects, and Two-Stage Anonymity will turn out to be satisfied by at most one social preorder. However, for such a social preorder to exist, we will need a condition on the individual preorder.[34]

Omega Independence. For any $P, P' \in \mathbb{P}^{\mathrm{v}}$ and rational number $\alpha \in (0, 1)$,

$$P \succsim_{\mathbb{P}^{\mathrm{v}}} P' \iff \alpha P + (1-\alpha)1_\Omega \succsim_{\mathbb{P}^{\mathrm{v}}} \alpha P' + (1-\alpha)1_\Omega.$$

We will present a defence of this condition, and discuss its relation to other independence axioms, in Section 3.4.

Let us comment on the justification for our three main axioms in the variable population context. Anteriority seems just as compelling as in the constant population case. And as in Section 2.3, our favored motivation for Two-Stage Anonymity is that it is entailed by Posterior Anonymity, now taking the following form.

Posterior Anonymity (Variable Population). Given $L, L' \in \mathbb{L}^{\mathrm{v}}$, suppose that $L(B) = L'(B)$ whenever B is a measurable, Σ^∞-invariant subset of $\mathbb{D}^{\mathrm{v}}$. Then $L \sim^{\mathrm{v}} L'$.

But Reduction to Prospects requires further comment. In Section 2.1, we endorsed the following interpretation of the constant population individual preorder:

(E) $P \succsim_{\mathbb{P}} Q$ if and only if, for any individuals i and j, facing prospect P is at least as good for i as facing Q is for j.

When the population consists of a single individual i, that is, when $\mathbb{I} = \{i\}$, it is also natural to suppose

(F) $P \succsim_{\mathbb{P}} Q$ if and only if $L \succsim L'$, where L and L' are the one-person lotteries in which the single person i faces the prospects P and Q respectively.

This is an instance of Reduction to Prospects. In the variable population case, one *could* extend **(E)** to the following:

(E$^{\mathrm{v}}$) $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$ if and only if, for any individuals i and j, facing prospect P is at least as good for i as facing Q is for j.

And similarly one has the following special case of Reduction to Prospects (Variable):

(F$^{\mathrm{v}}$) $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$ if and only if $L \succsim^{\mathrm{v}} L'$, where L and L' are the one-person lotteries in which the single person i faces the prospects P and Q respectively.

Granted the constant population **(E)**, **(F)** appears to be very plausible, although we discuss a possible anti-utilitarian objection to it at the end of Section 6.1. But the variable population case is more subtle.

Example 3.4. The following combination of views is well known, and often endorsed by utilitarians:

(a) A distribution consisting of one person with an excellent life is better, according to the social preorder, than a distribution containing no one at all; and, especially, a distribution consisting of one person with a terrible life, full of suffering, is worse than a distribution containing no one at all.
(b) Nonetheless, existing at a given welfare state cannot be better or worse *for an individual* than not existing at all.[35]

Granted **(E$^{\mathrm{v}}$)**, this combination of views violates **(F$^{\mathrm{v}}$)**, and hence Reduction to Prospects.

One option for defending Reduction to Prospects in the variable population case is to reject this combination of views. However, we prefer not to take a position on this. Instead, our assumption is that there is a preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ on prospects satisfying **(E)**, **(F)**, and **(F$^{\mathrm{v}}$)**, while remaining neutral on whether **(E$^{\mathrm{v}}$)** also holds.[36] However one settles this question, it seems to us that Anteriority and Two-Stage Anonymity, as well as Reduction to Prospects, retain their plausibility.

3.3. The aggregation theorem

Now we state the main variable population result. We assume given domains $\mathbb{I}^\infty$, $\mathbb{W}^{\mathrm{v}}$, $\mathbb{D}^{\mathrm{v}}$, $\mathbb{P}^{\mathrm{v}}$, and $\mathbb{L}^{\mathrm{v}}$ satisfying the domain conditions **(A$^{\mathrm{v}}$)**–**(D$^{\mathrm{v}}$)** of Section 3.1.

Theorem 3.5. *Given an arbitrary preorder* $\succsim_{\mathbb{P}^{\mathrm{v}}}$ *on* $\mathbb{P}^{\mathrm{v}}$, *there is at most one preorder* $\succsim^{\mathrm{v}}$ *on* $\mathbb{L}^{\mathrm{v}}$ *satisfying Anteriority, Reduction to Prospects, and Two-Stage Anonymity. When it exists, it is given by*

$$L \succsim^{\mathrm{v}} L' \iff p_L^{\mathbb{I}} \succsim_{\mathbb{P}^{\mathrm{v}}} p_{L'}^{\mathbb{I}} \tag{2}$$

for any finite non-empty $\mathbb{I} \subset \mathbb{I}^\infty$ *such that* L *and* L' *are lotteries in* $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$, *and where* $p_L^{\mathbb{I}}$ *(similarly* $p_{L'}^{\mathbb{I}}$*) is the prospect*

$$p_L^{\mathbb{I}} = \frac{1}{\#\mathbb{I}} \sum_{i \in \mathbb{I}} \mathcal{P}_i^{\mathrm{v}}(L).$$

It exists if and only if the individual preorder satisfies Omega Independence.

[34] Again, we can apply Theorem 2.2 to determine a unique social preorder on each $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ separately. The issue is whether these are compatible, in the sense of defining a social preorder on $\mathbb{L}^{\mathrm{v}}$ as a whole.

[35] This combination of views is endorsed, for example, by Broome (2004, cf. pp. 63–65) and Blackorby et al. (2005, cf. pp. 23–24); a classic argument for (b) is given in Broome (1999, p.168).

[36] For opposing utilitarian views, Broome (2004) and Blackorby et al. (2005) effectively reject **(E$^{\mathrm{v}}$)**, while, for example, Hammond (1991) would accept **(E$^{\mathrm{v}}$)** and **(F$^{\mathrm{v}}$)** as conceptual truths, serving to define individual welfare comparisons in terms of the social preorder.

Proof. Once we have fixed $\mathbb{I}$, the proof goes the same way as that of Theorem 2.2; for example, we define L_1 and L'_1 by summing over the group $\Sigma_{\mathbb{I}} \subset \Sigma^\infty$ of permutations of $\mathbb{I}$.

The only worry is that the comparison between L and L' defined by (2) might depend on $\mathbb{I}$, and that is where Omega Independence comes in. In detail, if $\mathbb{I} \subset \mathbb{I}'$ and $\#\mathbb{I} = m$ and $\#\mathbb{I}' = n$, then

$$p_L^{\mathbb{I}'} = \frac{m}{n} p_L^{\mathbb{I}} + \frac{n-m}{n} 1_\Omega.$$

Thus Omega Independence ensures the required independence of $\mathbb{I}$:

$$p_L^{\mathbb{I}'} \succsim p_{L'}^{\mathbb{I}'} \iff p_L^{\mathbb{I}} \succsim p_{L'}^{\mathbb{I}}.$$

To see that Omega Independence is a *necessary* condition, note that we can choose $\mathbb{I}$ and $\mathbb{I}'$ so that m/n equals any rational number $\alpha \in (0, 1)$. □

The following parallels Definition 2.3.

Definition 3.6. We say that a variable population social preorder $\succsim^{\mathrm{v}}$ is *generated* by the individual preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ whenever the variable population domain conditions **(A$^{\mathrm{v}}$)–(D$^{\mathrm{v}}$)** hold and $\succsim^{\mathrm{v}}$ satisfies (2). We call such social preorders *quasi utilitarian*.

The social preorders described by Theorem 2.2 turned out to automatically satisfy Posterior Anonymity. We can prove a similar result here, but we need a technical assumption. It would suffice to assume that $\mathbb{I}^\infty$ is countable—a modest limitation, given Anonymity and the fact that each lottery involves only finitely many individuals. However, we instead focus on a condition to the effect that there are plenty of measurable sets. Say that the sigma algebra on $\mathbb{D}^{\mathrm{v}}$ is *coherent* if the following holds: $B \subset \mathbb{D}^{\mathrm{v}}$ is measurable in $\mathbb{D}^{\mathrm{v}}$ if and only if, for every finite $\mathbb{I} \subset \mathbb{I}^\infty$, $B \cap \mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$ is measurable in $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$. (The left-to-right implication is automatic, since we defined the sigma algebra on $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$ to be the restriction of the one on $\mathbb{D}^{\mathrm{v}}$.) Note that coherence is a harmless assumption, in the sense that one can always expand the sigma algebra on $\mathbb{D}^{\mathrm{v}}$ to make it coherent without invalidating any of the domain conditions (see Lemma A.5 in Appendix A for details).

Proposition 3.7. *Suppose that the sigma algebra on $\mathbb{D}^{\mathrm{v}}$ is coherent, or that $\mathbb{I}^\infty$ is countable. If a variable population social preorder is generated by an individual preorder, then it satisfies Posterior Anonymity.*

In parallel to the constant population case, this shows that, granted coherence, Posterior Anonymity could be used in place of Two-Stage Anonymity in Theorem 3.5.

Remark 3.8. In the constant population case, there is no real difference between total and average utilitarianism. In this variable population setting, the fact that the definition of $p_L^{\mathbb{I}}$ involves 'averaging' over members of $\mathbb{I}$ may seem to suggest that (2) amounts to a form of average utilitarianism. But this impression is misleading: while $\mathbb{I}$ contains every individual who has a positive probability of existing under L or L', it is an arbitrary indexing set which may also contain individuals who are certain not to exist under L and L', and it can be replaced by any larger finite $\mathbb{I}' \supset \mathbb{I}$ without effect. In fact, one cannot say whether (2) should be seen as expressing a form of total utilitarianism, average utilitarianism, or something else, without more information about $\succsim_{\mathbb{P}^{\mathrm{v}}}$. Section 3.5 illustrates the extent to which theories with the form of total and average utilitarianism are compatible with (2), while Section 4 concludes that given (2), the social preorder has an expected total utility representation if and only if the individual preorder satisfies strong independence.

3.4. Omega independence

We now argue that Omega Independence is a fairly weak condition; in particular, it is compatible with any individual preorder on $\mathbb{P} \cup \{1_\Omega\}$.

To do this we need to be able to identify members of $\mathbb{P}$ with members of $\mathbb{P}^{\mathrm{v}}$. For this we assume that $\mathbb{P}$ is a (non-empty) convex set of probability measures on a measurable space $\mathbb{W}$, that $\mathbb{W}^{\mathrm{v}} = \mathbb{W} \cup \{\Omega\}$, and that $\mathbb{W}^{\mathrm{v}}$ has the sigma algebra generated by the one on $\mathbb{W}$. In other words, $A \subset \mathbb{W}^{\mathrm{v}}$ is measurable in $\mathbb{W}^{\mathrm{v}}$ if and only if $A \cap \mathbb{W}$ is measurable in $\mathbb{W}$; in particular, $\mathbb{W}$ and $\{\Omega\}$ are measurable in $\mathbb{W}^{\mathrm{v}}$. This enables us to identify members of $\mathbb{P}$ with probability measures on $\mathbb{W}^{\mathrm{v}}$ by the natural inclusion $P \mapsto P^{\mathrm{v}}$, where $P^{\mathrm{v}}(A) := P(A \cap \mathbb{W})$ for all measurable A in $\mathbb{W}^{\mathrm{v}}$.[37] We then identify $\mathbb{P}^{\mathrm{v}}$ with the convex hull of $\mathbb{P}_\Omega := \mathbb{P} \cup \{1_\Omega\}$. We summarize these assumptions by saying that $\mathbb{P}^{\mathrm{v}}$ *extends* $\mathbb{P}$. For any sets $X \subset Y$, we also say that a preorder $\succsim_Y$ on Y extends a preorder $\succsim_X$ on X if $x \succsim_X x' \iff x \succsim_Y x'$ for all $x, x' \in X$.

Proposition 3.9. *Assume that $\mathbb{P}^{\mathrm{v}}$ extends $\mathbb{P}$. Suppose given a preorder $\succsim_{\mathbb{P}}$. Let $\succsim_{\mathbb{P}_\Omega}$ be any preorder on $\mathbb{P}_\Omega$ that extends $\succsim_{\mathbb{P}}$. Then*

(i) There is a preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ on $\mathbb{P}^{\mathrm{v}}$ that extends $\succsim_{\mathbb{P}_\Omega}$ (and hence $\succsim_{\mathbb{P}}$) and satisfies Omega Independence.

(ii) There is a preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ on $\mathbb{P}^{\mathrm{v}}$ that extends $\succsim_{\mathbb{P}_\Omega}$ (and hence $\succsim_{\mathbb{P}}$) and violates Omega Independence.

The first part shows that Omega Independence is compatible with any preorder on $\mathbb{P}_\Omega$. For example, having fixed any $\succsim_{\mathbb{P}}$, Omega Independent $\succsim_{\mathbb{P}^{\mathrm{v}}}$ can be chosen so that for a given $P \in \mathbb{P}$, $1_\Omega \sim_{\mathbb{P}^{\mathrm{v}}} P$; alternatively, Omega Independent $\succsim_{\mathbb{P}^{\mathrm{v}}}$ can be chosen so that $1_\Omega \curlywedge_{\mathbb{P}^{\mathrm{v}}} P$ (or $1_\Omega \succ_{\mathbb{P}^{\mathrm{v}}} P$, or $P \succ_{\mathbb{P}^{\mathrm{v}}} 1_\Omega$) for all $P \in \mathbb{P}$. This provides the first sense in which Omega Independence is a weak condition.

The proposition as a whole shows that no matter how nonexistence is compared with other welfare states, Omega Independence of $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is logically independent of strong independence of $\succsim_{\mathbb{P}}$, despite the formal resemblance between these principles. In particular, because of the qualitative distinction between nonexistence and other welfare states, anyone who is moved by something like the Allais paradox to reject strong independence for $\succsim_{\mathbb{P}}$ might well accept Omega Independence for $\succsim_{\mathbb{P}^{\mathrm{v}}}$. In addition, even if $\succsim_{\mathbb{P}}$ satisfies strong independence, Omega Independence for $\succsim_{\mathbb{P}^{\mathrm{v}}}$ falls a long way short of implying strong independence for $\succsim_{\mathbb{P}^{\mathrm{v}}}$; Example 3.11 will illustrate this with a natural view about the value of nonexistence. These observations provide a second sense in which Omega Independence is weak.

Our variable population Theorem 3.5 shows that given a variable population domain, *any* Omega Independent individual preorder is compatible with our axioms of aggregation. That is a final sense in which Omega Independence is fairly weak, and in fact Theorem 3.5 gives strong reasons to accept Omega Independence. If the social preorder satisfies Anteriority, Reduction to Prospects and Two-Stage Anonymity, then the theorem tells us that the individual preorder must satisfy Omega Independence. This is an argument for Omega Independence from seemingly modest principles of aggregation.

Of course, someone who strongly objected to Omega Independence could instead interpret Theorem 3.5 as an impossibility theorem.[38]

[37] Another way to put this is that the sigma algebra on $\mathbb{D}^{\mathrm{v}}$ is the pushforward of the one on $\mathbb{D}$ by the inclusion of $\mathbb{W}$ in $\mathbb{W}^{\mathrm{v}}$, and members of $\mathbb{P}$ are identified with their pushforwards.

[38] We thank a referee for this observation.

3.5. Examples

In the following examples we assume that $\mathbb{P}^{\mathrm{v}}$ extends $\mathbb{P}$. In each example we give a general construction to show how a natural view about welfare comparisons involving nonexistence extends a given $\succsim_{\mathbb{P}}$ to an Omega Independent $\succsim_{\mathbb{P}^{\mathrm{v}}}$, illustrating Proposition 3.9(i). We then make the construction more concrete by further assuming the framework of Example 2.6, so that $\mathbb{P}$ is the set of finitely supported probability measures on $\mathbb{W} = \mathbb{R}$, implying that $\mathbb{P}^{\mathrm{v}}$ is the set of finitely supported probability measures on $\mathbb{W}^{\mathrm{v}} = \mathbb{W} \cup \{\Omega\}$, and $\succsim_{\mathbb{P}}$ is represented by expectations of a utility function $u\colon \mathbb{W} \to \mathbb{R}$.

Example 3.10 (*Total Utility and Critical Level Utilitarianism*)**.** One possibility for extending a given $\succsim_{\mathbb{P}}$ to $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is to identify some prospect $P_0 \in \mathbb{P}$ that is effectively interchangeable with Ω, in the sense that, for any $P \in \mathbb{P}$ and $\alpha \in [0, 1]$,

$$\alpha P + (1-\alpha)1_\Omega \sim_{\mathbb{P}^{\mathrm{v}}} \alpha P + (1-\alpha)P_0.$$

While this equivalence determines $\succsim_{\mathbb{P}^{\mathrm{v}}}$ in terms of P_0 and $\succsim_{\mathbb{P}}$, it does not guarantee that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies Omega Independence. But it does if $\succsim_{\mathbb{P}}$ satisfies strong independence.

To illustrate, in the framework of Example 2.6, suppose we extend the utility function $u\colon \mathbb{W} \to \mathbb{R}$ to a function $u\colon \mathbb{W}^{\mathrm{v}} \to \mathbb{R}$, and define $\succsim_{\mathbb{P}^{\mathrm{v}}}$ to be the individual preorder represented by expectations of this extension. This amounts to saying that Ω is interchangeable with any $P_0 \in \mathbb{P}$ that has expected utility $u(\Omega)$ (although there might not be such a P_0). The corresponding social preorder is represented by the expected value of $\sum_{i\in\mathbb{I}^\infty}(u \circ \mathcal{W}_i^{\mathrm{v}} - u(\Omega))$; see Theorem 4.4(iii).

Social orders of this type are also given by the 'critical level utilitarianism' of Blackorby et al. (2005), and the 'standardized total principle' of Broome (2004). These treatments do not formally give nonexistence a utility value. Instead, writing $\mathbb{I}(d)$ for the individuals who exist in distribution d, they posit some constant c such that the social preorder is represented by the expected value of $\sum_{i\in\mathbb{I}(d)}(u\circ\mathcal{W}_i - c)$. This constant is said to be a 'critical' or 'neutral' level of utility: an individual's existence in a given distribution contributes to social value to the extent that the utility of her welfare state exceeds c. Thus the social preorders described in the previous paragraph have a critical level utilitarian form with critical level $u(\Omega)$.

As we explain in Remark 4.5, we could normalize u so that $u(\Omega) = 0$. The stated representation of the social preorder would then have a total utility form; that is, $L \succsim^{\mathrm{v}} L'$ if and only if L has at least as much expected total utility. In Section 4 we consider very general expected total utility representations of the social preorder, but these might also be seen as general forms of critical level utilitarianism.

Example 3.11 (*Average Utilitarianism and Value Conditional on Existence*)**.** Here is a second way to extend a given $\succsim_{\mathbb{P}}$ to an Omega Independent $\succsim_{\mathbb{P}^{\mathrm{v}}}$. It works whether or not $\succsim_{\mathbb{P}}$ satisfies strong independence. The idea is that sure nonexistence is incomparable to any other prospect, while in other cases the value of a prospect P is to be identified with its value conditional on the existence of the individual.[39] So define $\succsim_{\mathbb{P}^{\mathrm{v}}}$ by the rule that, given $P, P' \in \mathbb{P}$ and $\alpha, \alpha' \in [0, 1]$,

$$\alpha P + (1-\alpha)1_\Omega \succsim_{\mathbb{P}^{\mathrm{v}}} \alpha' P' + (1-\alpha')1_\Omega \iff \begin{cases} \alpha, \alpha' > 0 \text{ and } P \succsim_{\mathbb{P}} P', \text{ or} \\ \alpha = \alpha' = 0. \end{cases}$$

Note that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ will violate strong independence (unless $\succsim_{\mathbb{P}}$ ranks all prospects as equal).

In the framework of Example 2.6, the resulting variable population social preorder can be seen as a version of *average utilitarianism*. It ranks lotteries by expected total utility divided by expected population size; call this function of lotteries Av_1. (Here and in the next example, the total utility of a distribution d is given by $\sum_{i\in\mathbb{I}(d)} u \circ \mathcal{W}_i$; nonexistence is not given a utility value.) The 'empty' lottery in which it is certain that no one exists is incomparable to the others.

This is an unusual version of average utilitarianism, but two more obvious versions are less well behaved. Ranking lotteries by expected average utility (Av_2) violates Anteriority. Alternatively, one could consider the expected utility conditional on existence for each individual who has a non-zero chance of existing, and then average over such individuals (Av_3). Ranking lotteries by this average then violates Two-Stage Anonymity.

To illustrate the differences between the three forms of average utilitarianism and total utilitarianism, consider the distributions $d_1 = (0, \Omega)$, $d_2 = (x, x)$, and $d_3 = (y, \Omega)$, written as utility profiles; thus one person exists in d_1 with utility 0, two people exist in d_2 with utility x, and one person exists in d_3 with utility y. Consider the lotteries $L = \frac{1}{2}1_{d_1} + \frac{1}{2}1_{d_2}$ and $L' = 1_{d_3}$. Let Tot compute expected total utility. We find $\mathrm{Av}_1(L) = \frac{2x}{3}$, $\mathrm{Av}_2(L) = \frac{x}{2}$, $\mathrm{Av}_3(L) = \frac{3x}{4}$, and $\mathrm{Tot}(L) = x$. On the other hand, $\mathrm{Av}_1(L') = \mathrm{Av}_2(L') = \mathrm{Av}_3(L') = \mathrm{Tot}(L') = y$. Thus the four different views can result in different judgments about L versus L'.

Example 3.12 (*Incomparability of Nonexistence*)**.** A third method of defining $\succsim_{\mathbb{P}^{\mathrm{v}}}$ may appeal to those who take to heart the view mentioned in Example 3.4(b) that nonexistence is incomparable to other welfare states. For $P, P' \in \mathbb{P}$ and $\alpha, \alpha' \in [0, 1]$, they may define

$$\alpha P + (1-\alpha)1_\Omega \succsim_{\mathbb{P}^{\mathrm{v}}} \alpha' P' + (1-\alpha')1_\Omega \iff \begin{cases} \alpha = \alpha' > 0 \text{ and } P \succsim_{\mathbb{P}} P', \text{ or} \\ \alpha = \alpha' = 0. \end{cases}$$

This invariably produces an individual preorder satisfying Omega Independence. However, it leads to widespread social incomparability: we will have $L \curlywedge^{\mathrm{v}} L'$ unless the expected population size under L equals that under L'. In the framework of Example 2.6, the social preorder ranks lotteries of the same expected population size by their expected total utility. For illustration, the lotteries L and L' introduced in the previous example are always incomparable. In the next subsection we give an example of a 'neutral-range' view that involves less widespread incomparability.

3.6. The Repugnant Conclusion

We now give some further examples organized around the 'Repugnant Conclusion' of Parfit (1986), which has played a central role in discussions of variable-population aggregation. This is the statement that for any distribution in which every individual has the same very high welfare state, there is a better distribution in which every individual has the same very low but positive welfare state, corresponding to a life barely worth living. For example, this is a consequence of critical level utilitarianism (Example 3.10), on the assumption that 'barely worth living' lives have utility above the critical level. Many people find the Repugnant Conclusion, or variations on it, as repugnant as the name suggests (see e.g. Parfit 1986, Hammond 1988, Blackorby et al. 1995).

Let w_0 be the welfare state of a life that is barely worth living, and W a much higher welfare state, representing an excellent quality of life. Let P_α be the prospect $\alpha 1_W + (1-\alpha)1_\Omega$, for

[39] Such an idea is emphasized, for example, by Fleurbaey and Voorhoeve (2016), and also seemingly endorsed by Harsanyi in correspondence reported in Ng (1983).

$\alpha \in [0, 1]$. Under the conditions of our variable population aggregation theorem, the Repugnant Conclusion amounts to the claim that $1_{w_0} \succ_{\mathbb{P}^\mathrm{v}} P_\alpha$, for some rational probability $\alpha \in (0, 1)$.

There are, at least formally, many ways in which this claim about prospects can be denied. Some we have already seen. The critical level utilitarianism of Example 3.10 holds that $P_\alpha \succ_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for any $\alpha \in (0, 1)$, as long as $u(\Omega)$ is above $u(w_0)$. The average utilitarianism of Example 3.11 similarly holds that $P_\alpha \succ_{\mathbb{P}^\mathrm{v}} 1_{w_0}$. And the highly incomplete social preorder of Example 3.12, ranking lotteries of the same expected population size by their expected total utility, holds that $1_{w_0} \curlywedge_{\mathbb{P}^\mathrm{v}} P_\alpha$.

In the first and third examples just mentioned, the individual preorder satisfies strong independence. As we have already advertised, this leads to a general form of expected total utility representation to be studied in Section 4. To illustrate the scope of this result, we now give two further examples of individual and social preorders that satisfy strong independence while avoiding the Repugnant Conclusion.

Example 3.13 (*Non-Archimedean Total Views*). In this example, people in welfare state w_0 contribute positively to the social value of a distribution, but no number of such people can contribute more than even one person in welfare state W.[40] The key condition on the individual preorder is that $P_\alpha \succ_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for every $\alpha \in (0, 1)$, even though, corresponding to $\alpha = 0$, $1_{w_0} \succ_{\mathbb{P}^\mathrm{v}} 1_\Omega$. This requires that the individual preorder violates mixture continuity and the closely related Archimedean axiom (see Section 4.1). As a concrete example, consider $\mathbb{V} = \mathbb{R}^2$, with the lexicographic ordering $\succsim_\mathbb{V}$; that is, $(x_1, x_2) \succsim_\mathbb{V} (y_1, y_2)$ if and only if either $x_1 > y_1$, or $x_1 = y_1$ and $x_2 \geq y_2$. Choose a utility function $u\colon \mathbb{W}^\mathrm{v} \to \mathbb{V}$ with $u(W) = (1, 0)$, $u(w_0) = (0, 1)$, and $u(\Omega) = 0$, and rank prospects by (component-wise) expectations of u. The corresponding social preorder ranks lotteries by expected total utility. Any distribution in which everyone has welfare state W is better than any distribution in which everyone has w_0.

Example 3.14 (*Neutral-Range Views*). In this example some welfare states, including w_0, are 'neutral' in the sense of being incomparable to Ω.[41] The Repugnant Conclusion is avoided because people in welfare state w_0 do not contribute positively to social value.[42] Such a view can be derived from an individual preorder satisfying the condition that $P_\alpha \curlywedge_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for α in some interval containing 0, while $P_\alpha \succ_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for α outside that interval. As a concrete example, suppose that $\mathbb{W} = \mathbb{R}$, with $w_0 = 1$ and $W = 100$. Let $\mathbb{V} = \mathbb{R}^2$ with the product preorder $\succsim_\mathbb{V}$: that is, $(x_1, x_2) \succsim_\mathbb{V} (y_1, y_2)$ if and only if $x_1 \geq y_1$ and $x_2 \geq y_2$. Define a utility function $u\colon \mathbb{W}^\mathrm{v} \to \mathbb{V}$ by $u(\Omega) = 0$ and $u(w) = (w+10, w-10)$ for $w \in \mathbb{W}$. Let the individual preorder be represented by (component-wise) expectations of u; note this is compatible with the natural ordering on $\mathbb{W}$. This makes w_0, and indeed any welfare state in the interval $(-10, 10)$, incomparable to Ω. In particular, P_α has expected utility $(110\alpha, 90\alpha)$ and 1_{w_0} has expected utility $(11, -9)$. Thus one finds that $P_\alpha \curlywedge_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for $\alpha \in [0, 1/10)$ and $P_\alpha \succ_{\mathbb{P}^\mathrm{v}} 1_{w_0}$ for $\alpha \in [1/10, 1]$. The corresponding social preorder ranks lotteries by expected total utility. From this one can deduce that a population of m people in welfare state W will be better than one of n people in welfare state w_0 as long as $n \leq 10m$; otherwise they are incomparable.

4. Expected utility

We now begin to explore more systematically the relationship between individual preorders and the social preorders they generate. What do natural constraints on the individual preorder tell us about the social preorder, and vice versa? In this section we focus on axioms related to expected utility theory, while in Section 5 we consider non-expected utility theory.

Section 4.1 presents the preliminary result that the social preorder inherits the most normatively central expected utility axioms from the individual preorder, in the sense that if the individual preorder satisfies a given axiom, then so does the social preorder it generates (and vice versa). This contrasts with common approaches in which the same expected utility axioms are imposed on the individual and social preorders; in our framework, this is often redundant.

Section 4.2 shows that if the individual preorder is represented by expected utility, then the social preorder it generates is represented by expected *total* utility. The continuity and completeness axioms of standard expected utility theory are often seen as normatively questionable, and some of the examples we have discussed may provide further reasons to drop them.[43] This is why we work with a vector-valued form of expected utility representation that relies only on strong independence, the most distinctive and normatively plausible axiom of expected utility theory.

Section 4.3 shows the equivalence (under the aggregation theorems) of various Pareto, independence, and separability axioms. Thus one might take Pareto or separability as fundamental and derive independence, since the former two axioms are arguably more central to the utilitarian project. It allows us to give our weakest axiomatization of an expected total utility representation of the social preorder, based solely on Two-Stage Anonymity and what we call Full Pareto, a natural strengthening of strong Pareto in the face of incompleteness.

4.1. Axioms

Let us review the main expected utility axioms before proving that they are inherited by the social preorder. At the heart of expected utility theory is the notion of independence. Several different independence axioms are possible, and, like other axioms from expected utility theory, they can be posited separately for either the individual or the social preorder. Thus we state them generically for a preorder $\succsim_X$ on a convex set X.

Independence axioms. Suppose given $p, p', q \in X$ and $\alpha \in (0, 1)$.

$(\mathrm{I_a})$ $p \sim_X p' \implies \alpha p + (1-\alpha)q \sim_X \alpha p' + (1-\alpha)q$.

$(\mathrm{I_b})$ $p \succ_X p' \implies \alpha p + (1-\alpha)q \succ_X \alpha p' + (1-\alpha)q$.

$(\mathrm{I_c})$ $p \curlywedge_X p' \implies \alpha p + (1-\alpha)q \curlywedge_X \alpha p' + (1-\alpha)q$.

Let $(\mathrm{I_1}) := (\mathrm{I_a})$, $(\mathrm{I_2}) := (\mathrm{I_a}) \wedge (\mathrm{I_b})$, and $(\mathrm{I_3}) := (\mathrm{I_a}) \wedge (\mathrm{I_b}) \wedge (\mathrm{I_c})$. These seem to be the reasonable packages of independence axioms. In particular, $(\mathrm{I_3})$ is equivalent to perhaps the best known independence axiom, strong independence, that is, $p \succsim_X p' \iff \alpha p+(1-\alpha)q \succsim_X \alpha p'+(1-\alpha)q$. Although the weaker independence axioms are often sufficient given other assumptions, Lemma 4.3

40 See e.g. Arrhenius and Rabinowicz (2015) and Thomas (2018) for recent discussions of such theories, which are often called 'non-Archimedean' or 'lexical'.

41 Example 3.12 is the extreme case in which all of $\mathbb{W}$ is in this neutral range.

42 This relatively popular kind of theory, often called a 'critical range' or 'neutral range' view, is developed by Broome (2004) and Blackorby et al. (2005), although these authors differ in how to interpret the relevant incomparability (see also Rabinowicz, 2009; Gustafsson, 2019). It is worth noting that these views are usually described using a set of real-valued utility functions, rather than a single vector-valued utility function (cf. Example 2.8); we connect the 'multi-utility' approach and our vectorial approach in Section 4.2.

43 We have in mind here especially the idea that Ω may be incomparable to other welfare states (Examples 3.11, 3.12 and 3.14) and the desire to avoid the Repugnant Conclusion (Section 3.6).

below strongly suggests that (I_3) should be seen as the core idea of expected utility.

Just as Omega Independence only quantified over scalars in $(0,1)\cap\mathbb{Q}$, we similarly define the *Rational Independence* axioms ($I_i^{\mathbb{Q}}$) for $i = 1,\dots,3$ as the corresponding independence axioms, but with α restricted to $(0,1)\cap\mathbb{Q}$. We will use these rational-coefficient axioms in Section 4.3.

Standard expected utility theory also assumes

Completeness (Comp). $\succsim_X$ is a complete preorder: for all $p,q\in X$, $p\succ_X q$ or $q\succ_X p$ or $p\sim_X q$.

The final main idea of standard expected utility is continuity, often understood to mean either one of the following two axioms.

Archimedean (Ar). For all $p,q,r\in X$, $p\succ_X q\succ_X r$ implies that there exist $\alpha,\beta\in(0,1)$ such that $\alpha p+(1-\alpha)r\succ_X q$ and $q\succ_X \beta p+(1-\beta)r$.

Mixture Continuity (MC). For all $p,q,r\in X$, the set $\{\alpha\in[0,1]:\alpha p+(1-\alpha)r\succsim_X q\}$ is closed in $[0,1]$, as is the set $\{\alpha\in[0,1]:q\succsim_X \alpha p+(1-\alpha)r\}$.[44]

Given (I_3) and (Comp), (Ar) is equivalent to (MC). But when $\succsim_X$ is incomplete, there is tension between the Archimedean and mixture continuity axioms, and one may have to choose between them.[45]

When X is equipped with a topology, many continuity conditions typically stronger than (MC) have been considered. The following is the most popular.

Continuity (Cont). $\{p\in X: p\succsim_X q\}$ and $\{p\in X: q\succsim_X p\}$ are closed for all $q\in X$.

One can only expect nice results about (Cont) if the basic operations on prospects and lotteries are themselves continuous. Say that *mixing is continuous* on X if for any $\lambda\in(0,1)$, $\lambda p+(1-\lambda)q$ is a continuous function of $p,q\in X$. In the constant population case, the basic assumption is as follows.

Topology (Top). $\mathbb{P}$ and $\mathbb{L}$ have topologies such that $\mathcal{L}$ and all the maps $\mathcal{P}_i$ are continuous, and mixing is continuous on $\mathbb{P}$.

In the variable population case, we need a further condition on the topology of $\mathbb{L}^{\mathrm{v}}$ that allows us to pass from continuity on each $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ to continuity on $\mathbb{L}^{\mathrm{v}}$ itself. Say that $\mathbb{L}^{\mathrm{v}}$ is *topologically coherent* if it satisfies the following condition: $X\subset\mathbb{L}^{\mathrm{v}}$ is closed if and only if $X\cap\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ is closed in $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ for every finite $\mathbb{I}\subset\mathbb{I}^\infty$, where $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ has a topology as a subspace of $\mathbb{L}^{\mathrm{v}}$. Thus in the variable population case we use

Topology (Variable Population) ($\mathrm{Top}^{\mathrm{v}}$). $\mathbb{P}^{\mathrm{v}}$ and $\mathbb{L}^{\mathrm{v}}$ have topologies such that all the maps $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}$ and $\mathcal{P}^{\mathrm{v}}_i$ are continuous, mixing is continuous on $\mathbb{P}^{\mathrm{v}}$, and $\mathbb{L}^{\mathrm{v}}$ is topologically coherent.

Example 4.1. Suppose that $\mathbb{W}^{\mathrm{v}}$ is a topological space, and give $\mathbb{D}^{\mathrm{v}}$ a topology as a subspace of $(\mathbb{W}^{\mathrm{v}})^{\mathbb{I}^\infty}$ with the product topology (cf. Example 3.2). Assuming that $\mathbb{P}^{\mathrm{v}}$ and $\mathbb{L}^{\mathrm{v}}$ consist of Borel measures, we can give them the weak topologies. That is, the topology on $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ is the coarsest one such that, for every bounded continuous $f:\mathbb{D}^{\mathrm{v}}\to\mathbb{R}$, the function $L\mapsto\int_{\mathbb{D}^{\mathrm{v}}} f\,\mathrm{d}L$ is continuous on $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$; similarly for $\mathbb{P}^{\mathrm{v}}$ with f bounded and continuous on $\mathbb{W}^{\mathrm{v}}$. Define a topology on $\mathbb{L}^{\mathrm{v}}$ by the condition that X is closed if and only if $X\cap\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ is closed in this weak topology on $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ for every finite $\mathbb{I}$.[46] It is then easy to check that ($\mathrm{Top}^{\mathrm{v}}$) holds.

Proposition 4.2 (*Inheritance*)**.** *Suppose that a (constant or variable population) social preorder is generated by an individual preorder. Then*

(i) Each of (Comp), (Ar), (MC), (I_i), *and* ($I_i^{\mathbb{Q}}$) *(for $i = 1,2,3$) is satisfied by the individual preorder if and only if it is satisfied by the social preorder.*

(ii) Assuming (Top) *or* ($\mathrm{Top}^{\mathrm{v}}$)*, the individual preorder satisfies* (Cont) *if and only the social preorder does.*

Thus the most normatively central expected utility axioms are all inherited by the social preorder. Similar results hold for many other normatively natural expected utility axioms.[47]

4.2. Expected utility representations

We saw in Proposition 4.2 that the standard axioms of expected utility theory are inherited by the social preorder. We now focus on the *conclusion* of expected utility theory, that is, on the existence of an expected utility representation. We show that such representations of the individual preorder yield expected *total* utility representations of the social preorder. This result works even for a very general kind of expected utility representation which, as we explain, requires only the independence axiom (I_3).

We again state the relevant conditions in terms of the generic preorder $\succsim_X$ on a convex set X, but in this subsection we further assume $X=\mathcal{P}(Y)$ for some convex set of probability measures $\mathcal{P}(Y)$ on a measurable space Y. In this case we say that $f:Y\to\mathbb{R}$ is $\mathcal{P}(Y)$*-integrable* if it is Lebesgue integrable with respect to all $p\in\mathcal{P}(Y)$.[48] Say that a function $U:\mathcal{P}(Y)\to\mathbb{R}$ is *expectational* if there is a $\mathcal{P}(Y)$-integrable function $u:Y\to\mathbb{R}$ such that $U(p)=\int_Y u\,\mathrm{d}p$. The basic form of an expected utility representation is as follows.

EUT. There is an expectational function $U:\mathcal{P}(Y)\to\mathbb{R}$ that represents $\succsim_X$. We say that U is an *EU representation* of $\succsim_X$.

Given the implausibility of completeness, however, there has been much interest in the following 'multi-utility' generalization of EUT.[49]

Multi EUT. There is a set $\mathcal{U}$ of expectational functions $\mathcal{P}(Y)\to\mathbb{R}$ such that for $p,q\in\mathcal{P}(Y)$, $p\succsim_X q\iff U(p)\geq U(q)$ for all $U\in\mathcal{U}$. We say that $\mathcal{U}$ is a *Multi EU representation* of $\succsim_X$.

44 This is the continuity axiom of Herstein and Milnor (1953).

45 For example, when incomplete $\succsim_X$ satisfies (I_3), it cannot satisfy both (MC) and a mild strengthening of (Ar) which is natural in the presence of incompleteness; see further Dubra (2011) and McCarthy and Mikkola (2018).

46 It does not follow automatically that the topology on $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ as a subspace of $\mathbb{L}^{\mathrm{v}}$ is the weak topology, as one might wish. But this does follow if $\{1_\Omega\}$ is closed in $\mathbb{P}^{\mathrm{v}}$, which is guaranteed e.g. if $\mathbb{W}^{\mathrm{v}}$ is metrizable (Bogachev, 2007, Cor. 8.2.4).

47 For example, the social preorder also inherits the strengthening of (Ar) mentioned in note 45, finite dominance axioms, and also countable dominance axioms (cf. Fishburn 1970, Hammond 1998) if $\mathbb{L}$ and $\mathbb{L}^{\mathrm{v}}$ are closed under countable mixing.

48 We follow Bogachev (2007, Def. 2.4.1) and other authors in not insisting that an integrable function must be measurable (although all results go through on that stronger notion of integrability). Still, if f is integrable with respect to p, then f coincides with a measurable function on some set of p-measure 1. (By the definition of integrability, there is a sequence (f_n) of simple functions converging to f on some set A of p-measure 1. It follows from (Bogachev, 2007, Thm. 2.1.5(v)) that the limit of $(f_n\chi_A)$ is a measurable function agreeing with f on A).

49 The general concept of a multi-representation was introduced in Ok (2002); see also Evren and Ok (2011). In the specific context of expected utility, see Dubra et al. (2004), Evren (2008, 2014), McCarthy et al. (2017a), and Gorno (2017).

However, if $\succsim_X$ has a Multi EU representation, it automatically satisfies (MC). Since our aggregation theorems allow for violations of all kinds of continuity axioms, we now consider a further kind of expected utility representation to cater for this possibility.

Here is the general set-up. A *preordered vector space* is a vector space $\mathbb{V}$ with a (possibly incomplete) preorder $\succsim_{\mathbb{V}}$ that is *linear* in the sense that $v \succsim_{\mathbb{V}} v' \iff \lambda v + w \succsim_{\mathbb{V}} \lambda v' + w$, for all $v, v', w \in \mathbb{V}$ and $\lambda > 0$.[50] So $\mathbb{R}$ with the standard ordering is one example; other examples for $(\mathbb{V}, \succsim_{\mathbb{V}})$ were described in Example 2.7 and Section 3.6. Given a preordered vector space $(\mathbb{V}, \succsim_{\mathbb{V}})$, we need a way of integrating $\mathbb{V}$-valued functions. Suppose we have a vector space $\mathbb{V}'$ of linear functionals on $\mathbb{V}$ that separates the points of $\mathbb{V}$.[51] A function $u: Y \to \mathbb{V}$ is *weakly $\mathcal{P}(Y)$-integrable with respect to* $\mathbb{V}'$ if there exists a function $U: \mathcal{P}(Y) \to \mathbb{V}$ such that $\int_Y \Lambda \circ u \, \mathrm{d}p = \Lambda \circ U(p)$ for all $\Lambda \in \mathbb{V}'$, $p \in \mathcal{P}(Y)$. In particular, every $\Lambda \circ u$ must be $\mathcal{P}(Y)$-integrable. We then define the *Pettis* or *weak* integral by setting $\int_Y u \, \mathrm{d}p := U(p)$. When $U: \mathcal{P}(Y) \to \mathbb{V}$ can be written in this form, we here also say that U is *expectational*.

Vector EUT. For some preordered vector space $(\mathbb{V}, \succsim_{\mathbb{V}})$ and some separating vector space $\mathbb{V}'$ of linear functionals on $\mathbb{V}$, there is an expectational function $U: \mathcal{P}(Y) \to \mathbb{V}$ that represents $\succsim_X$. We say that U is a *Vector EU* representation of $\succsim_X$.

An ordinary EU representation, as above, can be identified with a Vector EU representation with $(\mathbb{V}, \succsim_{\mathbb{V}}) = (\mathbb{R}, \geq)$, and with $\mathbb{V}'$ the set of linear maps from $\mathbb{R}$ to $\mathbb{R}$. We can also identify a Multi EU representation with a special kind of Vector EU representation. Indeed, for whatever index set I, equip the vector space $\mathbb{R}^I$ with the product preorder $\succsim_{\mathbb{R}^I}$, i.e. $x \succsim_{\mathbb{R}^I} y \iff x(i) \geq y(i)$ for all $i \in I$. Let $\mathbb{V}'$ be the span of the set of projections $x \mapsto x(i)$ of $\mathbb{R}^I$ onto $\mathbb{R}$; it clearly separates the points of $\mathbb{R}^I$. The weak integral with respect to $\mathbb{V}'$ with values in $\mathbb{R}^I$ is just the component-wise ordinary integral.[52] Then Multi EU representations of the form $\mathcal{U} = \{U_i: \mathcal{P}(Y) \to \mathbb{R} \mid i \in I\}$ correspond exactly to Vector EU representations $U: \mathcal{P}(Y) \to \mathbb{R}^I$; the correspondence is given by $U(p)(i) = U_i(p)$. However, since Multi EU representations imply both (I_3) *and* (MC), the following result shows that Vector EU representations are much more general.

Lemma 4.3. *Suppose $\succsim_X$ is a preorder on $\mathcal{P}(Y)$, a convex set of probability measures on a measurable space Y. Then $\succsim_X$ satisfies* (I_3) *if and only if it satisfies Vector EUT.*

This result shows that (I_3) is the only crucial axiom for expected utility theory in this setting, and makes it clear that the existence of a Vector EU representation is a normatively natural assumption.

Now let us apply these definitions in the context of our aggregation theorems. When combined with Lemma 4.3, the next theorem shows that if the individual preorder satisfies (I_3), then the social preorder is represented by *total expected utility*, or, equivalently, *expected total utility*.

Theorem 4.4 (*EUT Inheritance*)**.** *Suppose that a (constant or variable population) social preorder is generated by an individual preorder.*

(i) The individual preorder satisfies Vector EUT if and only if the social preorder does.

(ii) In the constant population case, if $\succsim_{\mathbb{P}}$ has a Vector EU representation

$$U(P) = \int_{\mathbb{W}} u \, \mathrm{d}P$$

then $\succsim$ has a Vector EU representation

$$V(L) = \sum_{i \in \mathbb{I}} U(\mathcal{P}_i(L)) = \int_{\mathbb{D}} \sum_{i \in \mathbb{I}} (u \circ \mathcal{W}_i) \, \mathrm{d}L.$$

(iii) In the variable population case, if $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has a Vector EU representation

$$U^{\mathrm{v}}(P) = \int_{\mathbb{W}^{\mathrm{v}}} u \, \mathrm{d}P$$

where U^{v} is normalized so that $U^{\mathrm{v}}(1_\Omega) = 0$,[53] *then $\succsim^{\mathrm{v}}$ has a Vector EU representation*

$$V^{\mathrm{v}}(L) = \sum_{i \in \mathbb{I}^\infty} U^{\mathrm{v}}(\mathcal{P}_i^{\mathrm{v}}(L)) = \int_{\mathbb{D}^{\mathrm{v}}} \sum_{i \in \mathbb{I}^\infty} (u \circ \mathcal{W}_i^{\mathrm{v}}) \, \mathrm{d}L. \quad (3)$$

Although stated for Vector EUT, the result holds for both ordinary EUT and Multi EUT as well (see the proof of part (i) for details).[54] Specialized to ordinary EUT, the claim that the constant population social preorder is represented by $V(L) = \int_{\mathbb{D}} \sum_{i \in \mathbb{I}} (u \circ \mathcal{W}_i) \, \mathrm{d}L$ is the conclusion of Harsanyi's utilitarian theorem, when translated into our framework, but resting on premises that are much weaker than his (see Section 6.4). Theorem 4.4 also includes the familiar fact that this expected value of the sum of individual utilities is identical to the sum of the expected values of individual utilities. This is sometimes put by saying that in Harsanyi's conclusion, *ex post* utilitarian social evaluation is equivalent to *ex ante* utilitarian social evaluation.[55] The general Vector EUT version allows for failures of continuity and completeness, but maintains the expected total utility form and *ex ante*/*ex post* equivalence. We have derived the same sort of expected total utility representation and *ex ante*/*ex post* equivalence in the variable population case.

Remark 4.5 (*Normalization*)**.** The main difference in the variable population case is the normalization condition on U^{v}. When utilities are values in a preordered vector space, one can add any constant to a utility function without changing the preorder it represents, allowing for different normalizations. Since we always have $1_\Omega \in \mathbb{P}^{\mathrm{v}}$ (Lemma 3.3(i)), the normalization $U^{\mathrm{v}}(1_\Omega) = 0$ used in Theorem 4.4 is always available. But other normalizations may be natural; for example, a utility value of zero is sometimes reserved for welfare states that are neutral, rather than good or bad, for the person in question.[56] Without imposing any normalization, (3) would become

$$V^{\mathrm{v}}(L) = \sum_{i \in \mathbb{I}^\infty} (U^{\mathrm{v}}(\mathcal{P}_i^{\mathrm{v}}(L)) - U^{\mathrm{v}}(1_\Omega)) = \int_{\mathbb{D}^{\mathrm{v}}} \sum_{i \in \mathbb{I}^\infty} (u \circ \mathcal{W}_i^{\mathrm{v}} - u(\Omega)) \, \mathrm{d}L. \quad (4)$$

---

50 A linear preorder, in our sense, is sometimes called a vector preorder.

51 Giving $\mathbb{V}$ the weak topology with respect to $\mathbb{V}'$ makes it a locally convex topological vector space whose dual is $\mathbb{V}'$ (Rudin, 1991, 3.10).

52 Note that in this case the weak topology on $\mathbb{V} := \mathbb{R}^I$ is just the product topology, and we can identify $\mathbb{V}'$ as the direct sum $\bigoplus_{i \in I} \mathbb{R}$ (see e.g. Kelley and Namioka (1963, Thm. 14.6)).

53 An equivalent normalization condition is $u(\Omega) = 0$, since $u(\Omega) = U^{\mathrm{v}}(1_\Omega)$. We explain the normalization more in Remark 4.5; it implies that, in the following formula, the sums have finitely many non-zero summands (see Lemma 3.3(i)).

54 In the case of Multi EUT, we can unwind parts (ii) and (iii) of the theorem in the following way. In the constant population case, if $\mathcal{U}$ is a Multi EU representation of $\succsim_{\mathbb{P}}$, then $\{\sum_{i \in \mathbb{I}} U \circ \mathcal{P}_i : U \in \mathcal{U}\}$ is a Multi EU representation of $\succsim$; in the variable case, if $\mathcal{U}^{\mathrm{v}}$ is a Multi EU representation of $\succsim_{\mathbb{P}^{\mathrm{v}}}$ with each $U^{\mathrm{v}} \in \mathcal{U}^{\mathrm{v}}$ normalized so that $U^{\mathrm{v}}(1_\Omega) = 0$, then $\{\sum_{i \in \mathbb{I}^\infty} U^{\mathrm{v}} \circ \mathcal{P}_i^{\mathrm{v}} : U^{\mathrm{v}} \in \mathcal{U}^{\mathrm{v}}\}$ is a Multi EU representation of $\succsim^{\mathrm{v}}$.

55 In frameworks in which the *ex ante* utilitarian evaluations are made using possibly differing individual subjective probabilities, an approach not considered here, it is well known that this equivalence can fail. We thank a referee for emphasizing this point.

56 The question of how these normalizations are related depends upon the interpretation of the individual preorder in cases of nonexistence as discussed at the end of Section 3.2.

Comparison with Example 3.10 shows that (4) can be seen as a very general version of the formula used to define critical level utilitarianism. It allows for failures of continuity and completeness, and can accommodate the popular view that there is a *range* of critical levels (see Section 3.6). In any case, we will continue to emphasize total utility representations like (3) rather than representations like (4) that make the critical level explicit.

Remark 4.6 (*Mixture-Preserving Representations*)**.** In Section 2.7 we noted that our aggregation theorems can be generalized to associative mixture sets of prospects and lotteries that do not necessarily consist of probability measures. In that general setting, expected utility representations do not make sense. However, one can still consider representations that are mixture preserving rather than expectational.[57] In analogy to Lemma 4.3 and Theorem 4.4, and with essentially the same proof, (I_3) for the individual preorder is still necessary and sufficient for the existence of a mixture-preserving representation with values in a preordered vector space (McCarthy et al., 2019, Lemma 16), and the social preorder then has a mixture-preserving representation by total utility. In addition to their generality and technical simplicity, an advantage of dealing with mixture-preserving representations is that, unlike Vector EU representations, they can always be given values in partially ordered rather than merely preordered vector spaces. This fits the natural thought that equally good prospects should have the same, rather than merely equally good, utilities.

Remark 4.7 (*Canonical Utility Spaces*)**.** One can use structure theorems for preordered and partially ordered vector spaces to make the utility spaces more concrete. In particular, mixture-preserving representations can always be taken into a product of what Hausner and Wendel (1952) call 'lexicographic function spaces'. Informally, this means that we can choose the utilities to be matrices of real numbers. The space of row-vectors is lexicographically ordered, and one matrix ranks higher than another if and only if it ranks higher in each row. Normatively natural constraints on the represented preorder correspond to dimensional restrictions on the matrices.[58]

4.3. Pareto, separability, and independence

In the previous subsection we indicated the power of strong independence (I_3) as a condition on the individual preorder: it allows us to derive an expected total utility representation of the social preorder. However, as we explained in Section 1 (especially note 6), it is not obvious that (I_3) is an axiom to which utilitarians are *conceptually* committed. We now show in Proposition 4.8 that, given the axioms of our aggregation theorems, the 'rational' independence axioms $(I_i^{\mathbb{Q}})$ introduced in Section 4.1 are equivalent to corresponding Pareto axioms, and also to corresponding separability axioms. As we will suggest, given the proximity of (I_3) and $(I_3^{\mathbb{Q}})$, this can be taken as an informal argument that (I_3) is a consequence of central utilitarian principles. Alternatively, Theorem 4.10 shows that Two-Stage Anonymity and a suitably strong Pareto principle are enough by themselves to yield a slightly more general kind of expected total utility representation without having to appeal to any independence axiom.

We will continue to consider both constant and variable population settings. However, in the constant population setting, the results are most striking, and easiest to state, if we consider a family of constant population models with populations of different sizes. If one accepts our constant population axioms for aggregation from Section 2.3 for one finite population, it is natural to accept them for every finite population. The same goes for various conditions like Pareto or separability.

Formally, a *constant population model* is any tuple $\mathbb{M} = (\mathbb{I}, \mathbb{W}, \mathbb{P}, \succsim_{\mathbb{P}}, \mathbb{D}, \mathbb{L}, \succsim)$ satisfying the constant population domain conditions **(A)**–**(C)** of Section 2.2. Similarly, a *variable population model* is any $\mathbb{M}^{\mathrm{v}} = (\mathbb{I}^{\infty}, \mathbb{W}^{\mathrm{v}}, \mathbb{P}^{\mathrm{v}}, \succsim_{\mathbb{P}^{\mathrm{v}}}, \mathbb{D}^{\mathrm{v}}, \mathbb{L}^{\mathrm{v}}, \succsim^{\mathrm{v}})$ satisfying the variable population domain conditions **(A$^{\mathrm{v}}$)**–**(D$^{\mathrm{v}}$)** of Section 3.1. And, given an infinite population $\mathbb{I}^{\infty}$, a *family* $\mathbb{F}$ of constant population models consists of a constant population model $(\mathbb{I}, \mathbb{W}, \mathbb{P}, \succsim_{\mathbb{P}}, \mathbb{D}_{\mathbb{I}}, \mathbb{L}_{\mathbb{I}}, \succsim_{\mathbb{I}})$ for each finite $\mathbb{I} \subset \mathbb{I}^{\infty}$. Note that $\mathbb{W}$, $\mathbb{P}$, and $\succsim_{\mathbb{P}}$ must be independent of $\mathbb{I}$.

We will present the following axioms in a way that applies to both variable population models and families of constant population models. Our convention so far has been to label variable population objects with the superscript 'v'. Here we will use the superscript '$*$' to cover both constant and variable cases: for example, $\mathbb{D}_{\mathbb{I}}^{*}$ stands for $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ if we are talking about a variable population model, and it stands for $\mathbb{D}_{\mathbb{I}}$ if we are talking about a family of constant population models. To make this work smoothly, given a variable population model, and finite $\mathbb{I} \subset \mathbb{I}^{\infty}$, we define $\succsim_{\mathbb{I}}^{\mathrm{v}}$ to be the restriction of $\succsim^{\mathrm{v}}$ to $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$. Thus in the new notation $\succsim_{\mathbb{I}}^{*}$ is invariably a preorder on $\mathbb{L}_{\mathbb{I}}^{*}$.[59]

With this background, suppose we are given either a variable population model, or a family of constant population ones. Let us state *Pareto* and *separability* axioms.

Because the individual preorder can be incomplete, Pareto axioms need to be stated with some care. We first define relations $\approx_{\mathbb{P}*}^{\mathbb{J}}$, $\rhd_{\mathbb{P}*}^{\mathbb{J}}$ and $\bowtie_{\mathbb{P}*}^{\mathbb{J}}$; these are ways of comparing lotteries with respect to a finite population $\mathbb{J}$. For any lotteries L and L' in $\mathbb{L}_{\mathbb{I}}^{*}$ and $\mathbb{J} \subset \mathbb{I}$:

$$L \approx_{\mathbb{P}*}^{\mathbb{J}} L' \iff \mathcal{P}_i^*(L) \sim_{\mathbb{P}*} \mathcal{P}_i^*(L') \text{ for all } i \in \mathbb{J}$$

$$L \rhd_{\mathbb{P}*}^{\mathbb{J}} L' \iff \mathcal{P}_i^*(L) \succ_{\mathbb{P}*} \mathcal{P}_i^*(L') \text{ for all } i \in \mathbb{J}$$

$$L \bowtie_{\mathbb{P}*}^{\mathbb{J}} L' \iff \mathcal{P}_i^*(L) \curlywedge_{\mathbb{P}*} \mathcal{P}_i^*(L'), \mathcal{P}_i^*(L) \sim_{\mathbb{P}*} \mathcal{P}_j^*(L), \text{ and}$$
$$\mathcal{P}_i^*(L') \sim_{\mathbb{P}*} \mathcal{P}_j^*(L') \text{ for all } i, j \in \mathbb{J}.$$

We might read $\approx_{\mathbb{P}*}^{\mathbb{J}}$, $\rhd_{\mathbb{P}*}^{\mathbb{J}}$ and $\bowtie_{\mathbb{P}*}^{\mathbb{J}}$ as, respectively, equally good, better, and equi-incomparable for all members of $\mathbb{J}$. To explain the last of these, suppose $\mathbb{I} = \{1, 2\}$ and consider the inference: $\mathcal{P}_i^*(L) \curlywedge_{\mathbb{P}*} \mathcal{P}_i^*(L')$ for $i = 1, 2 \implies L \curlywedge_{\mathbb{I}}^{*} L'$. This may seem natural: if L and L' are incomparable for both 1 and 2, they are incomparable. But suppose $\mathbb{W}^*$ includes welfare states v and w, and consider two distributions with two people each: $d = (v, w)$ and $d' = (w, v)$. Treating welfare states and distributions as degenerate prospects and lotteries, suppose $v \curlywedge_{\mathbb{P}*} w$. Then the inference just considered implies $d \curlywedge_{\mathbb{I}}^{*} d'$. But this violates any standard formulation of anonymity (in our framework, Two-Stage Anonymity). The use of $\bowtie_{\mathbb{P}*}^{\mathbb{J}}$ in the following axioms blocks this kind of inference.

Pareto axioms. Suppose given a variable population model, or a family of constant population ones. For finite $\mathbb{I} \subset \mathbb{I}^{\infty}$,

[57] See note 28 for the definition of 'mixture preserving'. Expectational functions are always mixture preserving, and the generalization is modest in the sense that mixture-preserving functions are expectational in the most commonly studied setting, where the domain X is a convex set of finitely supported probability measures on a measurable space Y with measurable singletons.

[58] Hausner and Wendel (1952) assumed completeness. In the case of incompleteness, representations involving lexicographic function spaces are given in Borie (2016), McCarthy et al. (2017b) (discussing dimensional restrictions), and Hara et al. (2019).

[59] The reader may notice that, given a variable population model, we can formally obtain a family of constant population models $(\mathbb{I}, \mathbb{W}^{\mathrm{v}}, \succsim_{\mathbb{P}^{\mathrm{v}}}, \mathbb{D}_{\mathbb{I}}^{\mathrm{v}}, \mathbb{L}_{\mathbb{I}}^{\mathrm{v}}, \succsim_{\mathbb{I}}^{\mathrm{v}})$, with the caveat that the set $\mathbb{W}^{\mathrm{v}}$ of welfare states in these models happens to contain Ω. So axioms and results about variable population models can sometimes be read directly off of axioms and results about families of constant population models. We find it clearer not to rely on this fact presentationally.

$L, L' \in \mathbb{L}^*_{\mathbb{I}}$, and any partition $\mathbb{I} = \mathbb{J} \sqcup \mathbb{K}$ with $\mathbb{J} \neq \emptyset$,

($\mathrm{P_a}$) $L \approx^{\mathbb{I}}_{\mathbb{P}*} L' \implies L \sim^*_{\mathbb{I}} L'$.

($\mathrm{P_b}$) $L \rhd^{\mathbb{J}}_{\mathbb{P}*} L'$ and $L \approx^{\mathbb{K}}_{\mathbb{P}*} L' \implies L \succ^*_{\mathbb{I}} L'$.

($\mathrm{P_c}$) $L \bowtie^{\mathbb{J}}_{\mathbb{P}*} L'$ and $L \approx^{\mathbb{K}}_{\mathbb{P}*} L' \implies L \curlywedge^*_{\mathbb{I}} L'$.

We will focus on the natural packages ($\mathrm{P_1}$) := ($\mathrm{P_a}$), ($\mathrm{P_2}$) := ($\mathrm{P_a}$) ∧ ($\mathrm{P_b}$), and ($\mathrm{P_3}$) := ($\mathrm{P_a}$) ∧ ($\mathrm{P_b}$) ∧ ($\mathrm{P_c}$). Of course, Pareto axioms are usually formulated with respect to a single finite population $\mathbb{I}$; we just apply them with respect to *every* finite $\mathbb{I} \subset \mathbb{I}^\infty$. Setting this aside, some of these packages have familiar names. ($\mathrm{P_1}$) is Pareto Indifference; ($\mathrm{P_2}$) is strong Pareto; but ($\mathrm{P_3}$) appears to be novel. We will call it Full Pareto.

The separability assumptions we consider only make sense under some further domain conditions. We want to be able to 'restrict' lotteries to a subpopulation $\mathbb{J}$. That is, suppose given finite populations $\mathbb{J} \subset \mathbb{I} \subset \mathbb{I}^\infty$. We first assume that for each $d \in \mathbb{D}^*_{\mathbb{I}}$, $\mathbb{D}^*_{\mathbb{J}}$ contains a (necessarily unique) distribution $\pi_{\mathbb{J}}(d)$ such that $\mathcal{W}^*_j(\pi_{\mathbb{J}}(d)) = \mathcal{W}^*_j(d)$ for each $j \in \mathbb{J}$. This defines a function $\pi_{\mathbb{J}}: \mathbb{D}^*_{\mathbb{I}} \to \mathbb{D}^*_{\mathbb{J}}$. We assume it is measurable.[60] Given $L \in \mathbb{L}^*_{\mathbb{I}}$, we can then define a pushforward probability measure $L|_{\mathbb{J}} = L \circ \pi_{\mathbb{J}}{}^{-1}$ on $\mathbb{D}^*_{\mathbb{J}}$. We further assume $L|_{\mathbb{J}}$ is in $\mathbb{L}^*_{\mathbb{J}}$. We thus have a restriction map $\mathbb{L}^*_{\mathbb{I}} \to \mathbb{L}^*_{\mathbb{J}}$, $L \mapsto L|_{\mathbb{J}}$. We summarize these assumptions by saying that *restrictions exist*.

Separability axioms. Suppose given a variable population model, or a family of constant population ones. Suppose that restrictions exist. For finite $\mathbb{I} \subset \mathbb{I}^\infty$, $L, L' \in \mathbb{L}^*_{\mathbb{I}}$, and any partition $\mathbb{I} = \mathbb{J} \sqcup \mathbb{K}$ with $\mathbb{J} \neq \emptyset$,

($\mathrm{S_a}$) $L|_{\mathbb{J}} \sim^*_{\mathbb{J}} L'|_{\mathbb{J}}$ and $L|_{\mathbb{K}} \sim^*_{\mathbb{K}} L'|_{\mathbb{K}} \implies L \sim^*_{\mathbb{I}} L'$.

($\mathrm{S_b}$) $L|_{\mathbb{J}} \succ^*_{\mathbb{J}} L'|_{\mathbb{J}}$ and $L|_{\mathbb{K}} \sim^*_{\mathbb{K}} L'|_{\mathbb{K}} \implies L \succ^*_{\mathbb{I}} L'$.

($\mathrm{S_c}$) $L|_{\mathbb{J}} \curlywedge^*_{\mathbb{J}} L'|_{\mathbb{J}}$ and $L|_{\mathbb{K}} \sim^*_{\mathbb{K}} L'|_{\mathbb{K}} \implies L \curlywedge^*_{\mathbb{I}} L'$.

We consider the natural combinations ($\mathrm{S_1}$) := ($\mathrm{S_a}$), ($\mathrm{S_2}$) := ($\mathrm{S_a}$) ∧ ($\mathrm{S_b}$), and ($\mathrm{S_3}$) := ($\mathrm{S_a}$) ∧ ($\mathrm{S_b}$) ∧ ($\mathrm{S_c}$). When $L|_{\mathbb{K}} \sim^*_{\mathbb{K}} L'|_{\mathbb{K}}$, ($\mathrm{S_3}$) says that the members of $\mathbb{K}$ can be ignored in the comparison between L and L'. That is to say, $L \succsim^*_{\mathbb{I}} L' \iff L|_{\mathbb{J}} \succsim^*_{\mathbb{J}} L'|_{\mathbb{J}}$. Thus ($\mathrm{S_3}$) can be seen as an axiom of strong separability across individuals.[61]

Separability is most interesting when the lotteries faced by $\mathbb{J}$ and $\mathbb{K}$ can vary independently. In the variable population case, it turns out that our basic domain conditions already ensure a supply of lotteries sufficient for our purposes. For a family of constant population models, the following suffices: say that the family is *compositional* if, for any partition $\mathbb{I} = \mathbb{J} \sqcup \mathbb{K}$, and any $P, Q \in \mathbb{P}$, there exists $L \in \mathbb{L}_{\mathbb{I}}$ such that $\mathcal{P}_j(L) = P$ for all $j \in \mathbb{J}$ and $\mathcal{P}_k(L) = Q$ for all $k \in \mathbb{K}$. For example, the family is compositional if each $\mathbb{D}_{\mathbb{I}}$ equals $\mathbb{W}^{\mathbb{I}}$ equipped with the product sigma algebra, and $\mathbb{L}_{\mathbb{I}}$ is the set of all lotteries on $\mathbb{D}_{\mathbb{I}}$ (Bogachev, 2007, Theorem 3.3.1).

Proposition 4.8 (*Equivalence of Pareto, Separability, and Independence*)**.** Constant Population. *Suppose given a compositional family $\mathbb{F}$ of constant population models, and that restrictions exist. Suppose that each social preorder $\succsim_{\mathbb{I}}$ is generated by $\succsim_{\mathbb{P}}$. Then, for $i = 1, 2, 3$:*

$\mathbb{F}$ *satisfies* (S_i) $\iff$ $\mathbb{F}$ *satisfies* (P_i) $\iff$

every $\succsim_{\mathbb{I}}$ *satisfies* ($\mathrm{I}^{\mathbb{Q}}_i$) $\iff$ $\succsim_{\mathbb{P}}$ *satisfies* ($\mathrm{I}^{\mathbb{Q}}_i$).

Variable Population. *Suppose given a variable population model $\mathbb{M}^{\mathrm{v}}$, and that restrictions exist. Suppose that the social preorder $\succsim^{\mathrm{v}}$ is generated by $\succsim_{\mathbb{P}^{\mathrm{v}}}$. Then, for $i = 1, 2, 3$:*

$\mathbb{M}^{\mathrm{v}}$ *satisfies* (S_i) $\iff$ $\mathbb{M}^{\mathrm{v}}$ *satisfies* (P_i) $\iff$

$\succsim^{\mathrm{v}}$ *satisfies* ($\mathrm{I}^{\mathbb{Q}}_i$) $\iff$ $\succsim_{\mathbb{P}^{\mathrm{v}}}$ *satisfies* ($\mathrm{I}^{\mathbb{Q}}_i$).

This result shows that, against the background of our aggregation theorems, there is little difference between Pareto, separability, and independence. It is true that Proposition 4.8 strictly speaking concerns rational independence axioms like ($\mathrm{I}^{\mathbb{Q}}_3$), but there is simply no plausible normative or descriptive theory that accepts ($\mathrm{I}^{\mathbb{Q}}_3$) while rejecting ($\mathrm{I_3}$). Examples mobilized against ($\mathrm{I_3}$), like the Allais paradox, are indeed always formulated using rational numbers as probabilities.

One could take this as an informal argument for ($\mathrm{I_3}$) from utilitarian principles such as ($\mathrm{P_3}$) and ($\mathrm{S_3}$), leading to the expected total utility representations of Theorem 4.4. Alternatively, we now show how to use Proposition 4.8 to derive a slightly more general kind of expected total utility representation of the social preorder directly from ($\mathrm{P_3}$) without assuming any independence condition.[62]

Say that $(\mathbb{V}, \succsim_{\mathbb{V}})$ is a $\mathbb{Q}$*-preordered vector space* if $\mathbb{V}$ is a real vector space and $\succsim_{\mathbb{V}}$ is a $\mathbb{Q}$-linear preorder, in the sense that for any $v, v', w \in \mathbb{V}$ and *rational* $\lambda > 0$, $v \succsim_{\mathbb{V}} v' \iff \lambda v + w \succsim_{\mathbb{V}} \lambda v' + w$. By allowing such a space of utilities, we can slightly generalize Vector EUT:

Rational Vector EUT. For some $\mathbb{Q}$-preordered vector space $(\mathbb{V}, \succsim_{\mathbb{V}})$ and some separating vector space $\mathbb{V}'$ of linear functionals on $\mathbb{V}$, there is an expectational function $U: \mathcal{P}(Y) \to \mathbb{V}$ that represents $\succsim_X$. We say that U is a *Rational Vector EU* representation of $\succsim_X$.

The significance of this definition is explained by the following analogue of Lemma 4.3.

Lemma 4.9. *Suppose $\succsim_X$ is a preorder on $\mathcal{P}(Y)$, a convex set of probability measures on a measurable space Y. Then $\succsim_X$ satisfies ($\mathrm{I}^{\mathbb{Q}}_3$) if and only if it satisfies Rational Vector EUT.*

Combined with Proposition 4.8, this allows us to derive an analogue of Theorem 4.4 that takes Full Pareto and Two-Stage Anonymity as the basic premises.[63]

Theorem 4.10. *Suppose given either a compositional family of constant population models or a variable population model, and that restrictions exist.*
Constant Population.

(i) Full Pareto and Two-Stage Anonymity, for each social preorder $\succsim_{\mathbb{I}}$ in the family, hold if and only if the individual preorder $\succsim_{\mathbb{P}}$ satisfies Rational Vector EUT and generates each $\succsim_{\mathbb{I}}$.

[60] Recall that each constant population space $\mathbb{D}_{\mathbb{I}}$ has its own sigma algebra, while each variable population space $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$ has the sigma algebra restricted from $\mathbb{D}^{\mathrm{v}}$. Recall also in what follows that, even in the variable population case, elements of $\mathbb{L}^{\mathrm{v}}_{\mathbb{J}}$ can be identified with probability measures on $\mathbb{D}^{\mathrm{v}}_{\mathbb{J}}$ (see note 33).

[61] A more common notion of strong separability says that, given $L, L', M, M' \in \mathbb{L}^*_{\mathbb{I}}$, with $L|_{\mathbb{J}} = M|_{\mathbb{J}}$, $L'|_{\mathbb{J}} = M'|_{\mathbb{J}}$, $L|_{\mathbb{K}} = L'|_{\mathbb{K}}$, and $M|_{\mathbb{K}} = M'|_{\mathbb{K}}$, one has $L \succsim^*_{\mathbb{I}} L'$ if and only if $M \succsim^*_{\mathbb{I}} M'$. Given a sufficiently rich domain of lotteries, our ($\mathrm{S_3}$) is equivalent to the slightly stronger claim that, in fact, $L \succsim^*_{\mathbb{I}} L'$ if and only if $L|_{\mathbb{J}} \succsim^*_{\mathbb{J}} L'|_{\mathbb{J}}$.

[62] That one can use Pareto *or* independence to derive an expected total utility representation, although in a somewhat different framework to ours, is emphasized by Mongin and Pivato (2015, p. 159); see also Pivato (2014, pp. 39–40). In Theorems 5.2 and 5.3 we show that, in one common setting, expected total utility representations follow from our aggregation theorems without assuming *any* independence, Pareto, or separability condition; we merely need monotonicity for the social preorder.

[63] In both the constant and variable population cases, Posterior Anonymity could be used in place of Two-Stage Anonymity, granted coherence (or countable $\mathbb{I}^\infty$) in the variable population case. See Proposition 3.7 and its preceding commentary, where coherence was defined.

(ii) If $\succsim_{\mathbb{P}}$ has a Rational Vector EU representation U, and $\succsim_{\mathbb{P}}$ generates each $\succsim_{\mathbb{I}}$, then each $\succsim_{\mathbb{I}}$ has a Rational Vector EU representation $\sum_{i\in\mathbb{I}} U \circ \mathcal{P}_i$.

Variable Population.

(iii) Full Pareto and Two-Stage Anonymity hold if and only if $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies Rational Vector EUT and generates $\succsim^{\mathrm{v}}$.

(iv) If $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has a Rational Vector EU representation U^{v}, normalized so that $U^{\mathrm{v}}(1_\Omega) = 0$, and $\succsim_{\mathbb{P}^{\mathrm{v}}}$ generates $\succsim^{\mathrm{v}}$, then $\succsim^{\mathrm{v}}$ has a Rational Vector EU representation $\sum_{i\in\mathbb{I}^\infty} U^{\mathrm{v}} \circ \mathcal{P}^{\mathrm{v}}_i$.

Just as in Theorem 4.4, the conveniently brief 'total expected utility' form of representation can be rewritten as expected total utility. So, to emphasize: this result shows that, given a rich enough domain, Full Pareto and Two-Stage Anonymity (or Full Pareto and Posterior Anonymity) are by themselves enough to yield an expected total utility representation of the social preorder (or of each one in the family), with an unusually general, but still well-behaved, space of utilities. However one feels about these general utility spaces, the fundamental point is that Full Pareto and Two-Stage Anonymity are enough to determine the social preorder in terms of the individual preorder, while guaranteeing separability (S_3) and at least the rational version of strong independence, ($I_3^{\mathbb{Q}}$). We give the proof of Theorem 4.10 in Appendix A, but a sketch will illustrate the perhaps surprising power of Full Pareto. Full Pareto entails both Anteriority and Reduction to Prospects, so Two-Stage Anonymity is the only one of our aggregation axioms then needed to show that the social preorder is generated by the individual preorder. Using Proposition 4.8, another application of Full Pareto implies that the individual preorder satisfies ($I_3^{\mathbb{Q}}$), and therefore (Lemma 4.9) has a Rational Vector EU representation. The derivation of the expected total utility representation of the social preorder then proceeds just as in Theorem 4.4.

We conclude with two further remarks about Proposition 4.8. First, the proposition lends some credence to our suggestion that Full Pareto, (P_3), is the right way of extending the usual strong Pareto condition (P_2) to say something 'Pareto-style' about incomparability. For the question of whether (P_3) is plausible, the crucial issue is the status of its component (P_c). Suppose first that $\mathbb{K}$ in the statement of (P_c) is empty. Then (P_c) is entailed by the conjunction of (P_1) and the following plausible principle (in, for concreteness, the variable-population framework): $P \curlywedge_{\mathbb{P}^{\mathrm{v}}} P' \implies \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P) \curlywedge^{\mathrm{v}} \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P')$. In the general case where $\mathbb{K}$ can be non-empty, (P_c) is then motivated by the kind of separability principle which underlies (P_b), that of ignoring groups of indifferent individuals. To this we now add that, since (P_3) is essentially equivalent to (I_3), given our axioms for aggregation, and since (I_3), as strong independence, is so well-established, (P_3) appears to be a very natural extension of (P_2).

Second, our aggregation Theorems 2.2 and 3.5 are compatible with the adoption of any non-expected utility theory for the individual preorder, provided only that Omega Independence is satisfied in the variable population case. This allows non-expected utility theory to be easily inserted into our approach to aggregation. But Proposition 4.8 reveals a potential cost. Non-expected utility theories typically reject every independence axiom. But given the assumptions of Theorem 3.5, rejecting any independence axiom requires rejecting the corresponding Pareto axiom. To its critics, this may be a further strike against non-expected utility theory; to its defenders, it may be evidence for a hidden problem with Pareto. We briefly address the options for someone with broadly utilitarian sympathies who wishes to adopt a non-expected utility theory without giving up Pareto in Section 6.6.

5. Non-expected utility

In this section we continue to explore the relationship between individual preorders and the social preorders they generate, but we now focus on non-expected utility theory. Although independence remains very popular as a normative principle, it continues to have its critics; see, for example, Buchak (2013). It is therefore natural to ask what typical non-expected utility conditions on the individual preorder imply about the social preorder, and vice versa.

Even if one accepts independence at the normative level, it is hard to ignore its widespread violation at the empirical level (see note 5), and the project we pursue here may have some relevance to empirical work. The literature has mostly focused on subjects who violate independence when only self-interest is at play. But such subjects may on occasion put themselves in the position of the social planner to make judgments about social outcomes. It is natural to ask whether their views about risk at the individual level are reflected in their views about welfare distributions, even in risk-free cases. Answering this first requires models of what independence-violating judgments about risk imply about social evaluation; that is what our aggregation theorems provide. We do not pursue this empirical angle here, but see Example 2.9 and Section 6.1 for discussion relevant to the natural idea that there is a connection between non-expected utility and egalitarian attitudes.

In what follows, we discuss two standard approaches to non-expected utility theory. The upshot is that ideas from non-expected utility theory provide two conceptually distinct paths from our aggregation theorems to something at least close to Harsanyi-style utilitarianism. First, Theorems 5.2 and 5.3 show that assuming monotonicity for the social preorder, along with some common background assumptions, is enough to guarantee that the social preorder is represented by expected total utility. Second, even if we deny monotonicity, Theorems 5.5 and 5.7 show that when the individual preorder has a 'local expected utility' representation in the style of Machina (1982), the social preorder has a 'local expected total utility' representation.

5.1. Axioms

One strand of non-expected utility theory has been to articulate axioms which mildly weaken independence in natural ways. Some non-expected utility axioms are straightforwardly inherited by the social preorder in both the constant and variable population cases. These include Betweenness, Quasiconcavity, Quasiconvexity, Very Weak Substitution, and Mixture Symmetry. In addition, Weak Substitution and Ratio Substitution are inherited in at least the constant population case.[64] These results follow easily from the fact that the map $L \mapsto p_L$ (or $L \mapsto p^{\mathbb{I}}_L$) is mixture preserving.

These conditions are typically combined with continuity and completeness in the non-expected utility literature, but there is work aimed at allowing for failures of each of those conditions. Just to give one example, Karni and Zhou (2019) propose an axiom they call Partial Substitution, a condition which relaxes Weak Substitution to accommodate incompleteness. At least in the constant population case, this is also inherited by the social preorder.

Inheritance of other non-expected utility axioms is less straightforward, as they are designed only for the case in which the set of outcomes is a compact interval of real numbers (Schmidt, 2004). Thus even if we assumed $\mathbb{W}$ was such an interval, the axioms would not make sense for $\mathbb{D} = \mathbb{W}^{\mathbb{I}}$. (And even

[64] For definitions and sources of these axioms see e.g. Schmidt (2004).

when the axioms make sense, representation theorems designed for an interval of outcomes may not apply.) That problem aside, the ease with which inheritance can be shown for the axioms so far discussed might lead one to guess that inheritance is the rule. Nevertheless, some important non-expected utility axioms are not inherited.

Suppose in general that $X = \mathcal{P}(Y)$ is a convex set of probability measures on a measurable space Y, and that X includes the delta-measure 1_y for every $y \in Y$. Suppose that a preorder $\succsim_X$ on $\mathcal{P}(Y)$ is *upper-measurable*, meaning that $U_y := \{z \in Y : 1_z \succsim_X 1_y\}$ is measurable for every $y \in Y$. Define a preorder $\succsim_X^{\mathrm{SD}}$ on $\mathcal{P}(Y)$ by $p \succsim_X^{\mathrm{SD}} q \iff p(U_y) \geq q(U_y)$ for all $y \in Y$. We say that p *stochastically dominates* q when $p \succsim_X^{\mathrm{SD}} q$. Consider the following axiom, which requires consistency with stochastic dominance.

Monotonicity (M). For an upper-measurable preorder $\succsim_X$,

(i) $p \sim_X^{\mathrm{SD}} q \implies p \sim_X q$; and
(ii) $p \succ_X^{\mathrm{SD}} q \implies p \succ_X q$.

This axiom is widely assumed in non-expected utility theory. But the next example shows that the social preorder does not always inherit (M) from the individual preorder, even in the constant population case.

Example 5.1. Make the assumptions of Example 2.9, where the individual preorder had a rank-dependent utility representation. Again make the concrete assumption that $r(x) = x^2$ and $u(x) = x$; equip $\mathbb{W} = \mathbb{R}$ and $\mathbb{D} = \mathbb{R}^n$ with the Borel sigma algebras. Assume a population of $n = 2$ people. Then $\succsim$ ranks a distribution $d = (w_1, w_2)$ with welfare states $w_1 \leq w_2$ according to the aggregate score $V(d) = \frac{3}{4}w_1 + \frac{1}{4}w_2$. Both $\succsim_{\mathbb{P}}$ and $\succsim$ are upper measurable, and $\succsim_{\mathbb{P}}$ satisfies (M). Consider three distributions $d_1 = (0, 0)$, $d_2 = (-1, 3)$ and $d_3 = (-2, 6)$. Then $1_{d_1} \sim 1_{d_2} \sim 1_{d_3}$, so that $1_{d_1} \sim^{\mathrm{SD}} L := \frac{1}{2}1_{d_2} + \frac{1}{2}1_{d_3}$. But $U(p_{1_{d_1}}) = 0$ and $U(p_L) = -\frac{1}{4}$, hence $1_{d_1} \succ L$, violating (M)(i). For a violation of (M)(ii), let $d_4 = (-\frac{1}{8}, -\frac{1}{8})$. Then $L \succ^{\mathrm{SD}} 1_{d_4}$ but $1_{d_4} \succ L$.

This example reveals tension in a common line of thought. For, in some variant, (M) has been seen as '[t]he most widely acknowledged principle of rational behavior under risk' (Schmidt, 2004, p. 19). But it is also sometimes said that rationality requires applying to the social preorder whatever conditions one imposes on the individual preorder (compare Harsanyi 1977a, p. 637).

One response would be insist that (M) does apply to the social preorder, and say so much the worse for non-expected utility theories that are forced to reject it there. But the following result suggests that this response places very strong restrictions on non-expected utility theories: indeed, given common background assumptions, (M) for the social preorder is *equivalent* to its having an EU representation (in which case it has an expected total utility form, by Theorem 4.4). We state the variable population version result first, and then note a version for a family of constant population models below.

Theorem 5.2. *Suppose that $\succsim^{\mathrm{v}}$ is upper-measurable and generated by $\succsim_{\mathbb{P}^{\mathrm{v}}}$. Suppose, moreover, that $\mathbb{D}^{\mathrm{v}}$ contains every possible distribution with finitely many people, i.e. $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}} = (\mathbb{W}^{\mathrm{v}})^{\mathbb{I}}$ for each finite $\mathbb{I} \subset \mathbb{I}^{\infty}$; that $\mathbb{P}^{\mathrm{v}}$ and each $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ consists of all finitely supported probability measures on $\mathbb{W}^{\mathrm{v}}$ and $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ respectively; and that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is complete and strongly continuous.*[65] *Then*

(i) The social preorder $\succsim^{\mathrm{v}}$ satisfies (M) *if and only if the individual preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies EUT.*
(ii) If $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has an EU representation U^{v}, normalized so that $U^{\mathrm{v}}(1_{\Omega}) = 0$, then $\succsim^{\mathrm{v}}$ has an EU representation $V^{\mathrm{v}} = \sum_{i \in \mathbb{I}^{\infty}} U^{\mathrm{v}} \circ \mathcal{P}_i^{\mathrm{v}}$.
(iii) In particular, if $\succsim^{\mathrm{v}}$ satisfies (M)*, then $\succsim_{\mathbb{P}^{\mathrm{v}}}$ and $\succsim^{\mathrm{v}}$ satisfy* (I_3)*,* (P_3) *and (if restrictions exist)* (S_3)*.*

Thus in this relatively simple setting, we obtain a total expected utility (or expected total utility) representation of the social preorder from our axioms for aggregation merely by assuming completeness and strong continuity for the individual preorder and monotonicity for the social preorder; such properties as strong independence (I_3), strong separability (S_3), and Full Pareto (P_3) are derived, not assumed. It is worth noting the analogue of Theorem 5.2 for families of constant population models (in the sense of Section 4.3). The proof is exactly the same, with constant population objects substituted for variable population ones. But the result is independently interesting because the hypotheses of completeness and strong continuity may both be more compelling when we exclude Ω from the set of welfare states.

Theorem 5.3. *Suppose that every social preorder $\succsim_{\mathbb{I}}$ in a family of constant population models is upper-measurable and generated by $\succsim_{\mathbb{P}}$. Suppose, moreover, that $\mathbb{D}_{\mathbb{I}} = \mathbb{W}^{\mathbb{I}}$ for each finite $\mathbb{I} \subset \mathbb{I}^{\infty}$; that $\mathbb{P}$ and each $\mathbb{L}_{\mathbb{I}}$ consists of all finitely supported probability measures on $\mathbb{W}$ and $\mathbb{D}_{\mathbb{I}}$ respectively; and that $\succsim_{\mathbb{P}}$ is complete and strongly continuous. Then*

(i) Every $\succsim_{\mathbb{I}}$ satisfies (M) *if and only if $\succsim_{\mathbb{P}}$ satisfies EUT.*
(ii) If $\succsim_{\mathbb{P}}$ has an EU representation U, then each $\succsim_{\mathbb{I}}$ has an EU representation $V = \sum_{i \in \mathbb{I}} U \circ \mathcal{P}_i$.
(iii) In particular, if every $\succsim_{\mathbb{I}}$ satisfies (M)*, then $\succsim_{\mathbb{P}}$ and every $\succsim_{\mathbb{I}}$ satisfy* (I_3)*,* (P_3) *and (if restrictions exist)* (S_3)*.*

These results make it seem unpromising (although perhaps not impossible) to pursue non-expected utility theory for social preorders based on (M). Of course, even if we maintain (M) for the individual preorder, denying it for social preorders also has its costs, not least that it rules out the application of representation theorems for social preorders that take (M) as a premise. We therefore now turn to a kind of non-expected utility representation that does not depend on (M) and which can apply to individual and social preorders alike.

5.2. Local expected utility

The axiomatic approach to non-expected utility theory tries to respect the normative plausibility of independence by focusing on axioms that weaken it only mildly. An alternative approach, pioneered by Machina (1982), abandons independence entirely while imposing technical conditions on preorders that are just strong enough to allow one to apply expected utility techniques locally in order to deduce important global properties. A number of technical conditions have been considered; we focus on one that allows us to elaborate on the utilitarian nature of our social preorders. We begin by explicating a sense, weaker than Machina's, in which a preorder of probability measures can be locally governed by expected utility.

Let $X = \mathcal{P}(Y)$ be a convex set of probability measures on Y. Recall from Section 4.2 that a function $f : Y \to \mathbb{R}$ is $\mathcal{P}(Y)$-*integrable* if it is Lebesgue integrable with respect to every element of $\mathcal{P}(Y)$; and that a function $U : \mathcal{P}(Y) \to \mathbb{R}$ is *expectational* if there is a $\mathcal{P}(Y)$-integrable function u such that, for any $q \in \mathcal{P}(Y)$, $U(q) = \int_Y u \, \mathrm{d}q$. Now, for any basepoint $p \in \mathcal{P}(Y)$, we can rewrite this as $U(p + t(q-p)) = \int_Y u \, \mathrm{d}(p + t(q-p))$ for all $q \in \mathcal{P}(Y)$ and $t \in [0, 1]$.

65 Say that a sequence (p_n) in a space $X = \mathcal{P}(Y)$ of probability measures *converges strongly* to $p \in X$ (written $p_n \xrightarrow{s} p$) whenever $p_n(A) \to p(A)$ for all measurable A in Y. A preorder $\succsim_X$ on X is *strongly continuous* if whenever $p_n \xrightarrow{s} p$, (i) $p_n \succsim_X q$ for all $n \implies p \succsim_X q$; and (ii) $q \succsim_X p_n$ for all $n \implies q \succsim_X p$. This is, of course, an instance of the continuity axiom (Cont): the topology is the one whose closed sets are precisely the subsets that contain the limit points of their strongly convergent sequences.

It is natural to say that U is a *locally expectational at p* if there is a function u_p satisfying this equation up to first order in t. To be precise, $U: \mathcal{P}(Y) \to \mathbb{R}$ is locally expectational at $p \in \mathcal{P}(Y)$ if there is a $\mathcal{P}(Y)$-integrable function u_p such that, for each[66] $q \in \mathcal{P}(Y)$,

$$U(p + t(q-p)) = \int_Y u_p \, \mathrm{d}(p + t(q-p)) + o(t) \quad \text{as } t \to 0^+. \tag{5}$$

Call such a u_p a *local utility function for U at p*. We say U is locally expectational on a subset $S \subset \mathcal{P}(Y)$ to mean that it is locally expectational at every $p \in S$.

Consider the following condition on a preorder $\succsim_X$ on $\mathcal{P}(Y)$. For a reason we will soon explain (see note 74), we state it relative to a subset $S \subset \mathcal{P}(Y)$; if not explicitly mentioned, $S = \mathcal{P}(Y)$.

Local EUT over *S*. There is a function $U: \mathcal{P}(Y) \to \mathbb{R}$ that represents $\succsim_X$ and that is locally expectational on S. We say that U is a *Local EU* representation of $\succsim_X$ over S.

While there are normatively natural axiomatizations of EUT, Multi EUT, and Vector EUT, the normative significance of Local EUT can be understood via a differentiability concept. A function $U: \mathcal{P}(Y) \to \mathbb{R}$ is said to be *Gâteaux differentiable at* $p \in \mathcal{P}(Y)$ if the one-sided limit

$$U'_p(q-p) := \lim_{t \to 0^+} \frac{U(p + t(q-p)) - U(p)}{t} \tag{6}$$

exists for all $q \in \mathcal{P}(Y)$.[67] Thus $U'_p(q-p)$ is a directional derivative of U at p in the direction $q - p$. Say that U is *integrally Gâteaux differentiable at* $p \in \mathcal{P}(Y)$ when it is Gâteaux differentiable at p and there exists a $\mathcal{P}(Y)$-integrable $u_p: Y \to \mathbb{R}$ such that

$$U'_p(q-p) = \int_Y u_p \, \mathrm{d}(q-p) \tag{7}$$

for all $q \in \mathcal{P}(Y)$.[68] Let ∇U_p be the set of such u_p; thus U is integrally Gâteaux differentiable at p if and only if $\nabla U_p \neq \varnothing$. It is well known that many normatively natural conditions on $\succsim_X$ are compatible with, and sometimes guarantee,[69] the assumption that $\succsim_X$ can be represented by an integrally Gâteaux differentiable function. But then $\succsim_X$ must satisfy Local EUT:

Lemma 5.4. *Suppose $\mathcal{P}(Y)$ is a convex set of probability measures on a measurable space Y. Then $U: \mathcal{P}(Y) \to \mathbb{R}$ is locally expectational at $p \in \mathcal{P}(Y)$ if and only if it is integrally Gâteaux differentiable at p. Specifically, the local utility functions for U at p are precisely those $u_p \in \nabla U_p$ such that $U(p) = \int_Y u_p \, \mathrm{d}p$.*

In parallel to the constant population claims of Theorem 4.4 we have

Theorem 5.5 (*Local EUT inheritance: constant population*)**.** *Suppose $\succsim_{\mathbb{P}}$ generates $\succsim$.*

(i) $\succsim_{\mathbb{P}}$ satisfies Local EUT if and only if $\succsim$ does.

(ii) In particular, if $\succsim_{\mathbb{P}}$ has a Local EU representation U, then $\succsim$ has a Local EU representation[70]

$$V(L) := \#\mathbb{I}U(p_L).$$

(iii) If u_L is a local utility function for U at p_L, then $\sum_{i \in \mathbb{I}} u_L \circ \mathcal{W}_i$ is a local utility function for V at L.

This result has two significant implications. First, the often justifiable assumption of Local EUT for the individual preorder guarantees that Local EUT techniques and results can be applied to the social preorder as well.[71] Second, in Theorem 4.4, we saw that if the individual preorder has an expected utility representation, then the social preorder has a representation by expected total utility. Correspondingly, the last part of Theorem 5.5 shows that if the individual preorder has a local expected utility representation, then the social preorder has what we can analogously call a *local expected total utility* representation.[72] This local version of Theorem 4.4(ii) bolsters the view that the social preorders described by our aggregation theorems are utilitarian in spirit even when they do not satisfy (I_3).[73]

We would like to extend Theorem 5.5 to the variable population case. As usual (see Section 3.1), there is no problem in doing so for each finite population $\mathbb{I} \subset \mathbb{I}^\infty$: for any Local EU representation U^{v} of $\succsim_{\mathbb{P}^{\mathsf{v}}}$, and any $L \in \mathbb{L}^{\mathsf{v}}_{\mathbb{I}}$, we can define $V^{\mathsf{v}}(L) = \#\mathbb{I}U^{\mathsf{v}}(p^{\mathbb{I}}_L)$, in parallel to Theorem 5.5(ii). This will be a Local EU representation of the restriction of $\succsim^{\mathsf{v}}$ to $\mathbb{L}^{\mathsf{v}}_{\mathbb{I}}$, with a local expected total utility interpretation as in Theorem 5.5(iii). The proofs are the same as in the constant population case. The only difficulty is that this $V^{\mathsf{v}}(L)$ is not a function of L independent of $\mathbb{I}$, so does not define a representation of the unrestricted social preorder.

To avoid this difficulty, we will focus on a narrower class of variable population social preorders, for which an unrestricted representation V^{v} is readily defined. We first explain why this class is still generated by a rich and normatively interesting set of individual preorders. Say that a function $U^{\mathsf{v}}: \mathbb{P}^{\mathsf{v}} \to \mathbb{R}$ is *Omega-linear* if for all $P \in \mathbb{P}^{\mathsf{v}}$ and $\alpha \in [0, 1]$,

$$U^{\mathsf{v}}(\alpha P + (1-\alpha)1_\Omega) = \alpha U^{\mathsf{v}}(P) + (1-\alpha)U^{\mathsf{v}}(1_\Omega). \tag{8}$$

Similarly, say that a function $V^{\mathsf{v}}: \mathbb{L} \to \mathbb{R}$ is *d_Ω-linear* if for all $L \in \mathbb{L}^{\mathsf{v}}$ and $\alpha \in [0, 1]$,

$$V^{\mathsf{v}}(\alpha L + (1-\alpha)1_{d_\Omega}) = \alpha V^{\mathsf{v}}(L) + (1-\alpha)V^{\mathsf{v}}(1_{d_\Omega}).$$

Recall from Lemma 3.3 that d_Ω here is the empty distribution.

Lemma 5.6. *Suppose $\mathbb{P}^{\mathsf{v}}$ extends $\mathbb{P}$ (see Section 3.4). Let $U: \mathbb{P} \to \mathbb{R}$.*

(i) For any $c \in \mathbb{R}$, U has a unique Omega-linear extension $U^{\mathsf{v}}: \mathbb{P}^{\mathsf{v}} \to \mathbb{R}$ that satisfies $U^{\mathsf{v}}(1_\Omega) = c$.

66 Recall that the expression '$f(t) = g(t) + o(t)$ as $t \to 0^+$' means $\lim_{t \to 0^+} \frac{f(t)-g(t)}{t} = 0$.

67 Our notion of Gâteaux differentiability is very weak, as it only requires a one-sided limit, only considers $q \in \mathcal{P}(Y)$, and does not make any topological assumptions. It coincides with what is sometimes known as semi-differentiability.

68 Similar definitions, but with more restrictions on u_p or Y, are found in Chew et al. (1987), Chew and Mao (1995), and Cerreia-Vioglio et al. (2017). For example, the latter, from whom we borrow notation, assume u_p is continuous and bounded, but we make no such assumption.

69 See Chew and Mao (1995) for a summary when Y is a real interval and $\mathcal{P}(Y)$ is the set of Borel probability measures.

70 The $\#\mathbb{I}$ in the definition of V is optional here, but it facilitates comparison with the variable population analogue Eq. (9) in which the corresponding factor is necessary.

71 See Cerreia-Vioglio et al. (2017) for extensive discussion of the global properties of preorders that satisfy (in our terminology) Local EUT in terms of their local utility functions, along with detailed applications.

72 A slightly different notion of local utilitarianism was discussed by Machina (1982, §5.2). His notion applies to social welfare functions on 'wealth distributions', which he idealizes as probability measures on $\mathbb{W}$.

73 It is worth noting that the individual preorder in Example 5.1 satisfies Local EUT, even though the social preorder does not satisfy (M). Thus Theorem 5.5 applies, even if nothing like Theorem 5.3 does. This illustrates the generality of the local expected utility-based methods. It also illustrates the strategy of using a standard non-expected utility theory (here, rank-dependent utility theory) for the individual preorder to derive Local EU representations of both the individual and social preorders. It would work less well to start from a standard non-expected utility theory for the social preorder, since such theories invariably assume (M), and, as we noted in Section 5.1, (M) for the social preorder is a stringent assumption.

(ii) If U is locally expectational on $\mathbb{P}$, such a U^{v} is locally expectational on $\mathbb{P}^{\mathsf{v}} \setminus \{1_\Omega\}$.

We have noted that there is a normatively interesting class of constant population individual preorders that satisfy Local EUT. Any such $\succsim_{\mathbb{P}}$ has a locally expectational representation U. By the lemma, U has an Omega-linear extension U^{v} which is locally expectational on $\mathbb{P}^{\mathsf{v}} \setminus \{1_\Omega\}$ with a free choice of the value of $U^{\mathsf{v}}(1_\Omega)$, implying flexibility in how nonexistence is compared with other welfare states. This U^{v} represents an Omega independent $\succsim_{\mathbb{P}^{\mathsf{v}}}$ that extends $\succsim$. Thus the individual preorders to which the next result applies form a rich class.[74]

Theorem 5.7 (*Local EUT Inheritance: Variable Population*)**.** *Suppose $\succsim_{\mathbb{P}^{\mathsf{v}}}$ generates $\succsim^{\mathsf{v}}$. Assume that the sigma algebra on $\mathbb{D}^{\mathsf{v}}$ is coherent.*[75]

(i) $\succsim_{\mathbb{P}^{\mathsf{v}}}$ satisfies Local EUT over $\mathbb{P}^{\mathsf{v}} \setminus \{1_\Omega\}$ with respect to an Omega-linear representation if and only if $\succsim^{\mathsf{v}}$ satisfies Local EUT over $\mathbb{L}^{\mathsf{v}} \setminus \{1_{d_\Omega}\}$ with respect to an d_Ω-linear representation.

(ii) In particular, if $\succsim_{\mathbb{P}^{\mathsf{v}}}$ has an Omega-linear representation U^{v} that is Local EU over $\mathbb{P}^{\mathsf{v}} \setminus \{1_\Omega\}$, then $\succsim_{\mathbb{P}^{\mathsf{v}}}$ has an d_Ω-linear representation V^{v} that is Local EU over $\mathbb{L}^{\mathsf{v}} \setminus \{1_{d_\Omega}\}$, defined for $L \in \mathbb{L}^{\mathsf{v}}_{\mathbb{I}}$ by[76]

$$V^{\mathsf{v}}(L) := \#\mathbb{I}U^{\mathsf{v}}(p_L^{\mathbb{I}}) - \#\mathbb{I}U^{\mathsf{v}}(1_\Omega). \tag{9}$$

(iii) If U^{v} is normalized so that $U^{\mathsf{v}}(1_\Omega) = 0$, then, for any $L \in \mathbb{L}^{\mathsf{v}}_{\mathbb{I}} \setminus \{1_{d_\Omega}\}$, if u_L is a local utility function for U^{v} at $p_L^{\mathbb{I}}$, $\sum_{i \in \mathbb{I}^\infty} u_L \circ \mathcal{W}_i^{\mathsf{v}}$ is a local utility function for V^{v} at L.

The normalization condition in (iii) has the usual significance: given one Omega-linear representation of $\succsim_{\mathbb{P}^{\mathsf{v}}}$ that is locally expectational on $\mathbb{P}^{\mathsf{v}} \setminus \{1_\Omega\}$, we can obtain another by adding a constant. Aside from the requirement of Omega-linearity, the implications of Theorem 5.7 parallel those of Theorem 5.5. In particular, a wide range of variable population social preorders are compatible with constant population individual preorders that satisfy normatively interesting non-expected utility conditions, and except perhaps at one point, these social preorders have local expected total utility representations.

6. Comparisons

We now relate our aggregation theorems to several standard topics: egalitarianism; the *ex ante* versus *ex post* distinction; utilitarianism; and Harsanyi's impartial observer theorem. We end with discussion of related literature.

6.1. Quasi utilitarianism and egalitarianism

Recall that we have defined quasi utilitarian social preorders to be precisely the social preorders that are compatible with our aggregation theorems (Definitions 2.3 and 3.6). We will defend this terminology in Section 6.3. But we know of no discussion of quasi utilitarian preorders, so our goal in this subsection is to discuss their properties in more detail, especially by contrasting them with egalitarian social preorders. First we show that they form a rich class: they are compatible with a wide variety of social preorders on distributions.

To simplify the discussion, let us assume that $\mathbb{D}^{\mathsf{v}}$ is the set of all possible distributions with finite populations, and that $\mathbb{D}$ is the set of all possible distributions with some given constant population. We also assume that the sigma algebras on $\mathbb{D}^{\mathsf{v}}$ and $\mathbb{D}$ separate points,[77] so that if $d \neq d'$ then $1_d \neq 1_{d'}$. We can therefore think of distributions as delta-measures. Finally, we assume that $\mathbb{L}^{\mathsf{v}}$ and $\mathbb{L}$ contain 1_d for any d in $\mathbb{D}^{\mathsf{v}}$ and $\mathbb{D}$ respectively.

Say that a preorder $\succsim_0^{\mathsf{v}}$ on $\mathbb{D}^{\mathsf{v}}$ is consistent with quasi utilitarianism if there exists some quasi utilitarian preorder $\succsim^{\mathsf{v}}$ on $\mathbb{L}^{\mathsf{v}}$ such that for all $d, d' \in \mathbb{D}^{\mathsf{v}}$, $d \succsim_0^{\mathsf{v}} d' \iff 1_d \succsim^{\mathsf{v}} 1_{d'}$. We can similarly ask whether a preorder $\succsim_0$ on $\mathbb{D}$ is consistent with quasi utilitarianism for the given finite population: whether $d \succsim_0 d' \iff 1_d \succsim 1_{d'}$ for all $d, d' \in \mathbb{D}$. Discussions of distributive views like utilitarianism or egalitarianism often focus solely on risk-free cases, so it is natural to ask which preorders on distributions are consistent with quasi utilitarianism. We answer this question in terms of the following two conditions.

Risk-Free Anonymity Given $d \in \mathbb{D}$ and $\sigma \in \Sigma$, we have $d \sim_0 \sigma d$.

Say that $c \in \mathbb{D}^{\mathsf{v}}$ is an *m-scaling* of $d \in \mathbb{D}^{\mathsf{v}}$ if it consists of 'm copies' of d—that is, there is an m-to-1 map s of $\mathbb{I}^\infty$ onto itself such that $\mathcal{W}_i^{\mathsf{v}}(c) = \mathcal{W}_{s(i)}^{\mathsf{v}}(d)$ for every individual i. For example, $(x, x, y, y, \Omega, \Omega, \dots)$ is a 2-scaling of $(x, y, \Omega, \dots)$.

Scale Invariance If, for some $m > 0$, $c, c' \in \mathbb{D}^{\mathsf{v}}$ are m-scalings of $d, d' \in \mathbb{D}^{\mathsf{v}}$ respectively, then $c \succsim_0^{\mathsf{v}} c' \iff d \succsim_0^{\mathsf{v}} d'$.

Risk-Free Anonymity is obviously a very weak and uncontroversial constraint, while Scale Invariance is a seemingly modest generalization (they are equivalent when $m = 1$). But these are the only constraints imposed by consistency with quasi utilitarianism.

Proposition 6.1. *Under the assumptions just made:*

(i) A preorder on $\mathbb{D}$ is consistent with constant population quasi utilitarianism if and only if it satisfies Risk-Free Anonymity.

(ii) A preorder on $\mathbb{D}^{\mathsf{v}}$ is consistent with variable population quasi utilitarianism if and only if it satisfies Scale Invariance.

This result shows that many apparently egalitarian (as well as inegalitarian) preorders of distributions are consistent with quasi utilitarianism. For example, as preorders of distributions, both maximin and leximin are consistent with quasi utilitarianism. This raises questions about the significance of quasi utilitarian preorders as a class. In particular, why do they merit the 'quasi utilitarian' name, if they include preorders with apparently egalitarian properties?

The answer is that the axioms of our aggregation theorems precisely rule out certain features of the social preorder that may be considered essential to standard egalitarian concerns, but which can only be seen once one introduces risk.[78] Thus, even if some quasi utilitarian preorders are egalitarian in *some* useful sense, quasi utilitarianism excludes what are arguably defining characteristics of egalitarianism.

To illustrate, suppose given welfare states x and z with $x \succ_{\mathbb{P}} z$, and a population consisting of Ann and Bob. Consider the

[74] It is easy to check from the definitions that, if U^{v} is Omega-linear, then the Gâteaux derivative $(U^{\mathsf{v}})'_{1_\Omega}(P - 1_\Omega) = U^{\mathsf{v}}(P) - U^{\mathsf{v}}(1_\Omega)$. It follows from (7) and Lemma 5.4 that if u is a local utility function for U^{v} at 1_Ω, then U^{v} is an EUT representation, equal to the expectation of u. This is why we do not insist on Local EUT over $\mathbb{P}^{\mathsf{v}}$ in the next theorem.

[75] We defined 'coherent' just before Proposition 3.7. Just as in that proposition, it would suffice to assume that $\mathbb{I}^\infty$ is countable, although we do not pursue this here.

[76] Since U^{v} is Omega-linear, it is easy to check that $V^{\mathsf{v}}(L)$ is a well-defined function of $L \in \mathbb{L}^{\mathsf{v}}$, independent of $\mathbb{I}$.

[77] This is weaker than the assumption that singletons are measurable.

[78] It is not obvious that egalitarians should accept scale invariance, since it is not clear how the value of equality scales with population size. But the tensions between our axioms for aggregation and versions of egalitarianism that we now discuss are more direct, and apply even in the constant population case.

following lotteries (in which each column displays a distribution with a 1/2 chance of occurring).

L_E	$\frac{1}{2}$	$\frac{1}{2}$
Ann	x	z
Bob	x	z

L_F	$\frac{1}{2}$	$\frac{1}{2}$
Ann	x	z
Bob	z	x

L_U	$\frac{1}{2}$	$\frac{1}{2}$
Ann	x	x
Bob	z	z

It is arguable that L_E is socially better than L_F on the grounds that while Ann and Bob face identical prospects (and therefore *ex ante* equality obtains) in both, L_E ensures *ex post* equality (Myerson, 1981). It is also arguable that L_F is better than L_U on the grounds that while there is nothing to chose between their outcomes, under L_F there is *ex ante* equality, so L_F is apparently fairer (Diamond, 1967).

In our view, suitable generalizations of '$L_E \succ L_F$' and '$L_F \succ L_U$' are essential to *ex post* and *ex ante* egalitarianism respectively.[79] In related work, we use this idea to develop accounts of *ex ante* and *ex post* egalitarianism that are compatible with any individual preorder. This leads to a taxonomy of the main distributive views in which quasi utilitarianism is distinguished from the other views by containing at its core the indifference to equality expressed by '$L_E \sim L_F \sim L_U$'. Quasi-utilitarianism is inconsistent with *ex ante* egalitarianism because it accepts Two-Stage Anonymity, and is inconsistent with *ex post* egalitarianism because it accepts Anteriority. Thus despite the egalitarian appearance of some quasi utilitarian preorders of distributions, there is a sharp distinction between quasi utilitarianism and standard forms of egalitarianism.

Because it accepts Reduction to Prospects, quasi utilitarianism also conflicts with a third type of view which has often been associated with egalitarianism. This is the distributive view widely discussed in moral philosophy known as 'prioritarianism' or 'the priority view' (see Tungodden 2003 for discussion). The core idea of prioritarianism is arguably a form of social risk aversion. For example, suppose that welfare state x is better for an individual than welfare state y, which is in turn better than z, and that the prospect of equal chances of x and z is exactly as good as getting welfare state y for sure. A prioritarian social preorder will consider the one-person distribution (y) to be strictly better than the one-person lottery that gives equal chances to the distributions (x) and (z). In Section 3.2 we stipulated that the individual preorder satisfies **(E)**; granted that, prioritarianism rejects **(F)**, and more generally, Reduction to Prospects, serving to distinguish quasi utilitarianism from prioritarianism.[80]

6.2. *Ex ante and ex post*

We now explain why there is a natural sense in which quasi utilitarian preorders are those social preorders that are weakly *ex ante* and anonymously *ex post*. This generalizes the contrast between quasi utilitarianism and *ex post* and *ex ante* egalitarianism developed in 6.1. We focus on the constant population case, the variable case being parallel. The Pareto conditions discussed below (and introduced in Section 4.3) are therefore understood relative to a fixed population. As throughout, we assume probabilities are given, and thus do not address problems that arise from defining *ex ante* principles in terms of possibly differing individual subjective probabilities.[81]

[79] For similar views, see Broome (1989, 1991), Ben-Porath et al. (1977), Fleurbaey (2010), Grant et al. (2012b), Saito (2013), and McCarthy (2015) among others. For surveys on *ex ante* and *ex post* egalitarianism see Mongin and Pivato (2016, §§25–26) and Fleurbaey (2018).

[80] Nevertheless, our aggregation theorems are still useful for the characterization of prioritarianism; see McCarthy (2017).

[81] For an entry into the huge literature on this topic, see Mongin and Pivato (2019); see also Section 6.7.4.

6.2.1. *Ex ante*

Let (Ant) and (RP) stand for Anteriority and Reduction to Prospects. The following irreversible implications are obtained by noting that (RP) is equivalent to the restriction of (P_3) to lotteries in $\mathcal{L}(\mathbb{P})$.

$$(RP) \Leftarrow (P_3) \Rightarrow (P_2) \Rightarrow (P_1) \Rightarrow (Ant)$$

Social preorders are commonly said to be *ex ante* if, in some sense, they respect unanimous 'before the event' judgments of individual welfare. Each of the above principles expresses some way of making this rough idea precise, which helps explain why '*ex ante*' is used quite flexibly. But the most popular interpretation sees social preorders as *ex ante* if they satisfy strong Pareto (P_2) (Mongin and d'Aspremont, 1998, §5.4). This corresponds to a relatively strong notion of unanimity: respect the unanimous judgments of non-indifferent individuals. But this notion of unanimity is more fully captured by (P_3), which strengthens (P_2) in cases where the individual preorder is incomplete. We therefore suggest that it is social preorders which satisfy (P_3) which should be seen as *ex ante*.

However, as long as the social preorder is impartial in the sense expressed by Two-Stage Anonymity, requiring it also to be *ex ante* in the sense of (P_3) carries an implicit commitment: it effectively means that the individual preorder has to satisfy strong independence (see Theorem 4.10). But that rules out a wide range of possibilities for individual welfare comparisons, so it is natural to ask which principle expresses the *ex ante* idea as strongly as possible while remaining neutral on the properties of the individual preorder. According to our aggregation Theorem 2.2, that principle is the conjunction of (Ant) and (RP). We will therefore say that social preorders satisfying that conjunct are *weakly ex ante*.

Similarly, in the variable population case, we will say that a social preorder satisfying (P_3) is *ex ante*, and one satisfying Anteriority and Reduction to Prospects is weakly *ex ante*.

6.2.2. *Ex post*

Social preorders are often said to be *ex post* when they satisfy expected utility theory (Mongin and d'Aspremont, 1998, §5.4). But this seems distant from the ordinary meaning of the term, which suggests that lotteries should be socially evaluated from some sort of 'after the event' perspective in which all risk has resolved. In particular, if two lotteries are in some natural sense equivalent from that perspective, then they should be ranked as equals.

To approach the matter more directly, let us suppose that $X = \mathcal{P}(Y)$ is a set of probability measures on a measurable space Y, and that the following domain condition $(*)$ is satisfied: the set $\{y\}$ is measurable for each $y \in Y$, and $\mathcal{P}(Y)$ contains each corresponding delta-measure 1_y. Say that a subset B of $\mathcal{P}(Y)$ is 'closed under indifference' if $y \in B$ and $1_y \sim_X 1_{y'}$ entail $y' \in B$. The following condition seems to capture a general sense in which a preorder $\succsim_X$ on $\mathcal{P}(Y)$ should count as *ex post*.

Posteriority. Given $p, p' \in \mathcal{P}(Y)$, suppose that $p(B) = p'(B)$ whenever B is a measurable subset of Y that is closed under indifference. Then $p \sim_X p'$.

For example, assume that the domain condition $(*)$ just stated holds for the set of lotteries $\mathbb{L}$. Say that a 'level of social welfare' is an equivalence class of distributions under the social indifference relation $\sim$. Two lotteries are naturally said to be equivalent from an 'after the event' perspective whenever they define the same probability measure over levels of social welfare. Posteriority then says that two such lotteries are equally good.[82] For this

[82] Compare the discussion of Posterior Anonymity in Section 2.3, and especially note 18.

reason, we will say that a social preorder is *ex post* if it satisfies Posteriority.[83]

Continuing with the domain conditions, if B is a measurable subset of $\mathbb{D}$ that is closed under indifference, Risk-Free Anonymity implies that B is permutation-invariant. Hence given Risk-Free Anonymity, Posterior Anonymity emerges as a much weaker, special case of Posteriority. It is therefore natural to call social preorders satisfying Posterior Anonymity *anonymously ex post*. The same applies in the variable population case.

An appealing feature of this terminology is that anonymously *ex post* social preorders rule out *ex ante* egalitarianism, and weakly *ex ante* social preorders rule out *ex post* egalitarianism (see Section 6.1).

It should be noted that this derivation of Posterior Anonymity does not always make sense in our very general framework, as the domain condition $(*)$ is not guaranteed to hold. For example, our framework does not require that 1_d is in $\mathbb{L}$ for all d in $\mathbb{D}$. Nevertheless, Posterior Anonymity has self-standing appeal, and is always well-defined in our framework.

6.2.3. Two-stage anonymity and the aggregation theorems redux

Two-Stage Anonymity is entailed by Posterior Anonymity, and although Posterior Anonymity is our conceptually favored principle, it was simpler to work with Two-Stage Anonymity. Nevertheless, granted the harmless assumption of coherence in the variable population case, Propositions 2.4 and 3.7 show that Two-Stage Anonymity and Posterior Anonymity are equivalent given our other axioms for aggregation. Thus we can recapitulate our aggregation theorems as follows.

Proposition 6.2. *Any social preorder (constant or variable population) that is generated by a given individual preorder is the unique social preorder that is weakly* ex ante *and anonymously* ex post. *(In the variable-population case, we assume that the sigma algebra on* $\mathbb{D}^{\mathrm{v}}$ *is coherent.)*

In Section 2.3 we raised the question of how Two-Stage Anonymity is related to strong independence as a condition on the social preorder. Assuming Anonymity for this discussion, it is clear that Two-Stage Anonymity follows more specifically from the weaker 'independence of indifference' axiom (I_1) we stated in Section 4.1. As we have just noted, Two-Stage Anonymity also follows from Posteriority, and also from the closely related part (i) of monotonicity. It is therefore natural to wonder whether Two-Stage Anonymity really represents a significant weakening of these three alternative principles.

In fact, it is much weaker than any of them, even when it is combined with our other axioms for aggregation. To see this, note two points. First, our axioms for aggregation are compatible with any constant population individual preorder $\succsim_{\mathbb{P}}$ (Theorem 2.2 and Proposition 3.9(i)). Second, no matter how severely $\succsim_{\mathbb{P}}$ violates the three alternative principles, the social preorder $\succsim$ it generates will always satisfy Two-Stage Anonymity, but will violate the alternative principles in a comparably severe manner. This is because, by Reduction to Prospects, the restriction of $\succsim$ to $\mathcal{L}(\mathbb{P})$ is a copy of $\succsim_{\mathbb{P}}$.[84] Examples like 5.1 also show that, even if the individual preorder satisfies (M)(i) and, relatedly, Posteriority, the social preorder need not do so. We noted in Section 5.1 that axiomatic approaches to non-expected utility weaken independence only mildly, and invariably impose monotonicity. Since Two-Stage Anonymity allows for major violations of those principles, Two-Stage Anonymity is compatible with violations that are far more severe than any which would be taken seriously by non-expected utility theorists. Thus although it is *ex post* enough to rule out Diamond's example of *ex ante* egalitarianism, Two-Stage Anonymity is much weaker than the conjunction of Anonymity with any of the standard *ex post* principles. The same conclusion applies to Posterior Anonymity, since, as we noted above, it is essentially equivalent to Two-Stage Anonymity given our other axioms for aggregation.

6.3. Utilitarianism

The case for our 'quasi utilitarian' terminology rests on the claim that our social preorders have enough properties traditionally associated with utilitarianism to merit the name (see note 4). The principal properties here are indifference to *ex ante* and *ex post* equality, anonymity, and the positive dependence of social welfare on individual welfare given by Reduction to Prospects. The reason for the 'quasi' is that our social preorders do not always have well-behaved total utility representations,[85] nor do they necessarily satisfy separability or Pareto conditions, even in risk-free cases; these conditions might fairly be seen as important utilitarian commitments.

Setting aside one subtlety to which we return below, it is only when our preorders satisfy strong independence (I_3) that they possess the full range of properties naturally associated with utilitarianism, including Pareto (P_3) and separability (S_3). In addition, (I_3) for the individual preorder is necessary and sufficient for the social preorder to have the well-behaved expected total utility representation of Section 4.2. Thus we propose to define as *utilitarian* precisely those social preorders that are generated by an individual preorder that satisfies (I_3). The discussion of Section 5.2 suggests that a fairly general and important range of social preorders should then be deemed *locally utilitarian*.

The subtlety is that one might identify utilitarianism with a slightly broader class of social preorders by replacing (I_3) with its 'rational' version $(\mathrm{I}_3^{\mathbb{Q}})$. As noted in Section 4.3, this would allow an especially parsimonious axiomatization that directly expresses classical utilitarian ideas without any appeal to independence. The most visibly utilitarian variant (Theorem 4.10 note 63) assumes only Posterior Anonymity and the Pareto principle (P_3). But Pareto and Risk-Free Anonymity are central utilitarian ideas, while the extension of Risk-Free Anonymity to Posterior Anonymity is a modest expression of the teleological basis of classical utilitarianism (compare Section 6.2.2). This proposal would still give utilitarian social preorders a reasonably well-behaved expected total utility representation (Theorem 4.10), and would imply the separability principle (S_3). Nonetheless, as we explained after Proposition 4.8, it is so implausible to violate (I_3) while satisfying $(\mathrm{I}_3^{\mathbb{Q}})$, and (I_3) is so standard and technically convenient, that for pragmatic reasons we recommend the slightly narrower identification.

---

83 Note that, when the social preorder is upper-measurable, so that the stochastic dominance relation is defined, Posteriority is implied by the first part of monotonicity, (M)(i), and these conditions often coincide in practice. Even then, though, Posteriority is logically weaker, and gets more directly at the *ex post* idea.

84 Here is a simple example of an individual preorder that violates (I_1), (M)(i), and Posteriority in an appropriately severe way. Set $\mathbb{W} = \{x, y\}$, and let $\mathbb{P}$ contain all probability measures on $\mathbb{W}$ (with every subset of $\mathbb{W}$ being measurable). Let the individual preorder be such that $1_x \sim_{\mathbb{P}} 1_y$, with all non-trivial mixtures of 1_x and 1_y ranked equally but strictly below 1_x and 1_y. This models a pure and extreme form of uncertainty-aversion.

85 There is a superficial sense in which every quasi utilitarian social preorder has an additive representation. Sticking to the constant population case, we could take $\mathbb{V} = \mathrm{Span}(\mathbb{P})$ as our utility space, and extend $\succsim_{\mathbb{P}}$ to a preorder on $\mathbb{V}$. Then $L \mapsto p_L \in \mathbb{V}$ is a vector-valued representation of $\succsim$, and since $p_L = \sum_{i \in \mathbb{I}} \frac{1}{\#\mathbb{I}} \mathcal{P}_i(L)$, one might say it is a total utility representation. But we cannot necessarily define $\succsim_{\mathbb{V}}$ in a way that validates the natural requirement that a sum is an increasing function of its summands. All of the 'total utility' representations we study in this paper satisfy this requirement, which is obviously related to Pareto and separability.

6.4. Harsanyi's utilitarian theorem

As we defined it in Section 1, Harsanyi's (1955) utilitarian theorem assumed interpersonal comparisons and derived a real-valued, expected total utility representation of the constant population social preorder. He used premises which, translated into our framework, amount to EUT for the individual preorder, EUT for the social preorder, Anonymity, and strong Pareto (P_2).[86] In Theorem 4.4 we showed how to derive the same result simply by adding EUT for the individual preorder to our basic axioms of Anteriority, Reduction to Prospects, and Two-Stage Anonymity. The premises we use, then, are weaker than Harsanyi's: Two-Stage Anonymity is much weaker than the conjunction of social EUT and Anonymity, while Anteriority and Reduction to Prospects are together much weaker than (P_2), given that the individual preorder satisfies EUT and is therefore complete.[87]

6.5. The veil of ignorance

Harsanyi (1953) gave a different argument for utilitarianism, often known as his impartial observer theorem. The distinctive premise is that social evaluation corresponds to self-interested evaluation by someone behind a veil of ignorance, uncertain who he is.

Surprisingly, Harsanyi seems not to have thought that his utilitarian theorem would extend to variable populations. He used the veil of ignorance in that case, and argued that it leads to average rather than total utilitarianism, a claim that has often been endorsed.[88] But appeals to the veil of ignorance have been criticized in the constant population case,[89] and they are especially difficult to interpret, let alone justify, in the variable case. For example, it is unclear whether individuals behind the veil are required to be certain of their existence.

However, while we do not endorse the veil of ignorance as a basic axiom, our Theorems 2.2 and 3.5 can nevertheless be seen as supporting a quite general version of the veil as a derived principle. In the variable population case, for example, the lottery $p_L^{\mathbb{I}}$ defined in Theorem 3.5 can be interpreted as the prospect faced by an individual behind a veil, in the sense that he has an equal chance of being any member of $\mathbb{I}$ under L. So the quasi utilitarian social preorders are precisely the ones that correspond to individual preorders behind this veil.

This version of the veil is compatible with average utilitarianism, as we saw in Example 3.11; indeed, Harsanyi endorsed the kind of individual preorder used in that example (Ng, 1983). But as illustrated in Sections 3.5 and 3.6, quasi utilitarianism is compatible with many other social preorders as well, including total utilitarianism.[90]

6.6. An alternative

We now remark on an alternative way of generalizing Harsanyi's utilitarian theorem. One of our aims has been to show that a Harsanyi-like approach to utilitarianism can be maintained even if strong independence is rejected. But as we now explain, the rejection of strong independence leads to tension between two ideas that may each seem central to Harsanyi's utilitarianism.

The discussion after Theorem 4.10 shows that Two-Stage (or Posterior) Anonymity and Full Pareto together imply that the social preorder is quasi utilitarian and that both the social and individual preorders satisfy strong independence, at least in its rational-coefficient form. Rejecting strong independence for the individual preorder in any plausible way therefore means abandoning either Two-Stage Anonymity or Full Pareto.[91] However, both Two-Stage Anonymity and Full Pareto seem strongly in the spirit of Harsanyi's approach. Two-Stage Anonymity is entailed by premises Harsanyi accepts, capturing (for one thing) indifference to *ex ante* equality. While he does not consider Full Pareto, Harsanyi (1977a) regards strong Pareto as a rather obvious assumption, and in Section 4.3 we suggested that Full Pareto is the natural way to extend strong Pareto in the face of incompleteness. Thus rejecting strong independence for the individual preorder involves a commitment to rejecting at least one assumption that is arguably integral to Harsanyi's utilitarianism.[92] In this paper we have taken Two-Stage Anonymity as a basic axiom, and allowed for the rejection of Full Pareto. But it would be natural to instead explore the possibilities for a Harsanyi-like approach that retains Full Pareto while allowing for the rejection of Two-Stage Anonymity, and thereby strong independence. The task would be to look for principles that sufficiently weaken Two-Stage Anonymity while preserving impartiality and indifference to *ex ante* equality. But we leave this for another time.

6.7. Related literature

We will not repeat comparisons made with Harsanyi in sections 1, 6.4, and 6.5. Instead, we briefly relate our results to some recent developments. These can be classified according to which aspects of the framework of Harsanyi's social aggregation theorem they preserve. Recall from Section 1 that this framework assumes a constant population, does not assume interpersonal comparisons, and involves risk rather than other forms of uncertainty.

6.7.1. Constant and variable population, no interpersonal comparisons, risk

For derivations of the conclusion of Harsanyi's social aggregation theorem using weaker assumptions, see Fleurbaey (2009, Thm. 1) and Mongin and Pivato (2015). For generalizations of the theorem still assuming no interpersonal comparisons and risk, see Hammond (1988) (variable populations); Zhou (1997) (allowing for infinite populations); Danan et al. (2015) (dropping completeness, allowing for infinite populations); and McCarthy et al. (2019) (dropping completeness and continuity, allowing for infinite populations).

86 In Section 2.1, we explained how utility levels in Harsanyi's framework can be interpreted as welfare states in ours, so that prospects are probability measures over utility levels; the individual preorder orders such prospects by expected utility.

87 We add the caveat that we have not discussed the exact domain conditions that Harsanyi's utilitarian theorem requires, and we do not claim to have reproduced his conclusion under exactly those conditions.

88 Harsanyi's views about the variable population case are mainly presented in correspondence quoted in Ng (1983), but see also Harsanyi (1977a, note 12). The claim that the veil supports average utilitarianism was also made by Rawls (1971, §27), who coined the phrase 'veil of ignorance', and is also implicit in Vickrey (1945); see Kavka (1975), Barry (1977), and Ng (1983) for skepticism.

89 See for example Scanlon (1982) and Broome (1991, §3.3).

90 For a modification of the constant population veil of ignorance that is incompatible with our approach, see the defense of 'generalized utilitarianism' in Grant et al. (2010, 2012a). Their approach is designed to accommodate different attitudes to risk within the population in a way that in our framework would clash with Two-Stage Anonymity.

91 It is worth noting that this observation applies directly to the version of the veil of ignorance that is vindicated by our aggregation theorems (see Section 6.5). If the impartial observer's preferences over prospects violate strong independence, her social preferences over lotteries essentially cannot satisfy both Two-Stage (or Posterior) Anonymity and Full Pareto.

92 Theorem 5.3 makes a similar point: monotonicity for the social preorder entails strong independence, albeit with somewhat restrictive background conditions. Given that standard ways of rejecting strong independence typically maintain monotonicity, the latter might also be seen as a core commitment of Harsanyi's utilitarianism.

6.7.2. Constant population, interpersonal comparisons, risk

Harsanyi-like results that, like ours, explicitly assume a single individual preorder, and a form of anonymity, but weaken the premises of Harsanyi's utilitarian theorem, are given in Fleurbaey (2009) and Pivato (2013).

Fleurbaey (2009, Thm. 2) derives a constant population, finite support, total expected utility representation based on ordinary expected utility theory for the individual preorder, completeness for the social preorder, strong Pareto, anonymity, and a dominance condition similar to (M). This weakens the assumptions of Harsanyi's utilitarian theorem by not requiring continuity or independence for the social preorder. Differences in framework make difficult a strict comparison to our own results, but suffice to note that the ordinary EUT version of our Theorem 4.4(ii) rests on significantly weaker premises: it does not require social completeness, uses only Reduction to Prospects and Anteriority instead of strong Pareto, and uses only Two-Stage Anonymity instead of anonymity and dominance (cf. Section 6.2). Our Theorems 4.10 and 5.3 derive similar results based on Pareto and (M) without even assuming EUT for the individual preorder.

In the result that is closest to ours, Pivato (2013, Thm. 2.1) assumes Pareto and independence axioms for the individual and social preorders,[93] along with Anonymity. He shows that the social preorder must extend the one generated by the individual preorder. This shows that if the individual preorder is complete, then so is the social preorder; more generally, it restricts the degree of incompleteness of the social preorder in terms of that of the individual preorder.

Our Theorem 2.2 improves this picture. Its conclusion is that the social preorder is identical to the one generated by the individual preorder, showing precisely how incompleteness of the individual preorder determines that of the social preorder. Moreover, it makes no assumptions at all about the individual preorder, and rests on premises that are much weaker than Pivato's in most respects. The only premise that is not implied by any of Pivato's is a component of Reduction to Prospects: for any P, P' in $\mathbb{P}$, $P \curlywedge_{\mathbb{P}} P' \implies \mathcal{L}(P) \curlywedge \mathcal{L}(P')$. This principle expresses a natural connection between individual and social incompleteness, and we suggest that it is very plausible.

6.7.3. Variable population, interpersonal comparisons, risk

An extension of Harsanyi's utilitarian theorem to the variable population case, resulting in critical level utilitarianism, is given in Blackorby et al. (1998, 2007). Along with the full expected utility framework, this result assumes that at least some distributions have a critical level (roughly, a utility level at which creating an additional person would be a matter of social indifference). Under one of several important interpretations, Pivato (2014, Thm. 1) shows, roughly, that for variable but finite populations, there is a Harsanyi-like mixture-preserving total utility representation into a linearly ordered abelian group if and only if the social preorder is complete, anonymous, and satisfies a separability condition. Under this interpretation, the separability condition implies both strong independence and strong separability across individuals. Thus the main advance in terms of ethical assumptions is to have dispensed with continuity.[94] In contrast, our Theorem 3.5 neither assumes nor implies completeness, continuity, strong independence, strong separability across people, or the existence of a critical level for some distribution. When we further assume that the individual preorder is strongly independent, we obtain an expectational, and therefore mixture-preserving (note 57), total utility representation into a preordered vector space. As Pivato notes, a linearly ordered abelian group can always be embedded in an ordered vector space (and specifically into a lexicographic function space, by the Hahn embedding theorem), so our use of preordered vector spaces as utility spaces allows his kind of representation as a special case (compare Remark 4.7). One difference is that Pivato's framework is designed to allow for infinitesimal probabilities, whereas we have assumed standard real-valued probabilities. But as discussed in Section 2.7, this assumption plays no real role in either of our aggregation theorems.

6.7.4. Constant population, no interpersonal comparisons, uncertainty

In Harsanyi's social aggregation theorem, each individual's preorder and the social preorder apply to lotteries, understood as probability measures on a shared outcome space. His result is therefore most obviously relevant to situations in which probabilities are objective or universally agreed. If instead lotteries are interpreted as such things as Anscombe–Aumann or Savage acts, we obtain a framework in which individuals can apparently have differing subjective probabilities. But a number of results suggest that the gain in generality is illusory: strong Pareto plus a version of subjective expected utility for all the preorders implies that the subjective probabilities must be identical.[95] This seems to show that Harsanyi-style social aggregation is not possible for individuals whose subjective probabilities disagree.[96] These results do not assume interpersonal comparisons, but parallel difficulties will emerge if and when they are introduced.

In either case, one well-known reaction is to conclude that uncertainty should first be given a single representation before social aggregation takes place.[97] Such a representation could have many interpretations, such as expert consensus, social consensus, or the view of the social planner (compare Section 2.1). It could fall far short of being a single probability measure. But as illustrated in Examples 2.10 and 2.11, our aggregation theorems can still cope.

Acknowledgments

We are grateful to Peter Wakker and John Weymark for valuable advice. We have benefited greatly from sets of comments from two anonymous referees. The comments from one referee were exceptionally detailed and generous. David McCarthy thanks the Research Grants Council of the Hong Kong Special Administrative Region, China (HKU 750012H) for support. Teruji Thomas thanks the Arts and Humanities Research Council, United Kingdom for a research studentship, and the Leverhulme Trust, United Kingdom (RPG-2014-064).

Appendix A. Proofs

Some lemmas

We first collect together a few basic results about probability measures to which we will frequently appeal.

93 These are intermediate in strength between $(I_1)/(P_1)$ and $(I_2)/(P_2)$.

94 However, the main aim of Pivato (2014) is to provide an extension of his Theorem 1 to the infinite population setting, a topic not considered here; see note 32.

95 See e.g. Broome (1990), Mongin (1995, 1998), Mongin and Pivato (2015), and Zuber (2016).

96 For further discussion and results, see e.g. Mongin (1998), Gilboa et al. (2004), Chambers and Hayashi (2006, 2014), Danan et al. (2016), Mongin (2016), Gajdos et al. (2008), Crès et al. (2011), Gilboa et al. (2014), Alon and Gayer (2016), Billot and Vergopoulos (2016), Mongin and Pivato (2019), and Qu (2017).

97 See e.g. Mongin (1998, pp. 352–3).

Lemma A.1. *Let p be a probability measure on a measurable space Y, and let $A_1, \ldots, A_m$ and $B_1, \ldots, B_n$ be measurable subsets of Y. Suppose that each $y \in Y$ appears the same number of times in the sets A_i as in the B_i. Then*

$$\sum_{i=1}^{m} p(A_i) = \sum_{i=1}^{n} p(B_i).$$

Proof. Writing χ_{A_i} and χ_{B_i} for the characteristic functions of A_i and B_i, we have $\sum_{i=1}^{m} p(A_i) = \int_Y (\sum_{i=1}^{m} \chi_{A_i})\,\mathrm{d}p$ and $\sum_{i=1}^{n} p(B_i) = \int_Y (\sum_{i=1}^{n} \chi_{B_i})\,\mathrm{d}p$. By hypothesis, for each $y \in Y$, $\left(\sum_{i=1}^{m} \chi_{A_i}\right)(y) = \left(\sum_{i=1}^{n} \chi_{B_i}\right)(y)$. The two integrands are therefore identical. □

Lemma A.2. *Let X and Y be measurable spaces. Let $f: X \to Y$ be a measurable function, and let μ be a nonnegative measure on X. Then $\mu \circ f^{-1}$ is a measure on Y. Moreover, suppose $g: Y \to \mathbb{R}$ is integrable with respect to $\mu \circ f^{-1}$. Then $g \circ f$ is integrable with respect to μ, and one has*

$$\int_Y g \,\mathrm{d}(\mu \circ f^{-1}) = \int_X g \circ f \,\mathrm{d}\mu.$$

Proof. The proof that $\mu \circ f^{-1}$ is a measure is routine. For the second claim, we spell out a remark Bogachev (2007) makes after his Theorem 3.6.1. Since g is $(\mu \circ f^{-1})$-integrable, it agrees with some measurable function G on some set of $(\mu \circ f^{-1})$-measure 1 (see note 48). Bogachev's Theorem 3.6.1 shows that $G \circ f$ is integrable with respect to μ and that

$$\int_Y G \,\mathrm{d}(\mu \circ f^{-1}) = \int_X G \circ f \,\mathrm{d}\mu.$$

However, $g \circ f$ agrees with $G \circ f$ on a set of μ-measure 1; therefore $g \circ f$ is integrable with respect to μ, and replacing G by g in the displayed equation does not change the value of either side. □

The following two lemmas concern the notion of 'support' introduced in footnote 33.

Lemma A.3. *Let p be a probability measure on a measurable space Y, and let A be a subset of Y. The following conditions are equivalent:*

(1) The measure p is supported on A, in the sense that $p(B) = 0$ whenever $B \subset Y$ is measurable and disjoint from A.
(2) If $B_1, B_2 \subset Y$ are measurable and $B_1 \cap A = B_2 \cap A$, then $p(B_1) = p(B_2)$.

Proof. Suppose p is supported on A. Suppose $B_1 \cap A = B_2 \cap A$. Then $p(B_1) = p(B_1 \cap B_2) + p(B_1 \setminus B_2)$. $B_1 \setminus B_2$ is disjoint from A. So $p(B_1) = p(B_1 \cap B_2)$; this equals $p(B_2)$ by parallel reasoning. Conversely, suppose $p(B_1) = p(B_2)$ whenever $B_1 \cap A = B_2 \cap A$. Then, if B is disjoint from A, $p(B) = p(\emptyset) = 0$. □

The following lemma shows that any finitely supported probability measure can be written as a convex combination of delta-measures, regardless of whether singletons are measurable.

Lemma A.4. *Let p be a probability measure on a measurable space Y. Then p is finitely supported (i.e. supported on a finite set) if and only if $p = \sum_{i=1}^{n} \alpha_i 1_{y_i}$ for some $n \geq 1$, some distinct $y_1, \ldots, y_n \in Y$, and some $\alpha_1, \ldots, a_n > 0$ with $\sum_{i=1}^{n} \alpha_i = 1$. Moreover, if the sigma algebra on Y separates points,[98] then the y_i and α_i are uniquely determined (up to re-ordering).*

Proof. The right to left direction is obvious: if p is a weighted sum of delta-measures then it is supported on the set of points occurring in the sum. For the left to right, suppose p is supported on a finite set of n elements, $S = \{y_1, \ldots, y_n\}$. We can assume n is as small as possible. Now, suppose that every measurable set containing a given y_i also contains some distinct y_j; then p is in fact supported on $S \setminus \{y_j\}$, contradicting the minimality of n. We can therefore find measurable sets $A_1, \ldots, A_n$ such that $A_i \cap S = \{y_i\}$, and by excising the intersections of these sets, we can ensure that they are disjoint. Note that each $p(A_i) \neq 0$ by the minimality of n. We claim that $p = \sum_{i=1}^{n} p(A_i) 1_{y_i}$. To see this, note that (by Lemma A.3) two measurable sets have the same measure with respect to p if they contain the same elements of S; therefore, for any measurable A,

$$p(A) = p\left(\bigcup_{y_i \in A \cap S} A_i\right) = \sum_{y_i \in A \cap S} p(A_i) = \sum_{i=1}^{n} p(A_i) 1_{y_i}(A).$$

Of course, $\sum_{i=1}^{n} p(A_i) = p(Y) = 1$.

As to the uniqueness claim, suppose that $p = \sum_{i=1}^{n} \alpha_i 1_{y_i}$ can also be written as $\sum_{i=1}^{m} \beta_i 1_{x_i}$, with the x_i mutually distinct, the y_i mutually distinct, and all $\alpha_i, \beta_i > 0$. It suffices to show that each y_j is equal to some x_k, and that $\alpha_j = \beta_k$. By the hypothesis that the sigma algebra separates points, we can find a measurable set B_j whose intersection with $\{y_1, \ldots, y_n, x_1, \ldots, x_m\}$ is precisely $\{y_j\}$. Thus, from the second expression for p, $p(B_j) = 0$ unless y_j is equal to some x_k, in which case $p(B_j) = \beta_k$. But from the first expression $p(B_j) = \alpha_j > 0$. This shows that some x_k must equal y_j, and then $\beta_k = \alpha_j$, as claimed. □

Proof of Proposition 2.4. Suppose that $\succsim$ is generated by $\succsim_{\mathbb{P}}$. Suppose that $L(B) = L'(B)$ for every measurable and permutation-invariant $B \subset \mathbb{D}$. We want to show $L \sim L'$. Suppose given measurable $A \subset \mathbb{W}$. We can write

$$\begin{aligned} \#\mathbb{I} \cdot p_L(A) &= \sum_{i \in \mathbb{I}} \mathcal{P}_i(L)(A) \\ &= \sum_{i \in \mathbb{I}} L(\mathcal{W}_i^{-1}(A)) = \sum_{n=1}^{\#\mathbb{I}} L\left(\bigcup_{\substack{I \subset \mathbb{I} \\ \#I = n}} \bigcap_{i \in I} \mathcal{W}_i^{-1}(A)\right). \end{aligned} \tag{10}$$

The first equation is from the definition of p_L, the second from the definition of $\mathcal{P}_i$, and the last equation is a direct application of Lemma A.1. (Note that if a distribution is an element of the argument of L in exactly k terms of the left-hand sum, then it also an element of the argument of L in exactly k terms of the right-hand sum, namely those with $n = 1, 2, \ldots, k$.) On the right hand side, all arguments of L are permutation-invariant. We therefore find that

$$\#\mathbb{I} \cdot p_L(A) = \#\mathbb{I} \cdot p_{L'}(A)$$

for arbitrary measurable A. Hence $p_L = p_{L'}$. Since $\succsim_{\mathbb{P}}$ generates $\succsim$, we must have $L \sim L'$, as required. □

Proof of Lemma 3.3. For (i), suppose given $L \in \mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ and $i \in \mathbb{I}^{\infty} \setminus \mathbb{I}$. Let A be measurable in $\mathbb{W}^{\mathrm{v}}$ with $\Omega \in A$. Then $\mathbb{D}_{\mathbb{I}}^{\mathrm{v}} \subset (\mathcal{W}_i^{\mathrm{v}})^{-1}(A)$, hence $\mathcal{P}_i^{\mathrm{v}}(L)(A) = L((\mathcal{W}_i^{\mathrm{v}})^{-1}(A)) = 1$. Since this is true for every such A, we must have $\mathcal{P}_i^{\mathrm{v}}(L) = 1_\Omega$. Since $\mathbb{L}^{\mathrm{v}}$ is nonempty, and hence by domain condition **($\mathbf{D}^{\mathrm{v}}$)** some $\mathbb{L}_{\mathbb{I}}^{\mathrm{v}}$ is nonempty, by domain condition **($\mathbf{A}^{\mathrm{v}}$)** we must have $1_\Omega \in \mathbb{P}^{\mathrm{v}}$.

For (ii), for any finite $\mathbb{I}$, we have $d_\Omega = \mathcal{D}_{\mathbb{I}}^{\mathrm{v}}(\Omega) \in \mathbb{D}^{\mathrm{v}}$. Now invoke assumption **($\mathbf{B}^{\mathrm{v}}$)**.

For (iii), we have $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(1_\Omega) \in \mathbb{L}^{\mathrm{v}}$, and we claim that $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(1_\Omega) = 1_{d_\Omega}$. Indeed, for any measurable $B \subset \mathbb{D}^{\mathrm{v}}$ with $d_\Omega \in B$, we have $\Omega \in (\mathcal{D}_{\mathbb{I}}^{\mathrm{v}})^{-1}(B)$. Therefore $\mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(1_\Omega)(B) = 1_\Omega((\mathcal{D}_{\mathbb{I}}^{\mathrm{v}})^{-1}(B)) = 1$, as desired. Moreover, for any measurable $A \subset \mathbb{W}^{\mathrm{v}}$ with $\Omega \in A$, we

[98] I.e. for all $y, y' \in Y$, there is a measurable set containing y but not y'.

have, for any $i \in \mathbb{I}^\infty$, $d_\Omega \in (\mathcal{W}_i^\mathsf{v})^{-1}(A)$. Therefore $\mathcal{P}_i^\mathsf{v}(1_{d_\Omega})(A) = 1_{d_\Omega}((\mathcal{W}_i^\mathsf{v})^{-1}(A)) = 1$. Therefore $\mathcal{P}_i^\mathsf{v}(1_{d_\Omega}) = 1_\Omega$.

For (iv), suppose we have $L \in \mathbb{L}_\mathbb{I}^\mathsf{v}$, and $\mathcal{P}_i^\mathsf{v}(L) = 1_\Omega$ for all $i \in \mathbb{I}^\infty$. Then $L((\mathcal{W}_i^\mathsf{v})^{-1}(\mathbb{W})) = \mathcal{P}_i^\mathsf{v}(L)(\mathbb{W}) = 0$ for all i. (Note that $\mathbb{W}$ is measurable in $\mathbb{W}^\mathsf{v}$, being the complement of $\{\Omega\}$, which is measurable by hypothesis.) Defining $B := \bigcup_{i\in\mathbb{I}}(\mathcal{W}_i^\mathsf{v})^{-1}(\mathbb{W})$, we must have $L(B) = 0$. Suppose given measurable $B' \subset \mathbb{D}^\mathsf{v}$ with $d_\Omega \notin B'$. We have $B' \cap \mathbb{D}_\mathbb{I}^\mathsf{v} \subset B \cap \mathbb{D}_\mathbb{I}^\mathsf{v}$, so $B' \cap \mathbb{D}_\mathbb{I}^\mathsf{v} = B' \cap B \cap \mathbb{D}_\mathbb{I}^\mathsf{v}$. Since L is supported on $\mathbb{D}_\mathbb{I}^\mathsf{v}$, Lemma A.3 gives $L(B') = L(B' \cap B) \le L(B)$, so $L(B') = 0$. Therefore $L = 1_{d_\Omega}$. □

Lemma A.5. *Suppose given a tuple $(\mathbb{I}^\infty, \mathbb{W}^\mathsf{v}, \mathbb{P}^\mathsf{v}, \mathbb{D}^\mathsf{v}, \mathbb{L}^\mathsf{v})$ satisfying the variable population domain conditions* **(A^v)–(D^v)**, *and let $\mathcal{F}$ denote the sigma algebra on $\mathbb{D}^\mathsf{v}$. Then $\mathcal{F}$ is contained in a sigma algebra $\overline{\mathcal{F}}$ such that (a) $\overline{\mathcal{F}}$ is coherent; (b) every $L \in \mathbb{L}^\mathsf{v}$ extends naturally to a probability measure $\overline{L}$ with respect to $\overline{\mathcal{F}}$; (c) if we write $\overline{\mathbb{D}^\mathsf{v}}$ for $\mathbb{D}^\mathsf{v}$ equipped with the sigma algebra $\overline{\mathcal{F}}$, and $\overline{\mathbb{L}^\mathsf{v}} := \{\overline{L} : L \in \mathbb{L}^\mathsf{v}\}$, then $(\mathbb{I}^\infty, \mathbb{W}^\mathsf{v}, \mathbb{P}^\mathsf{v}, \overline{\mathbb{D}^\mathsf{v}}, \overline{\mathbb{L}^\mathsf{v}})$ again satisfies the domain conditions* **(A^v)–(D^v)**.

Proof. Define $\overline{\mathcal{F}}$ by the rule that $\overline{B} \subset \mathbb{D}^\mathsf{v}$ is in $\overline{\mathcal{F}}$ if and only if, for every finite $\mathbb{I} \subset \mathbb{I}^\infty$, there exists some $B \in \mathcal{F}$ such that $\overline{B} \cap \mathbb{D}_\mathbb{I}^\mathsf{v} = B \cap \mathbb{D}_\mathbb{I}^\mathsf{v}$. It is easy to check that $\overline{\mathcal{F}}$ is a coherent sigma algebra containing $\mathcal{F}$, and the restriction of $\overline{\mathcal{F}}$ to each $\mathbb{D}_\mathbb{I}^\mathsf{v} \subset \mathbb{D}^\mathsf{v}$ is the same as the restriction of $\mathcal{F}$.

Given $L \in \mathbb{L}^\mathsf{v}$, we can extend L to a probability measure $\overline{L}$ on $\overline{\mathcal{F}}$: if $L \in \mathbb{L}_\mathbb{I}^\mathsf{v}$, then, for B and $\overline{B}$ related as above, $\overline{L}(\overline{B}) := L(B)$. To see that $\overline{L}(\overline{B})$ is independent of the choice of $\mathbb{I}$ and B, use Lemma A.3.

It remains to verify that the domain conditions **(A^v)–(D^v)** are satisfied. For **(A^v)**, it is obvious that each $\mathcal{W}_i^\mathsf{v}$ is measurable with respect to $\overline{\mathcal{F}}$, since $\overline{\mathcal{F}}$ contains $\mathcal{F}$. For clarity, let us retain the notation $\mathcal{P}_i^\mathsf{v}$ for the map from $\mathbb{L}^\mathsf{v}$ to prospects defined in terms of $\mathcal{F}$, and write $\overline{\mathcal{P}_i^\mathsf{v}}$ for analogous map on $\overline{\mathbb{L}^\mathsf{v}}$ defined in terms of $\overline{\mathcal{F}}$. Then $\overline{\mathcal{P}_i^\mathsf{v}}(\overline{L}) = \overline{L} \circ (\mathcal{W}_i^\mathsf{v})^{-1} = L \circ (\mathcal{W}_i^\mathsf{v})^{-1} = \mathcal{P}_i^\mathsf{v}(L)$, which shows that $\overline{\mathcal{P}_i^\mathsf{v}}(\overline{\mathbb{L}^\mathsf{v}}) \subset \mathbb{P}^\mathsf{v}$.

For **(B^v)**, each $\mathcal{D}_\mathbb{I}^\mathsf{v}$ is measurable with respect to $\overline{\mathcal{F}}$: we have $(\mathcal{D}_\mathbb{I}^\mathsf{v})^{-1}(\overline{B}) = (\mathcal{D}_\mathbb{I}^\mathsf{v})^{-1}(B)$, for $B, \overline{B}$ related as above, showing that $(\mathcal{D}_\mathbb{I}^\mathsf{v})^{-1}(\overline{B})$ is measurable in $\mathbb{W}^\mathsf{v}$. Again distinguishing $\mathcal{L}_\mathbb{I}^\mathsf{v} : \mathbb{P}^\mathsf{v} \to \mathbb{L}^\mathsf{v}$ defined in terms of $\mathcal{F}$ from the map $\overline{\mathcal{L}_\mathbb{I}^\mathsf{v}}$ defined in terms of $\overline{\mathcal{F}}$, we find $\overline{\mathcal{L}_\mathbb{I}^\mathsf{v}}(P)(\overline{B}) = P((\mathcal{D}_\mathbb{I}^\mathsf{v})^{-1}(\overline{B})) = P((\mathcal{D}_\mathbb{I}^\mathsf{v})^{-1}(B)) = \mathcal{L}_\mathbb{I}^\mathsf{v}(P)(B)$. This shows that $\overline{\mathcal{L}_\mathbb{I}^\mathsf{v}}(P) = \overline{\mathcal{L}_\mathbb{I}^\mathsf{v}(P)} \in \overline{\mathbb{L}^\mathsf{v}}$, so $\overline{\mathcal{L}_\mathbb{I}^\mathsf{v}}(\mathbb{P}^\mathsf{v}) \subset \overline{\mathbb{L}^\mathsf{v}}$.

For **(C^v)**, given $\sigma \in \Sigma^\infty$ and $\overline{B} \in \overline{\mathcal{F}}$, we can find $B \in \mathcal{F}$ with $\overline{B} \cap \mathbb{D}_{\sigma\mathbb{I}}^\mathsf{v} = B \cap \mathbb{D}_{\sigma\mathbb{I}}^\mathsf{v}$. Then $(\sigma^{-1}\overline{B}) \cap \mathbb{D}_\mathbb{I}^\mathsf{v} = (\sigma^{-1}B) \cap \mathbb{D}_\mathbb{I}^\mathsf{v}$. Since this holds for any $\mathbb{I}$, $\sigma^{-1}\overline{B}$ is in $\overline{\mathcal{F}}$, and the action of Σ^∞ is measurable with respect to $\overline{\mathcal{F}}$. Moreover, for $L \in \mathbb{L}_\mathbb{I}^\mathsf{v}$, we find $\sigma\overline{L}(\overline{B}) = \overline{L}(\sigma^{-1}\overline{B}) = L(\sigma^{-1}B) = \sigma L(B) = \overline{\sigma L}(\overline{B})$, so $\overline{\mathbb{L}^\mathsf{v}}$ is Σ^∞-invariant.

Finally, for **(D^v)**, it is easy to check that, given $L \in \mathbb{L}_\mathbb{I}^\mathsf{v}$, we have also $\overline{L} \in \overline{\mathbb{L}_\mathbb{I}^\mathsf{v}}$. □

Proof of Proposition 3.7. Suppose that $\succsim^\mathsf{v}$ is generated by $\succsim_{\mathbb{P}^\mathsf{v}}$. Suppose that $L(B) = L'(B)$ for every measurable and Σ^∞-invariant $B \subset \mathbb{D}^\mathsf{v}$. We want to show $L \sim^\mathsf{v} L'$. Pick finite non-empty $\mathbb{I} \subset \mathbb{I}^\infty$ such that $L, L' \in \mathbb{L}_\mathbb{I}^\mathsf{v}$. It suffices to show that $p_L^\mathbb{I} = p_{L'}^\mathbb{I}$. Now, for any measurable $A' \subset \mathbb{W}^\mathsf{v}$, $p_L^\mathbb{I}(A') = 1 - p_L^\mathbb{I}(\mathbb{W}^\mathsf{v} \setminus A')$. Since either $\Omega \notin A'$ or $\Omega \notin \mathbb{W}^\mathsf{v} \setminus A'$, it suffices to show that $p_L^\mathbb{I}(A) = p_{L'}^\mathbb{I}(A)$ for every measurable $A \subset \mathbb{W}^\mathsf{v}$ such that $\Omega \notin A$.

For each number n, $1 \le n \le \#\mathbb{I}$, let B_n be the set of distributions in which at least n individuals have welfare states in A; and, for any finite population $\mathbb{J}$, let $B_n^\mathbb{J}$ be the set of distributions in which at least n individuals in $\mathbb{J}$ have welfare states in A. That is:

$$B_n := \bigcup_{\substack{I\subset\mathbb{I}^\infty \\ \#I=n}} \bigcap_{i\in I} (\mathcal{W}_i^\mathsf{v})^{-1}(A) \qquad B_n^\mathbb{J} := \bigcup_{\substack{I\subset\mathbb{J} \\ \#I=n}} \bigcap_{i\in I} (\mathcal{W}_i^\mathsf{v})^{-1}(A).$$

Each set $B_n^\mathbb{J}$ is measurable in $\mathbb{D}^\mathsf{v}$ (since the $\mathcal{W}_i^\mathsf{v}$ are measurable functions, and we take finitely many intersections and unions); therefore $B_n^\mathbb{J} \cap \mathbb{D}_\mathbb{J}^\mathsf{v}$ is measurable in $\mathbb{D}_\mathbb{J}^\mathsf{v}$ (since its sigma algebra is the restriction of the one on $\mathbb{D}^\mathsf{v}$). Given the assumption that $\Omega \notin A$, we have $B_n \cap \mathbb{D}_\mathbb{J}^\mathsf{v} = B_n^\mathbb{J} \cap \mathbb{D}_\mathbb{J}^\mathsf{v}$. Therefore, if the sigma algebra on $\mathbb{D}^\mathsf{v}$ is coherent, B_n is also measurable in $\mathbb{D}^\mathsf{v}$. If, instead of coherence, we assume that $\mathbb{I}^\infty$ is countable, then the definition of B_n involves only countable unions and finite intersections of measurable sets, so B_n is again measurable.

In exact parallel to (10) in the proof of Proposition 2.4, we have

$$\#\mathbb{I} \cdot p_L^\mathbb{I}(A) = \sum_{i\in\mathbb{I}} \mathcal{P}_i^\mathsf{v}(A) = \sum_{i\in\mathbb{I}} L((\mathcal{W}_i^\mathsf{v})^{-1}(A)) = \sum_{n=1}^{\#\mathbb{I}} L(B_n^\mathbb{I})$$

again using Lemma A.1 for the last equation. Now, because L is supported on $\mathbb{D}_\mathbb{I}^\mathsf{v}$, and, for each n, $B_n^\mathbb{I} \cap \mathbb{D}_\mathbb{I}^\mathsf{v} = B_n \cap \mathbb{D}_\mathbb{I}^\mathsf{v}$, Lemma A.3 gives $L(B_n^\mathbb{I}) = L(B_n)$. Therefore $\#\mathbb{I} \cdot p_L^\mathbb{I}(A) = \sum_{n=1}^{\#\mathbb{I}} L(B_n)$, and similarly $\#\mathbb{I} \cdot p_{L'}^\mathbb{I}(A) = \sum_{n=1}^{\#\mathbb{I}} L'(B_n)$.

Now since each set B_n is Σ^∞-invariant, $L(B_n) = L'(B_n)$. Therefore $p_L^\mathbb{I}(A) = p_{L'}^\mathbb{I}(A)$, establishing Posterior Anonymity. □

The following lemma is used in the proof of Proposition 3.9. Preordered vector spaces are defined in Section 4.2.

Lemma A.6. *Every preorder has a representation with values in a preordered vector space.*

Proof. The following construction is inspired by Conrad (1953); see McCarthy et al. (2017b, Thm. 11) for an alternative. Suppose $\succsim_X$ is a preorder on a set X. Set $\overline{X} := X/\!\sim_X$, and for $x \in X$ let $\overline{x}$ be its class in $\overline{X}$. There is a partial order $\succsim_{\overline{X}}$ on $\overline{X}$ defined by $\overline{x} \succsim_{\overline{X}} \overline{y} \iff x \succsim_X y$. Let $\mathbb{V}$ be the vector space of functions $f : \overline{X} \to \mathbb{R}$ such that $\text{supp}(f) := \{\overline{x} \in \overline{X} : f(\overline{x}) \neq 0\}$ satisfies the ascending chain condition, i.e. every nonempty subset has a $\succsim_{\overline{X}}$-maximal element. Define a relation $\succsim_\mathbb{V}$ on $\mathbb{V}$ by the rule that $f \succsim_\mathbb{V} g \iff f(\overline{x}) > g(\overline{x})$ for all $\overline{x}$ maximal in $\text{supp}(f-g)$. This makes $\succsim_\mathbb{V}$ into a preordered vector space; the only non-trivial claim to prove is that $\succsim_\mathbb{V}$ is transitive.

Suppose, for this, that $f \succsim_\mathbb{V} g \succsim_\mathbb{V} h$. We want to show that, given $\overline{x}$ maximal in $\overline{X}_{fh} := \text{supp}(f-h)$, we have $(f-h)(\overline{x}) > 0$. Now, $(f-h)(\overline{x}) = (f-g)(\overline{x}) + (g-h)(\overline{x})$, so at least one of the latter two terms must be non-zero. Correspondingly, $\overline{x}$ is in $\overline{X}_{fg} \cup \overline{X}_{gh}$. This union also satisfies the ascending chain condition, so we can find $\overline{y}$ maximal in $\{\overline{y} \in \overline{X}_{fg} \cup \overline{X}_{gh} : \overline{y} \succsim_{\overline{X}} \overline{x}\}$. This $\overline{y}$ is automatically maximal in $\overline{X}_{fg} \cup \overline{X}_{gh}$. So if $\overline{y}$ is in $\overline{X}_{fg}$, it must be maximal there, and $(f-g)(\overline{y}) > 0$; if $\overline{y}$ is in $\overline{X}_{gh}$, it must be maximal *there*, and $(g-h)(\overline{y}) > 0$; therefore, either way, $(f-h)(\overline{y}) = (f-g)(\overline{y}) + (g-h)(\overline{y}) > 0$. Thus $\overline{y} \in \overline{X}_{fh}$. Since $\overline{x}$ is maximal in $\overline{X}_{fh}$, and $\overline{y} \succsim_{\overline{X}} \overline{x}$, and $\succsim_{\overline{X}}$ is a partial order, we must actually have $\overline{x} = \overline{y}$, and $(f-h)(\overline{x}) > 0$. Therefore $\succsim_\mathbb{V}$ is transitive.

Finally, consider the function that maps $x \in X$ to the characteristic function of $\{\overline{x}\}$; this is a representation of $\succsim_X$ with values in $\mathbb{V}$. □

Proof of Proposition 3.9. (i) Applying Lemma A.6 to $X = \mathbb{P}_\Omega$, we have a representation $U : \mathbb{P}_\Omega \to \mathbb{V}$ of $\succsim_{\mathbb{P}_\Omega}$, for some preordered vector space $(\mathbb{V}, \succsim_\mathbb{V})$. Since $\mathbb{P}^\mathsf{v}$ extends $\mathbb{P}$, each member of $\mathbb{P}^\mathsf{v}$ can be written in the form $P_\alpha := \alpha P + (1-\alpha)1_\Omega$ for some $P \in \mathbb{P}$, $\alpha \in [0,1]$. This presentation is unique except when $\alpha = 0$. Define a function $\overline{U} : \mathbb{P}^\mathsf{v} \to \mathbb{V}$ by the rule

$$\overline{U}(P_\alpha) = \alpha U(P) + (1-\alpha)U(1_\Omega).$$

Let $\succsim_{\mathbb{P}^\mathsf{v}}$ be the preorder on $\mathbb{P}^\mathsf{v}$ represented by $\overline{U}$. We claim that $\succsim_{\mathbb{P}^\mathsf{v}}$ is Omega Independent and extends $\succsim_{\mathbb{P}_\Omega}$.

For all $P \in \mathbb{P}_\Omega$, $\overline{U}(P) = U(P)$, so $\succsim_{\mathbb{P}^v}$ extends $\succsim_{\mathbb{P}_\Omega}$. To show that $\succsim_{\mathbb{P}^v}$ satisfies Omega Independence, suppose given $P, P' \in \mathbb{P}^v$ and $\alpha \in (0,1) \cap \mathbb{Q}$. We wish to show that $P \succsim_{\mathbb{P}^v} P' \iff P_\alpha \succsim_{\mathbb{P}^v} P'_\alpha$. We have $P = Q_\beta$ and $P' = Q'_\gamma$ for some $Q, Q' \in \mathbb{P}$, $\beta, \gamma \in [0,1]$. Then $P_\alpha = Q_{\alpha\beta}$ and $P'_\alpha = Q'_{\alpha\gamma}$. Thus:

$$\begin{aligned} P \succsim_{\mathbb{P}^v} P' &\iff \overline{U}(P) \succsim_{\mathbb{V}} \overline{U}(P') \\ &\iff \beta U(Q) + (1-\beta)U(1_\Omega) \succsim_{\mathbb{V}} \gamma U(Q') \\ &\quad + (1-\gamma)U(1_\Omega) \\ &\iff \alpha\beta U(Q) + (1-\alpha\beta)U(1_\Omega) \succsim_{\mathbb{V}} \alpha\gamma U(Q') \\ &\quad + (1-\alpha\gamma)U(1_\Omega) \\ &\iff \overline{U}(P_\alpha) \succsim_{\mathbb{V}} \overline{U}(P'_\alpha) \\ &\iff P_\alpha \succsim_{\mathbb{P}^v} P'_\alpha \end{aligned}$$

Here the third line is obtained from the second by applying the order-preserving transformation of $\mathbb{V}$ given by $v \mapsto \alpha v + (1-\alpha)U(1_\Omega)$. This establishes that $\succsim_{\mathbb{P}^v}$ is Omega Independent and extends $\succsim_{\mathbb{P}_\Omega}$.

(ii) Given $\succsim_{\mathbb{P}_\Omega}$, there is a unique preorder $\succsim_{\mathbb{P}^v}$ on $\mathbb{P}^v$ that extends $\succsim_{\mathbb{P}_\Omega}$ and satisfies the following property: for any $P, Q \in \mathbb{P}$ and $\alpha, \beta \in (0,1)$,

$$P_\alpha \succ_{\mathbb{P}^v} Q \quad \text{and} \quad P_\alpha \succ_{\mathbb{P}^v} 1_\Omega \quad \text{and} \quad P_\alpha \sim_{\mathbb{P}^v} Q_\beta.$$

In other words, elements of $\mathbb{P}^v$ not in $\mathbb{P}_\Omega$ are ranked as equals above all elements of $\mathbb{P}_\Omega$. This $\succsim_{\mathbb{P}^v}$ does not satisfy Omega Independence: for $P, Q \in \mathbb{P}, \alpha, \beta \in (0,1)$ as before, we have $\alpha P + (1-\alpha)1_\Omega = P_\alpha \succsim_{\mathbb{P}^v} Q_{\alpha\beta} = \alpha Q_\beta + (1-\alpha)1_\Omega$, but we do not have $P \succsim_{\mathbb{P}^v} Q_\beta$ as Omega Independence requires for rational values of α. □

Proof of Proposition 4.2. For (i), we present only the variable population case, the constant population case being exactly parallel. The general strategy is to use the assumption that $\succsim_{\mathbb{P}^v}$ generates $\succsim^v$ as follows. We use it directly to derive each condition on $\succsim^v$ from the same condition on $\succsim_{\mathbb{P}^v}$; conversely, we use Reduction to Prospects (which holds by Theorem 3.5) to derive the condition on $\succsim_{\mathbb{P}^v}$ from the condition on $\succsim^v$. Moreover, we can use the fact that the maps $L \mapsto p^{\mathbb{I}}_L$ and $P \mapsto \mathcal{L}^v_{\mathbb{I}}(P)$ are mixture preserving (in the sense of note 28). The arguments for the different conditions are very similar, so we only present the proof for (I_3), or strong independence. Recall that for a preorder $\succsim_X$ on a convex set X, (I_3) is equivalent to the condition

$p \succsim_X p'$ if and only if $\alpha p + (1-\alpha)q \succsim_X \alpha p' + (1-\alpha)q$ for all $p, p', q \in X$, $\alpha \in (0,1)$.

Suppose first that $\succsim_{\mathbb{P}^v}$ satisfies (I_3). Suppose given $L, L', M \in \mathbb{L}^v$, $\alpha \in (0,1)$. There is some finite, nonempty $\mathbb{I} \subset \mathbb{I}^\infty$ with L, L', $M \in \mathbb{L}^v_{\mathbb{I}}$. Then

$$\begin{aligned} L \succsim^v L' &\iff p^{\mathbb{I}}_L \succsim_{\mathbb{P}^v} p^{\mathbb{I}}_{L'} \\ &\qquad (\succsim_{\mathbb{P}^v} \text{ generates } \succsim^v) \\ &\iff \alpha p^{\mathbb{I}}_L + (1-\alpha)p^{\mathbb{I}}_M \succsim_{\mathbb{P}^v} \alpha p^{\mathbb{I}}_{L'} + (1-\alpha)p^{\mathbb{I}}_M \\ &\qquad ((I_3) \text{ for } \succsim_{\mathbb{P}^v}) \\ &\iff p^{\mathbb{I}}_{\alpha L + (1-\alpha)M} \succsim_{\mathbb{P}^v} p^{\mathbb{I}}_{\alpha L' + (1-\alpha)M} \\ &\qquad (L \mapsto p^{\mathbb{I}}_L \text{ is mixture preserving}) \\ &\iff \alpha L + (1-\alpha)M \succsim^v \alpha L' + (1-\alpha)M \\ &\qquad (\succsim_{\mathbb{P}^v} \text{ generates } \succsim^v). \end{aligned}$$

Therefore $\succsim^v$ satisfies (I_3), as claimed. Conversely, suppose $\succsim^v$ satisfies (I_3), and suppose given $P, Q, R \in \mathbb{P}^v$. Then

$$\begin{aligned} P \succsim_{\mathbb{P}^v} Q &\iff \mathcal{L}^v_{\mathbb{I}}(P) \succsim^v \mathcal{L}^v_{\mathbb{I}}(Q) \\ &\qquad (\text{Reduction to Prospects}) \\ &\iff \alpha\mathcal{L}^v_{\mathbb{I}}(P) + (1-\alpha)\mathcal{L}^v_{\mathbb{I}}(R) \succsim^v \alpha\mathcal{L}^v_{\mathbb{I}}(Q) + (1-\alpha)\mathcal{L}^v_{\mathbb{I}}(R) \\ &\qquad ((I_3) \text{ for } \succsim^v) \\ &\iff \mathcal{L}^v_{\mathbb{I}}(\alpha P + (1-\alpha)R) \succsim^v \mathcal{L}^v_{\mathbb{I}}(\alpha Q + (1-\alpha)R) \\ &\qquad (\mathcal{L}^v_{\mathbb{I}} \text{ is mixture preserving}) \\ &\iff \alpha P + (1-\alpha)R \succsim_{\mathbb{P}^v} \alpha P + (1-\alpha)R \\ &\qquad (\text{Reduction to Prospects}). \end{aligned}$$

So (I_3) for $\succsim^v$ implies (I_3) for $\succsim_{\mathbb{P}^v}$.

Now let us turn to part (ii) of the proposition, beginning with the constant population case. First a general observation. Suppose given topological spaces X, Y with preorders $\succsim_X, \succsim_Y$, and a function $f: X \to Y$. Assume (1) that f is continuous, and (2) that for all $a, b \in X$, $a \succsim_X b \iff f(a) \succsim_Y f(b)$. Then, we claim, if $\succsim_Y$ is continuous, so is $\succsim_X$. Indeed, for any $q \in X$, we find

$$\begin{aligned} \{p \in X : p \succsim_X q\} &= \{p \in X : f(p) \succsim_Y f(q)\} \\ &= f^{-1}\{y \in Y : y \succsim_Y f(q)\}. \end{aligned}$$

The right-hand side is the inverse image of a closed set under a continuous function, so it is closed. A similar calculation shows that $\{p \in X : q \succsim_X p\}$ is closed; hence $\succsim_X$ is continuous.

Taking $f = \mathcal{L}: \mathbb{P} \to \mathbb{L}$, assumption (1) in the previous paragraph is part of (Top), and assumption (2) follows from Reduction to Prospects, which itself follows by Theorem 2.2 from the hypothesis that $\succsim_{\mathbb{P}}$ generates $\succsim$. We conclude that, if $\succsim$ is continuous, so is $\succsim_{\mathbb{P}}$. Conversely, define $f: \mathbb{L} \to \mathbb{P}$ by $f(L) = p_L$. Assumption (1) follows from the continuity of mixing and of every $\mathcal{P}_i$, whereas (2) is part of what it means for $\succsim$ to be generated by $\succsim_{\mathbb{P}}$. We conclude that, if $\succsim_{\mathbb{P}}$ is continuous, so is $\succsim$.

As for the variable population case, let $\succsim^v_{\mathbb{I}}$ be the restriction of $\succsim^v$ to $\mathbb{L}^v_{\mathbb{I}}$, and equip the latter with a topology as a subspace of $\mathbb{L}^v$. It suffices to prove two claims: first, that $\succsim_{\mathbb{P}^v}$ is continuous on $\mathbb{P}^v$ if and only if $\succsim^v_{\mathbb{I}}$ is continuous on $\mathbb{L}^v_{\mathbb{I}}$, for every finite $\mathbb{I} \subset \mathbb{I}^\infty$; second, that the latter holds if and only if $\succsim^v$ is continuous on $\mathbb{L}^v$.

The first claim follows from the logic just used for the constant population case. As for the second claim, suppose (from right to left) that $\succsim^v$ is continuous on $\mathbb{L}^v$. By definition of $\succsim^v_{\mathbb{I}}$, for any $M \in \mathbb{L}^v_{\mathbb{I}}$, $\{L \in \mathbb{L}^v_{\mathbb{I}}: L \succsim^v_{\mathbb{I}} M\} = \{L \in \mathbb{L}^v: L \succsim^v M\} \cap \mathbb{L}^v_{\mathbb{I}}$. Since $\succsim^v$ is continuous, $\{L \in \mathbb{L}^v: L \succsim^v M\}$ is closed in $\mathbb{L}^v$, and since $\mathbb{L}^v_{\mathbb{I}}$ has the subspace topology, this intersection is closed in $\mathbb{L}^v_{\mathbb{I}}$. A similar argument shows $\{L \in \mathbb{L}^v_{\mathbb{I}}: M \succsim^v_{\mathbb{I}} L\}$ is closed in $\mathbb{L}^v_{\mathbb{I}}$ as well. Therefore $\succsim^v_{\mathbb{I}}$ is continuous. For left to right, it suffices to show that, for any $L_0 \in \mathbb{L}^v$, the set $X = \{L \in \mathbb{L}^v : L \succsim^v L_0\}$ is closed in $\mathbb{L}^v$ (and similarly that $\{L \in \mathbb{L}^v : L_0 \succsim^v L\}$ is closed). By topological coherence, it is enough to show that $X \cap \mathbb{L}^v_{\mathbb{I}}$ is closed in $\mathbb{L}^v_{\mathbb{I}}$, for every finite $\mathbb{I}$. Pick finite $\mathbb{J}$ such that L_0 is in $\mathbb{L}^v_{\mathbb{J}}$, and let $\mathbb{K} = \mathbb{I} \cup \mathbb{J}$, so L_0 is also in $\mathbb{L}^v_{\mathbb{K}}$. Then $X \cap \mathbb{L}^v_{\mathbb{K}}$ is closed in $\mathbb{L}^v_{\mathbb{K}}$, by continuity of $\succsim^v_{\mathbb{K}}$. That means there is some closed $V \subset \mathbb{L}^v$ such that $V \cap \mathbb{L}^v_{\mathbb{K}} = X \cap \mathbb{L}^v_{\mathbb{K}}$. But then $X \cap \mathbb{L}^v_{\mathbb{I}} = V \cap \mathbb{L}^v_{\mathbb{I}}$ is closed in $\mathbb{L}^v_{\mathbb{I}}$, as desired. □

Proof of Lemma 4.3. We first check that $\succsim_X$ satisfies (I_3) if it satisfies Vector EUT, that is, if it has a Vector EU representation U with respect to some preordered vector space $(\mathbb{V}, \succsim_{\mathbb{V}})$ and separating vector space $\mathbb{V}'$ of linear functionals on $\mathbb{V}$. The main point is that U is mixture preserving; as defined in note 28, this means that $U(\alpha p + (1-\alpha)q) = \alpha U(p) + (1-\alpha)U(q)$ for $p, q \in \mathcal{P}(Y)$, $\alpha \in [0,1]$. This is just the linearity of the integral; in detail, for any $\Lambda \in \mathbb{V}'$ we have $\Lambda(\alpha U(p) + (1-\alpha)U(q)) = \alpha\Lambda(U(p)) + (1-\alpha)\Lambda(U(q)) = \alpha\int_Y \Lambda \circ u \, dp + (1-\alpha)\int_Y \Lambda \circ u \, dq = \int_Y \Lambda \circ u \, d(\alpha p + (1-\alpha)q) = \Lambda(U(\alpha p + (1-\alpha)q))$. Since this works for any Λ in $\mathbb{V}'$, we obtain the mixture preservation equation. Now, to derive (I_3), for $p, p', q \in \mathcal{P}(Y)$ and $\alpha \in (0, 1)$,

$$\begin{aligned} p \succsim_X p' &\iff U(p) \succsim_{\mathbb{V}} U(p') \\ &\iff \alpha U(p) + (1-\alpha)U(q) \succsim_{\mathbb{V}} \alpha U(p') + (1-\alpha)U(q) \\ &\iff U(\alpha p + (1-\alpha)q) \succsim_{\mathbb{V}} U(\alpha p' + (1-\alpha)q) \\ &\iff \alpha p + (1-\alpha)q \succsim_X \alpha p' + (1-\alpha)q. \end{aligned}$$

The first and last biconditionals hold because U is a representation of $\succsim_X$. The second biconditional is immediate from the definition of 'preordered vector space', and the third follows from the fact that U is mixture preserving.

Conversely, suppose that $\succsim_X$ satisfies (I_3). Let $\mathbb{V}$ be the vector space of finite signed measures on Y (it is the span of the set of probability measures). For each measurable $A \subset Y$, define $F_A: \mathbb{V} \to \mathbb{R}$ by $F_A(p) = p(A)$. Let $\mathbb{V}'$ be the span of all such F_A. Then $\mathbb{V}'$ is a separating vector space of linear functionals on $\mathbb{V}$.

Let U be the inclusion of $\mathcal{P}(Y)$ into $\mathbb{V}$. Define $u: Y \to \mathbb{V}$ by $u(y) = 1_y$. For each $F_A \in \mathbb{V}'$, $(F_A \circ u)(y) = F_A(1_y) = 1_y(A)$, so $F_A \circ u$ is the characteristic function of A. Let $\Lambda \in \mathbb{V}'$. We may write $\Lambda = \sum_{i=1}^n \lambda_i F_{A_i}$ for some $\lambda_i \in \mathbb{R}$ and measurable $A_i \subset Y$. Therefore, for each $p \in \mathcal{P}(Y)$, we have $\int_Y \Lambda \circ u \, dp = \sum_{i=1}^n \lambda_i \int_Y F_{A_i} \circ u \, dp = \sum_{i=1}^n \lambda_i p(A_i) = \sum_{i=1}^n \lambda_i F_{A_i}(p) = (\Lambda \circ U)(p)$. Therefore, u is weakly $\mathcal{P}(Y)$-integrable with respect to $\mathbb{V}'$, and $\int_Y u \, dp = U(p)$. It only remains to define a linear preorder on $\mathbb{V}$ making U into a representation of $\succsim_X$.

Define $C \subset \mathbb{V}$ by $C := \{\lambda(q - q') : \lambda > 0;\, q, q' \in \mathcal{P}(Y);\, q \succsim_X q'\}$. Define a binary relation $\succsim_{\mathbb{V}}$ on $\mathbb{V}$ by $v \succsim_{\mathbb{V}} v' \iff v - v' \in C$. It is well known, and easy to check, that this construction makes $\mathbb{V}$ into a preordered vector space. The only non-trivial claim is that $\succsim_{\mathbb{V}}$ is transitive, which follows from $\succsim_X$ satisfying (I_3). Indeed, suppose that $v \succsim_{\mathbb{V}} v' \succsim_{\mathbb{V}} v''$. We can then write $v - v' = \lambda_1(q_1 - q_1')$ and $v' - v'' = \lambda_2(q_2 - q_2')$ for some $\lambda_1, \lambda_2 > 0$, $q_1 \succsim_X q_1'$, and $q_2 \succsim_X q_2'$. Setting $\alpha = \lambda_1/(\lambda_1 + \lambda_2)$, straightforward rearrangement shows

$$\begin{aligned} v - v'' &= (v - v') + (v' - v'') \\ &= 2(\lambda_1 + \lambda_2)\left[\left(\frac{\alpha q_1 + (1-\alpha)q_2}{2} + \frac{\alpha q_1' + (1-\alpha)q_2'}{2} \right) \right. \\ &\quad \left. - \left(\alpha q_1' + (1-\alpha) q_2' \right) \right]. \end{aligned}$$

Now two applications of (I_3) show $q_3 := \alpha q_1 + (1-\alpha)q_2 \succsim_X \alpha q_1' + (1-\alpha)q_2 \succsim_X \alpha q_1' + (1-\alpha)q_2' =: q_3'$, and another application shows $q_4 := q_3/2 + q_3'/2 \succsim_X q_3'$. However, the displayed equation says $v - v'' = 2(\lambda_1 + \lambda_2)(q_4 - q_3')$, so $v - v''$ is in C, so $v \succsim_{\mathbb{V}} v''$.

Finally, we check that U is a representation of $\succsim_X$. Since U is the inclusion of $\mathcal{P}(Y)$ into $\mathbb{V}$, the claim is just that $p \succsim_X p' \iff p \succsim_{\mathbb{V}} p'$. First, $p \succsim_X p' \implies p - p' \in C \implies p \succsim_{\mathbb{V}} p'$. Conversely, suppose $p \succsim_{\mathbb{V}} p'$. Then there must be $\lambda > 0$ and $q, q' \in \mathcal{P}(Y)$ with $q \succsim_X q'$ and $p - p' = \lambda(q - q')$. Let $\alpha := \frac{1}{1+\lambda}$. Then $\alpha p + (1-\alpha)q' = \alpha p' + (1-\alpha)q$. This, together with the fact that $\succsim_X$ satisfies (I_3), yields $q \succsim_X q' \implies \alpha p' + (1-\alpha)q \succsim_X \alpha p' + (1-\alpha)q' \implies \alpha p + (1-\alpha)q' \succsim_X \alpha p' + (1-\alpha)q' \implies p \succsim_X p'$, as desired. □

Proof of Theorem 4.4. Part (i) follows from Proposition 4.2 and Lemma 4.3. However, a more explicit argument is useful. Suppose given a Vector EU representation U of $\succsim_{\mathbb{P}}$. Then the formula in part (ii) will define a Vector EU representation of $\succsim$ (and similarly for part (iii) in the variable population case). Note that the representation V has values in the same space as U. Just before Lemma 4.3, we explained how both ordinary and Multi EU representations can be identified with Vector EU representations whose value spaces $\mathbb{V}$ have certain forms. So with these identifications in mind, if U is an ordinary EU representation, V will be one too, and if U is a Multi EU representation, so is V. Conversely, given a Vector EU representation V of $\succsim$, with $V(L) = \int_{\mathbb{D}} v \, dL$, we claim that $U := V \circ \mathcal{L}$ is a Vector EU representation of $\succsim_{\mathbb{P}}$. (And similarly in the variable population case: if V^{v} is a Vector EU representation of $\succsim^{\mathsf{v}}$, then any $V^{\mathsf{v}} \circ \mathcal{L}_{\mathbb{I}}^{\mathsf{v}}$ is a Vector EU representation of $\succsim_{\mathbb{P}^{\mathsf{v}}}$.) That this U is a representation follows from Reduction to Prospects, which holds by Theorem 2.2. To see that it is expectational, consider any $\Lambda \in \mathbb{V}'$, where $\mathbb{V}'$ is the separating vector space of linear functional used to define the weak integral. Using, in turn, the definition of U, the definition of the weak integral, the definition of $\mathcal{L}$, and Lemma A.2, we find

$$\begin{aligned} \Lambda \circ U(P) = \Lambda \circ V \circ \mathcal{L}(P) &= \int_{\mathbb{D}} \Lambda \circ v \, d\mathcal{L}(P) \\ &= \int_{\mathbb{D}} \Lambda \circ v \, d(P \circ \mathcal{D}^{-1}) = \int_{\mathbb{W}} \Lambda \circ v \circ \mathcal{D} \, dP. \end{aligned}$$

By definition of the weak integral, we find that $U(P) = \int_{\mathbb{W}} v \circ \mathcal{D} \, dP$, showing that U is expectational. Again, since U has values in the same space as the given V, if V is (up to identification) either an ordinary EU representation or a Multi EU representation, then so is U.

The proofs of parts (ii) and (iii) are parallel, so we present only the variable population case, part (iii). We begin with a general observation. Suppose that $\succsim_{\mathbb{P}^{\mathsf{v}}}$ is represented by a function $U^{\mathsf{v}}: \mathbb{P}^{\mathsf{v}} \to \mathbb{V}$, where $(\mathbb{V}, \succsim_{\mathbb{V}})$ is a preordered vector space, normalized so that $U^{\mathsf{v}}(1_\Omega) = 0$. Suppose that U^{v} is mixture preserving. (As defined in note 28, this means that $U^{\mathsf{v}}(\alpha P + (1-\alpha)P') = \alpha U^{\mathsf{v}}(P) + (1-\alpha)U^{\mathsf{v}}(P')$ for $P, P' \in \mathbb{P}^{\mathsf{v}}$, $\alpha \in (0, 1)$.) For $L, L' \in \mathbb{L}_{\mathbb{I}}^{\mathsf{v}}$, we have

$$\begin{aligned} L \succsim^{\mathsf{v}} L' &\iff p_L^{\mathbb{I}} \succsim_{\mathbb{V}} p_{L'}^{\mathbb{I}} \\ &\qquad (\succsim_{\mathbb{P}^{\mathsf{v}}} \text{ generates } \succsim^{\mathsf{v}}) \\ &\iff U^{\mathsf{v}}(p_L^{\mathbb{I}}) \succsim_{\mathbb{V}} U^{\mathsf{v}}(p_{L'}^{\mathbb{I}}) \\ &\qquad (U^{\mathsf{v}} \text{ represents } \succsim_{\mathbb{P}^{\mathsf{v}}}) \\ &\iff \tfrac{1}{\#\mathbb{I}} \textstyle\sum_{i\in\mathbb{I}} U^{\mathsf{v}}(\mathcal{P}_i^{\mathsf{v}}(L)) \succsim_{\mathbb{V}} \tfrac{1}{\#\mathbb{I}} \textstyle\sum_{i\in\mathbb{I}} U^{\mathsf{v}}(\mathcal{P}_i^{\mathsf{v}}(L')) \\ &\qquad (U^{\mathsf{v}} \text{ is mixture preserving}) \\ &\iff \textstyle\sum_{i\in\mathbb{I}^\infty} U^{\mathsf{v}}(\mathcal{P}_i^{\mathsf{v}}(L)) \succsim_{\mathbb{V}} \textstyle\sum_{i\in\mathbb{I}^\infty} U^{\mathsf{v}}(\mathcal{P}_i^{\mathsf{v}}(L')). \end{aligned} \tag{11}$$

The last line incorporates two moves: multiplying both sides of the previous line by $\#\mathbb{I}$ (this is an order-preserving transformation of $\mathbb{V}$) and then extending the sum from $\mathbb{I}$ to $\mathbb{I}^\infty$: by Lemma 3.3(i), the additional terms are all zero.

Suppose now that $\succsim_{\mathbb{P}^{\mathsf{v}}}$ satisfies Vector EUT with respect to some $(\mathbb{V}, \succsim_{\mathbb{V}}, \mathbb{V}')$. Let $U^{\mathsf{v}}: \mathbb{P}^{\mathsf{v}} \to \mathbb{V}$ provide a Vector EU representation, so that $U^{\mathsf{v}}(P) = \int_{\mathbb{W}^{\mathsf{v}}} u \, dP$. By adding a constant to U^{v}, we may assume $U^{\mathsf{v}}(1_\Omega) = 0$. As shown in the proof of Lemma 4.3, U^{v} is mixture preserving. From (11) we therefore find that $\succsim^{\mathsf{v}}$ is represented by the function $L \mapsto \sum_{i\in\mathbb{I}^\infty} U^{\mathsf{v}}(\mathcal{P}_i^{\mathsf{v}}(L)) = \sum_{i\in\mathbb{I}^\infty} \int_{\mathbb{W}^{\mathsf{v}}} u \, d\mathcal{P}_i^{\mathsf{v}}(L)$.

To establish part (iii) of the theorem, the only thing left to prove is the identity

$$\sum_{i\in\mathbb{I}^\infty} \int_{\mathbb{W}^{\mathsf{v}}} u \, d\mathcal{P}_i^{\mathsf{v}}(L) = \int_{\mathbb{D}^{\mathsf{v}}} \sum_{i\in\mathbb{I}^\infty} (u \circ \mathcal{W}_i^{\mathsf{v}}) \, dL$$

stated in the definition of $V^{\mathsf{v}}(L)$. Again, each sum over $i \in \mathbb{I}^\infty$ can be replaced by a finite sum over $i \in \mathbb{I}$. Considering each $i \in \mathbb{I}$ separately, we are reduced to proving

$$\int_{\mathbb{W}^{\mathsf{v}}} u \, d\mathcal{P}_i^{\mathsf{v}} = \int_{\mathbb{D}^{\mathsf{v}}} u \circ \mathcal{W}_i^{\mathsf{v}} \, dL. \tag{12}$$

For any $\Lambda \in \mathbb{V}'$, Lemma A.2 yields

$$\int_{\mathbb{W}^{\mathsf{v}}} \Lambda \circ u \, d\mathcal{P}_i^{\mathsf{v}}(L) = \int_{\mathbb{W}^{\mathsf{v}}} \Lambda \circ u \, d(L \circ (\mathcal{W}_i^{\mathsf{v}})^{-1}) = \int_{\mathbb{D}^{\mathsf{v}}} \Lambda \circ u \circ \mathcal{W}_i^{\mathsf{v}} \, dL.$$

Eq. (12) then follows from the definition of the $\mathbb{V}$-valued integral. □

Proof of Proposition 4.8. We will treat the constant and variable population cases simultaneously. In either case, the equivalence between $(\mathrm{I}_i^{\mathbb{Q}})$ for the individual preorder and $(\mathrm{I}_i^{\mathbb{Q}})$ for the social preorder (or preorders) follows from Proposition 4.2. So it remains to show that $(\mathrm{S}_i) \iff (\mathrm{I}_i^{\mathbb{Q}}) \iff (\mathrm{P}_i)$ for $i = 1, 2, 3$, where $(\mathrm{I}_i^{\mathbb{Q}})$ is understood as a condition on $\succsim_{\mathbb{P}^*}$.

We first argue that $(\mathrm{S}_i) \impliedby (\mathrm{I}_i^{\mathbb{Q}}) \implies (\mathrm{P}_i)$ for $i = 1, 2, 3$. It will be sufficient to show that $(\mathrm{S}_i) \impliedby [(\mathrm{I}_\mathrm{a}^{\mathbb{Q}})\&(\mathrm{I}_i^{\mathbb{Q}})] \implies (\mathrm{P}_i)$ for $i = \mathrm{a, b, c}$. So suppose we have $(\mathrm{I}_\mathrm{a}^{\mathbb{Q}})$ and $(\mathrm{I}_i^{\mathbb{Q}})$. Let the symbol $\diamond$ stand for $\sim$, $\succ$, or $\curlywedge$, corresponding to $i = \mathrm{a, b, c}$. We claim

(D1) The antecedent of each of (S_i) and (P_i) implies $p_L^{\mathbb{K}} \sim_{\mathbb{P}^*} p_{L'}^{\mathbb{K}}$.
(D2) The antecedent of each of (S_i) and (P_i) implies $p_L^{\mathbb{J}} \diamond_{\mathbb{P}^*} p_{L'}^{\mathbb{J}}$.

Granted (D1) and (D2), we can deduce $p_L^{\mathbb{I}} \diamond_{\mathbb{P}^*} p_{L'}^{\mathbb{I}}$ by assuming the antecedent of either (S_i) or (P_i):

$$\begin{aligned} p_L^{\mathbb{I}} &= \tfrac{\#\mathbb{J}}{\#\mathbb{I}} p_L^{\mathbb{J}} + \tfrac{\#\mathbb{K}}{\#\mathbb{I}} p_L^{\mathbb{K}} \\ &\sim_{\mathbb{P}^*} \tfrac{\#\mathbb{J}}{\#\mathbb{I}} p_L^{\mathbb{J}} + \tfrac{\#\mathbb{K}}{\#\mathbb{I}} p_{L'}^{\mathbb{K}} && (\mathrm{I}_\mathrm{a}^{\mathbb{Q}}) \text{ and (D1)} \\ &\diamond_{\mathbb{P}^*} \tfrac{\#\mathbb{J}}{\#\mathbb{I}} p_{L'}^{\mathbb{J}} + \tfrac{\#\mathbb{K}}{\#\mathbb{I}} p_{L'}^{\mathbb{K}} && (\mathrm{I}_i^{\mathbb{Q}}) \text{ and (D2)} \\ &= p_{L'}^{\mathbb{I}} \end{aligned}$$

Since $\succsim_{\mathbb{P}^*}$ generates $\succsim_{\mathbb{I}}^*$, we find $L \diamond_{\mathbb{I}}^* L'$, validating both (S_i) and (P_i). It remains to prove (D1) and (D2).

Suppose the antecedent of (S_i) is satisfied, so that $L|_{\mathbb{K}} \sim_{\mathbb{K}}^* L'|_{\mathbb{K}}$. Then $p_L^{\mathbb{K}} = p_{L|_{\mathbb{K}}}^{\mathbb{K}} \sim_{\mathbb{P}^*} p_{L'|_{\mathbb{K}}}^{\mathbb{K}} = p_{L'}^{\mathbb{K}}$, as claimed by (D1). Similar reasoning shows $p_L^{\mathbb{J}} = p_{L|_{\mathbb{J}}}^{\mathbb{J}} \diamond_{\mathbb{P}^*} p_{L'|_{\mathbb{J}}}^{\mathbb{J}} = p_{L'}^{\mathbb{J}}$, as claimed by (D2).

Suppose instead that the antecedent of (P_i) is satisfied, so that $L \approx_{\mathbb{P}^*}^{\mathbb{K}} L'$. This means that $\mathcal{P}_k^*(L) \sim_{\mathbb{P}^*} \mathcal{P}_k^*(L')$ for all $k \in \mathbb{K}$. We obtain $p_L^{\mathbb{K}} \sim_{\mathbb{P}^*} p_{L'}^{\mathbb{K}}$, as claimed by (D1), by repeatedly applying $(\mathrm{I}_\mathrm{a}^{\mathbb{Q}})$. If $i = \mathrm{a}$ or $i = \mathrm{b}$, then $p_L^{\mathbb{J}} \diamond_{\mathbb{P}^*} p_{L'}^{\mathbb{J}}$ follows by a similar method. The case $i = \mathrm{c}$ is slightly more complicated. Choose any $j \in \mathbb{J}$. Since $\mathcal{P}_k^*(L) \sim_{\mathbb{P}^*} \mathcal{P}_j^*(L)$ for any other $k \in \mathbb{J}$, we can deduce $p_L^{\mathbb{J}} \sim_{\mathbb{P}^*} \mathcal{P}_j^*(L)$ by repeatedly applying $(\mathrm{I}_\mathrm{a}^{\mathbb{Q}})$. Similarly, $p_{L'}^{\mathbb{J}} \sim_{\mathbb{P}^*} \mathcal{P}_j^*(L')$. Since $\mathcal{P}_j^*(L) \curlywedge_{\mathbb{P}^*} \mathcal{P}_j^*(L')$, we obtain $p_L^{\mathbb{J}} \curlywedge_{\mathbb{P}^*} p_{L'}^{\mathbb{J}}$, completing the proof of (D2).

We now argue that $(\mathrm{S}_i) \implies (\mathrm{I}_i^{\mathbb{Q}}) \impliedby (\mathrm{P}_i)$ for $i = 1, 2, 3$, and indeed for each of $i = \mathrm{a, b, c}$. Suppose given $p, p', q \in \mathbb{P}^*$ and $\alpha \in (0, 1) \cap \mathbb{Q}$. Let $P := \alpha p + (1-\alpha)q$, $P' := \alpha p' + (1-\alpha)q$. It suffices to show that, given $i \in \{\mathrm{a, b, c}\}$, each of (S_i) and (P_i) implies the conditional $p \diamond_{\mathbb{P}^*} p' \implies P \diamond_{\mathbb{P}^*} P'$.

Let $\mathbb{J}, \mathbb{K} \subset \mathbb{I}^\infty$ be finite with $\mathbb{J} \cap \mathbb{K} = \varnothing$ such that $\frac{\#\mathbb{J}}{\#\mathbb{K}} = \frac{\alpha}{1-\alpha}$. Let $\mathbb{I} := \mathbb{J} \cup \mathbb{K}$. We first specialize to the case of a family $\mathbb{F}$ of constant population models. Given the hypothesis that $\mathbb{F}$ is compositional, we can find $L, L' \in \mathbb{L}_{\mathbb{I}}$ such that $\mathcal{P}_j(L) = p$ and $\mathcal{P}_j(L') = p'$ for all $j \in \mathbb{J}$, and $\mathcal{P}_k(L) = \mathcal{P}_k(L') = q$ for all $k \in \mathbb{K}$. Then $p_L^{\mathbb{I}} = P$ and $p_{L'}^{\mathbb{I}} = P'$, while $p_L^{\mathbb{J}} = p$, $p_{L'}^{\mathbb{J}} = p'$, and $p_L^{\mathbb{K}} = p_{L'}^{\mathbb{K}} = q$. We claim that the antecedent of each of (P_i) and (S_i) holds if and only if $p \diamond_{\mathbb{P}} p'$; for (P_i) this is immediate, while for (S_i) it holds because $\succsim_{\mathbb{P}}$ generates $\succsim_{\mathbb{J}}$ and $\succsim_{\mathbb{K}}$. In addition, the consequent of each of (P_i) and (S_i) holds if and only if $P \diamond_{\mathbb{P}} P'$; this follows from the fact that $\succsim_{\mathbb{P}}$ generates $\succsim_{\mathbb{I}}$. Therefore, as we wanted to show, each of (S_i) and (P_i) yields the implication $p \diamond_{\mathbb{P}} p' \implies P \diamond_{\mathbb{P}} P'$.

Turning now to a variable population model $\mathbb{M}$, we can argue in a similar way, but now defining $L := \frac{1}{2}\mathcal{L}_{\mathbb{J}}^{\mathrm{v}}(p) + \frac{1}{2}\mathcal{L}_{\mathbb{K}}^{\mathrm{v}}(q)$ and $L' := \frac{1}{2}\mathcal{L}_{\mathbb{J}}^{\mathrm{v}}(p') + \frac{1}{2}\mathcal{L}_{\mathbb{K}}^{\mathrm{v}}(q)$. In this case, $\mathcal{P}_j^{\mathrm{v}}(L) = \frac{1}{2}p + \frac{1}{2}1_\Omega = p_L^{\mathbb{J}}$ and $\mathcal{P}_j^{\mathrm{v}}(L') = \frac{1}{2}p' + \frac{1}{2}1_\Omega = p_{L'}^{\mathbb{J}}$ for all $j \in \mathbb{J}$, and $\mathcal{P}_k^{\mathrm{v}}(L) = \mathcal{P}_k^{\mathrm{v}}(L') = \frac{1}{2}q + \frac{1}{2}1_\Omega = p_L^{\mathbb{K}} = p_{L'}^{\mathbb{K}}$ for all $k \in \mathbb{K}$; moreover, $p_L^{\mathbb{I}} = \frac{1}{2}P + \frac{1}{2}1_\Omega$ and $p_{L'}^{\mathbb{I}} = \frac{1}{2}P' + \frac{1}{2}1_\Omega$. As before, the antecedent of each of (P_i) and (S_i) holds if and only if $p \diamond_{\mathbb{P}^{\mathrm{v}}} p'$; this now uses Omega Independence. In addition, the consequent of each of (P_i) and (S_i) holds if and only if $P \diamond_{\mathbb{P}^{\mathrm{v}}} P'$; this also uses Omega Independence. Therefore, as we wanted to show, each of (S_i) and (P_i) yields the implication $p \diamond_{\mathbb{P}^{\mathrm{v}}} p' \implies P \diamond_{\mathbb{P}^{\mathrm{v}}} P'$. □

Proof of Lemma 4.9. The proof is exactly the same as the proof of Lemma 4.3, with '(I_3)' replaced by '$(\mathrm{I}_3^{\mathbb{Q}})$', 'Vector EU' replaced by 'Rational Vector EU', 'preordered vector space' replaced by '$\mathbb{Q}$-preordered vector space', 'linear preorder' replaced by '$\mathbb{Q}$-linear preorder', and the coefficients $\alpha, \lambda, \lambda_1, \lambda_2$ restricted to rational numbers throughout. □

Proof of Theorem 4.10. We consider the variable population case, the constant population case being parallel. For part (iii) of the theorem, Full Pareto implies Anteriority and Reduction to Prospects as special cases. So, by the aggregation Theorem 3.5, Full Pareto and Two-Stage Anonymity imply that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ generates $\succsim^{\mathrm{v}}$. Indeed, this shows that Full Pareto and Two-Stage Anonymity hold if and only if Full Pareto holds and $\succsim_{\mathbb{P}^{\mathrm{v}}}$ generates $\succsim^{\mathrm{v}}$. By Proposition 4.8, therefore, Full Pareto and Two-Stage Anonymity hold if and only if $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies $(\mathrm{I}_3^{\mathbb{Q}})$ and generates $\succsim^{\mathrm{v}}$. Appealing to Lemma 4.9, we find that, as claimed, Full Pareto and Two-Stage Anonymity hold if and only if $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has a Rational Vector EU representation and generates $\succsim^{\mathrm{v}}$.

The proof of part (iv) of the theorem, giving a total utility representation of $\succsim^{\mathrm{v}}$, is exactly parallel to the proof of Theorem 4.4(iii). □

Proof of Theorem 5.2. Suppose we establish statement (i) of the theorem, that $\succsim^{\mathrm{v}}$ satisfies (M) if and only if $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has an EU representation. The explicit form for V^{v} in statement (ii) was derived in Theorem 4.4. As for statement (iii), preorders with EU representations satisfy (I_3) (by Lemma 4.3), while (P_3) and (S_3) follow easily from the total expected utility form of V^{v}. (The assumption that restrictions exist is required for (S_3) to make sense.)

So it remains to establish statement (i). Suppose first that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has an EU representation. By Theorem 4.4, $\succsim^{\mathrm{v}}$ also has an EU representation V^{v} and in particular satisfies (Comp). Suppose lottery L stochastically dominates L'. By Lemma A.4 we can write both L and L' as convex combinations of delta-measures: $L = \sum_{i=1}^m \alpha_i 1_{d_i}$ and $L' = \sum_{i=1}^n \alpha_i' 1_{d_i'}$ with each $\alpha_i, \alpha_i' \in (0, 1]$ and $d_i, d_i' \in \mathbb{D}^{\mathrm{v}}$. By (Comp), we can assume $1_{d_1} \succsim^{\mathrm{v}} \cdots \succsim^{\mathrm{v}} 1_{d_m}$ and $1_{d_1'} \succsim^{\mathrm{v}} \cdots \succsim^{\mathrm{v}} 1_{d_n'}$; recombining terms as necessary, we can assume $m = n$ and each $\alpha_i = \alpha_i'$. Stochastic dominance then means that the first sum dominates the second term-wise, i.e. $1_{d_i} \succsim^{\mathrm{v}} 1_{d_i'}$, so $V^{\mathrm{v}}(1_{d_i}) \geq V^{\mathrm{v}}(1_{d_i'})$. But then $V^{\mathrm{v}}(L) = \sum_{i=1}^n \alpha_i V^{\mathrm{v}}(1_{d_i}) \geq \sum_{i=1}^n \alpha_i V^{\mathrm{v}}(1_{d_i'}) = V^{\mathrm{v}}(L')$. Therefore $L \succsim^{\mathrm{v}} L'$, and we find that $\succsim^{\mathrm{v}}$ satisfies (M).

For the converse, suppose that $\succsim^{\mathrm{v}}$ satisfies (M). In Steps 1–5 below, we show that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (I_3), and then in Step 6 use this to construct an EU representation. Suppose given $P, Q, R \in \mathbb{P}^{\mathrm{v}}$ and $\alpha \in (0, 1)$. Write $[P, R]$ for the mixture $\alpha P + (1-\alpha)R$. To establish (I_3) for $\succsim_{\mathbb{P}^{\mathrm{v}}}$, we want to show that $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q \iff [P, R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q, R]$.

Step 1. According to Lemma A.4, we can write P and Q as convex combinations of delta-measures. Suppose for a first step that the coefficients of these delta-measures are rational numbers. It follows that, for some common denominator N, any population $\mathbb{I}$ of size N, and some $v_i, w_i \in \mathbb{W}^{\mathrm{v}}$, we can write

$$P = \frac{1}{N}\sum_{i \in \mathbb{I}} 1_{v_i} \quad Q = \frac{1}{N}\sum_{i \in \mathbb{I}} 1_{w_i}.$$

By hypothesis there exists some $d_P \in \mathbb{D}_{\mathbb{I}}^{\mathrm{v}}$ with $\mathcal{W}_i^{\mathrm{v}}(d_P) = v_i$ for all $i \in \mathbb{I}$, and therefore a lottery $L_P := 1_{d_P}$ with $p_{L_P}^{\mathbb{I}} = P$. Similarly for Q.

Since, by hypothesis, $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is complete, either $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$ or $Q \succsim_{\mathbb{P}^{\mathrm{v}}} P$; since $\succsim_{\mathbb{P}^{\mathrm{v}}}$ generates $\succsim^{\mathrm{v}}$, this means that either $L_P \succsim^{\mathrm{v}} L_Q$ or $L_Q \succsim^{\mathrm{v}} L_P$. Since L_P and L_Q are delta-measures, $[L_P, \mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(R)]$ stochastically dominates $[L_Q, \mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(R)]$ if $L_P \succsim^{\mathrm{v}} L_Q$, and vice versa if $L_Q \succsim^{\mathrm{v}} L_P$. Applying (M), we find

$$[L_P, \mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(R)] \succsim^{\mathrm{v}} [L_Q, \mathcal{L}_{\mathbb{I}}^{\mathrm{v}}(R)] \iff L_P \succsim^{\mathrm{v}} L_Q.$$

Now, $p^{\mathbb{I}}_{[L_P,\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(R)]} = [P,R]$ and $p^{\mathbb{I}}_{[L_Q,\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(R)]} = [Q,R]$, and we already have $p^{\mathbb{I}}_{L_P} = P$ and $p^{\mathbb{I}}_{L_Q} = Q$. Since $\succsim_{\mathbb{P}^{\mathrm{v}}}$ generates $\succsim^{\mathrm{v}}$, these equations yield in addition the first and third biconditionals:

$$[P,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q,R] \iff [L_P,\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(R)] \succsim^{\mathrm{v}} [L_Q,\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(R)] \iff L_P \succsim^{\mathrm{v}} L_Q \iff P \succsim_{\mathbb{P}^{\mathrm{v}}} Q.$$

This establishes (I$_3$) for $\succsim_{\mathbb{P}^{\mathrm{v}}}$ under the restriction that P and Q have rational coefficients.

Step 2. Suppose now that P,Q are general—that is, they are arbitrary finitely supported prospects. In this step, we show as a preliminary that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is upper-measurable and satisfies (M). If x is any welfare state, then

$$U_x := \{y \in \mathbb{W}^{\mathrm{v}} : 1_y \succsim_{\mathbb{P}^{\mathrm{v}}} 1_x\} = (\mathcal{D}^{\mathrm{v}}_{\mathbb{I}})^{-1}(\{d \in \mathbb{D}^{\mathrm{v}} : 1_d \succsim^{\mathrm{v}} 1_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(x)}\}) = (\mathcal{D}^{\mathrm{v}}_{\mathbb{I}})^{-1}(U_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(x)}) \tag{13}$$

for any finite non-empty population $\mathbb{I}$. Here we use Reduction to Prospects, which follows from the fact that $\succsim^{\mathrm{v}}$ is generated by $\succsim_{\mathbb{P}^{\mathrm{v}}}$. Since $\succsim^{\mathrm{v}}$ is upper-measurable by hypothesis, Eq. (13) presents U_x as the inverse image of a measurable set by a measurable function; therefore U_x is measurable, and $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is upper-measurable.

To show that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (M), suppose that $P \succsim^{\mathrm{SD}}_{\mathbb{P}^{\mathrm{v}}} Q$. We have to show that $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$, and that if $P \succ^{\mathrm{SD}}_{\mathbb{P}^{\mathrm{v}}} Q$ then $P \succ_{\mathbb{P}^{\mathrm{v}}} Q$.

Let $\mathbb{I} \subset \mathbb{I}^{\infty}$ be finite and nonempty. First, we claim that $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)$ stochastically dominates $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$. If so, it follows from (M) for $\succsim^{\mathrm{v}}$ that $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P) \succsim^{\mathrm{v}} \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$, and by Reduction to Prospects that $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$, as desired. Since $\succsim^{\mathrm{v}}$ is upper-measurable, the claim is that $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)(U_d) \geq \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)(U_d)$ for all $d \in \mathbb{D}^{\mathrm{v}}$. Fix d. By definition of $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}$, $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)(U_d) = P(A)$ where

$$A := (\mathcal{D}^{\mathrm{v}}_{\mathbb{I}})^{-1}(U_d) = \{w \in \mathbb{W}^{\mathrm{v}} \colon 1_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(w)} \succsim^{\mathrm{v}} 1_d\}.$$

Similarly, $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)(U_d) = Q(A)$. The claim, then, is that $P(A) \geq Q(A)$. To show this, let S_Q be a finite set supporting Q. If $S_Q \cap A = \emptyset$, then $Q(A) = 0$, so $P(A) \geq Q(A)$ as claimed. Otherwise, since $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is complete, there is a minimal element s of $S_Q \cap A$, in the sense that $w \in S_Q \cap A \implies 1_w \succsim_{\mathbb{P}^{\mathrm{v}}} 1_s$. Now, if $1_w \succsim_{\mathbb{P}^{\mathrm{v}}} 1_s$, then, by Reduction to Prospects, $1_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(w)} \succsim^{\mathrm{v}} 1_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(s)} \succsim^{\mathrm{v}} 1_d$; this shows that $U_s \subset A$. Therefore $P(A) \geq P(U_s)$; since $P \succsim^{\mathrm{SD}}_{\mathbb{P}^{\mathrm{v}}} Q$, $P(U_s) \geq Q(U_s)$; and since $U_s \cap S_Q = A \cap S_Q$, Lemma A.3 gives $Q(U_s) = Q(A)$. Therefore $P(A) \geq Q(A)$, as claimed. So we conclude that $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)$ stochastically dominates $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$ and $P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$.

Now suppose that, more strongly, $P \succ^{\mathrm{SD}}_{\mathbb{P}^{\mathrm{v}}} Q$. That is, $P \succsim^{\mathrm{SD}}_{\mathbb{P}^{\mathrm{v}}} Q$ but $P(U_x) > Q(U_x)$ for some $x \in \mathbb{W}^{\mathrm{v}}$. By (13) and the definition of $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}$, we find $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)(U_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(x)}) > \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)(U_{\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(x)})$. The previous argument showed that $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)$ stochastically dominates $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$; this strict inequality shows that the domination is strict, i.e. $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$ does not stochastically dominate $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)$. By (M) for $\succsim^{\mathrm{v}}$, this means $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P) \succ^{\mathrm{v}} \mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(Q)$, and by Reduction to Prospects, $P \succ_{\mathbb{P}^{\mathrm{v}}} Q$. This establishes (M) for $\succsim_{\mathbb{P}^{\mathrm{v}}}$.

Step 3. Now we claim we can find a sequence (P_i) in $\mathbb{P}^{\mathrm{v}}$ strongly converging to P such that each P_i has rational coefficients and stochastically dominates P. To see this, using Lemma A.4 write P as a sum of delta-measures, $P = \alpha_1 1_{v_1} + \cdots + \alpha_n 1_{v_n}$. Since by hypothesis $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is complete, we can assume $1_{v_1} \succsim_{\mathbb{P}^{\mathrm{v}}} 1_{v_2} \succsim_{\mathbb{P}^{\mathrm{v}}} \cdots \succsim_{\mathbb{P}^{\mathrm{v}}} 1_{v_n}$. In the simplest case, $n = 1$, and then we can take $P_i := P$ for all i. For $n > 1$, let P' be the prospect $P' = \frac{\alpha_1}{1-\alpha_n} 1_{v_1} + \cdots + \frac{\alpha_{n-1}}{1-\alpha_n} 1_{v_{n-1}}$, so that $P = (1-\alpha_n)P' + \alpha_n 1_{v_n}$. By induction on n, we can find a sequence (P'_i) of prospects with rational values, each stochastically dominating P', strongly converging to P'. Let (β_i) be a sequence from $[0,1] \cap \mathbb{Q}$ approaching α_n from below. Then it is easy to check that the sequence of prospects given by $P_i := (1-\beta_i)P'_i + \beta_i 1_{v_n}$ has the required properties.

Step 4. By a similar construction, we can find a sequence (Q_i) strongly converging to Q such that each Q_i has rational values and Q stochastically dominates each Q_i.

Step 5. Since, as we proved in Step 2, $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (M), $P_i \succsim_{\mathbb{P}^{\mathrm{v}}} P$ and $Q \succsim_{\mathbb{P}^{\mathrm{v}}} Q_i$. Using this, the result for rational-coefficient prospects in Step 1, and strong continuity (applied once for $P_i \xrightarrow{s} P$ and a second time for $Q_i \xrightarrow{s} Q$), we have

$$P \succsim_{\mathbb{P}^{\mathrm{v}}} Q \implies \forall ij.P_i \succsim_{\mathbb{P}^{\mathrm{v}}} Q_j \iff \forall ij.[P_i,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q_j,R] \implies [P,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q,R].$$

To complete the derivation of (I$_3$), it suffices to show that the first and last implications displayed are reversible. For the first one, strong continuity yields $\forall ij.P_i \succsim_{\mathbb{P}^{\mathrm{v}}} Q_j \implies P \succsim_{\mathbb{P}^{\mathrm{v}}} Q$. For the last one, for each i, $[P_i,R]$ stochastically dominates $[P,R]$, so, by (M), $[P_i,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [P,R]$. Similarly, for any j, $[Q,R]$ stochastically dominates $[Q_j,R]$, so $[Q,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q_j,R]$. Therefore, if $[P,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q,R]$, we must also have $[P_i,R] \succsim_{\mathbb{P}^{\mathrm{v}}} [Q_j,R]$ for any i,j. This establishes (I$_3$).

Step 6. It remains to show that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has an EU representation. We first show that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (MC). Suppose that β is a limit point of $\{\alpha \in [0,1] : \alpha P + (1-\alpha)R \succsim_{\mathbb{P}^{\mathrm{v}}} Q\}$. Then there is a sequence (β_n) in $[0,1]$ converging to β with $\beta_n P + (1-\beta_n)R \succsim_{\mathbb{P}^{\mathrm{v}}} Q$. It is clear that $\beta_n P + (1-\beta_n)R$ converges strongly to $\beta P + (1-\beta)R$, so by strong continuity, $\beta P + (1-\beta)R \succsim_{\mathbb{P}^{\mathrm{v}}} Q$, implying that $\{\alpha \in [0,1] : \alpha P + (1-\alpha)R \succsim_{\mathbb{P}^{\mathrm{v}}} Q\}$ is closed. A similar argument shows that $\{\alpha \in [0,1] : Q \succsim_{\mathbb{P}^{\mathrm{v}}} \alpha P + (1-\alpha)R\}$ is closed. Therefore $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (MC).

Given that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ satisfies (I$_3$), (MC), and (Comp), the main result of Herstein and Milnor (1953, Theorem 8) is that $\succsim_{\mathbb{P}^{\mathrm{v}}}$ has a mixture-preserving representation $U^{\mathrm{v}} \colon \mathbb{P}^{\mathrm{v}} \to \mathbb{R}$. Set $u(y) = U^{\mathrm{v}}(1_y)$ for any $y \in \mathbb{W}^{\mathrm{v}}$. We want to show that, for any $P \in \mathbb{P}^{\mathrm{v}}$, $U^{\mathrm{v}}(P) = \int_{\mathbb{W}^{\mathrm{v}}} u \,\mathrm{d}P$. We can again use Lemma A.4 to write P in the form $P = \sum \alpha_i 1_{v_i}$. Since U^{v} is mixture preserving, we have $U^{\mathrm{v}}(P) = \sum \alpha_i u(v_i)$. It remains to show that $u(v_i) = \int_{\mathbb{W}^{\mathrm{v}}} u \,\mathrm{d}1_{v_i}$, which is automatic if u is measurable. To show that u is measurable, it suffices to show that $A_x := u^{-1}([x,\infty))$ is a measurable subset of $\mathbb{W}^{\mathrm{v}}$, for all $x \in \mathbb{R}$. First, if $u(v) < x$ for all $v \in \mathbb{W}^{\mathrm{v}}$, then $A_x = \emptyset$ is measurable. Second, if $\{u(w) : w \in \mathbb{W}^{\mathrm{v}}, u(w) \geq x\}$ has a minimal element $u(w)$, then A_x is the upper set U_w, and so is measurable since (as proved in Step 2) $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is upper-measurable. Otherwise, choose a sequence (w_i) in $\mathbb{W}^{\mathrm{v}}$ with $u(w_i) \geq x$ and $u(w_i)$ converging to $\inf\{u(w) : w \in \mathbb{W}^{\mathrm{v}}, u(w) \geq x\}$. Then A_x is the countable union of the upper sets U_{w_i}, and therefore measurable. □

Proof of Lemma 5.4. Let $U \colon \mathcal{P}(Y) \to \mathbb{R}$. Suppose first that U is integrally Gâteaux differentiable at $p \in \mathcal{P}(Y)$. In other words, there exists some $v_p \in \nabla U_p$. We claim that there is at least one $u_p \in \nabla U_p$ satisfying

$$U(p) = \int_Y u_p \,\mathrm{d}p. \tag{14}$$

Using the fact that $\mathcal{P}(Y)$ consists of probability measures, it is easy to check that ∇U_p is closed under the addition of constant functions; thus $u_p := v_p + U(p) - \int_Y v_p \,\mathrm{d}p$ is also in ∇U_p. By integrating both sides with respect to p, we find that u_p satisfies (14). We conclude that U is integrally Gâteaux differentiable at $p \in \mathcal{P}(Y)$ if and only if there exists $u_p \in \nabla U_p$ satisfying (14). To prove the lemma, it remains to show that u_p is in ∇U_p, and satisfies (14), if and only if it is a local utility function for U at p, in the sense of satisfying (5).

Suppose given any $u_p \in \nabla U_p$ satisfying (14). Being in ∇U_p means that $U'_p(q-p) = \int_Y u_p \,\mathrm{d}(q-p)$ for all $q \in \mathcal{P}(Y)$. By definition of $U'_p(q-p)$, this is equivalent to

$$U(p + t(q-p)) = U(p) + t \int_Y u_p \,\mathrm{d}(q-p) + o(t) \text{ as } t \to 0^+. \tag{15}$$

Using (14), we obtain Eq. (5).

Conversely, suppose u_p satisfies (5) for all $q \in \mathcal{P}(Y)$. Putting $t \to 0^+$ in (5), we recover (14). Together, (5) and (14) entail (15). As in the previous paragraph, (15) means that u_p is in ∇U_p. □

Proof of Theorem 5.5. For the right to left direction of part (i) of the theorem, suppose that $\succsim$ satisfies Local EUT. In particular, suppose that $V:\mathbb{L}\to\mathbb{R}$ represents $\succsim$ and is locally expectational on $\mathbb{L}$. Since $\succsim_{\mathbb{P}}$ generates $\succsim$, Reduction to Prospects holds, by the aggregation theorem Theorem 2.2. It follows that $U:=V\circ\mathcal{L}$ is a representation of $\succsim_{\mathbb{P}}$.

It remains to show that U is locally expectational. Fix $P\in\mathbb{P}$. By Lemma 5.4, V is integrally Gâteaux differentiable at $\mathcal{L}(P)$; that is, $\nabla V_{\mathcal{L}(P)}\neq\varnothing$. Since the map $P\mapsto\mathcal{L}(P)$ is mixture-preserving, $\mathcal{L}(P+t(Q-P))=\mathcal{L}(P)+t(\mathcal{L}(Q)-\mathcal{L}(P))$, so $U(P+t(Q-P))=V(\mathcal{L}(P)+t(\mathcal{L}(Q)-\mathcal{L}(P)))$. Applying the definition (6) of the Gâteaux derivative, we find

$$U'_P(Q-P)=V'_{\mathcal{L}(P)}(\mathcal{L}(Q)-\mathcal{L}(P)). \tag{16}$$

Thus U is Gâteaux differentiable at P. Now fix $v_P\in\nabla V_{\mathcal{L}(P)}$. For any $Q\in\mathbb{P}$, v_P is integrable with respect to $\mathcal{L}(Q)=Q\circ\mathcal{D}^{-1}$. By Lemma A.2, $v_P\circ\mathcal{D}$ is integrable with respect to Q, for any Q, and is hence $\mathbb{P}$-integrable. Lemma A.2 also gives

$$\int_{\mathbb{D}} v_P\,\mathrm{d}\mathcal{L}(Q)=\int_{\mathbb{W}} v_P\circ\mathcal{D}\,\mathrm{d}Q.$$

Combining this with (16) we find

$$U'_P(Q-P)=V'_{\mathcal{L}(P)}(\mathcal{L}(Q)-\mathcal{L}(P))=\int_{\mathbb{D}} v_P\,\mathrm{d}(\mathcal{L}(Q)-\mathcal{L}(P))$$
$$=\int_{\mathbb{W}} v_P\circ\mathcal{D}\,\mathrm{d}(Q-P).$$

Thus U is integrally Gâteaux differentiable at P. By Lemma 5.4, U is locally expectational at P, as claimed.

Conversely, suppose $\succsim_{\mathbb{P}}$ satisfies Local EUT, with a Local EU representation $U:\mathcal{P}(Y)\to\mathbb{R}$. Note that $\#\mathbb{I}U$ also represents $\succsim_{\mathbb{P}}$. Since $\succsim_{\mathbb{P}}$ generates $\succsim$, $\succsim$ is therefore represented by

$$L\mapsto V(L):=\#\mathbb{I}U(p_L).$$

We want to show that V is locally expectational.

Fix $L\in\mathbb{L}$. By Lemma 5.4, U is integrally Gâteaux differentiable at p_L; that is, $\nabla U_{p_L}\neq\varnothing$. Since the map $L\mapsto p_L$ is mixture-preserving, $p_{L+t(M-L)}=p_L+t(p_M-p_L)$, so $V(L+t(M-L))=\#\mathbb{I}U(p_L+t(p_M-p_L))$. Applying the definition (6) of the Gâteaux derivative, we find

$$V'_L(M-L)=\#\mathbb{I}U'_{p_L}(p_M-p_L). \tag{17}$$

Thus V is Gâteaux differentiable at L.

Fix $u_L\in\nabla U_{p_L}$. For any $M\in\mathbb{L}$ and $i\in\mathbb{I}$, u_L is integrable with respect to $\mathcal{P}_i(M)=M\circ\mathcal{W}_i^{-1}$. Using Lemma A.2, we find that $u_L\circ\mathcal{W}_i$ is integrable with respect to M, implying that $\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i$ is $\mathbb{L}$-integrable, and also that

$$\#\mathbb{I}\int_{\mathbb{W}} u_L\,\mathrm{d}p_M=\int_{\mathbb{W}}\sum_{i\in\mathbb{I}}u_L\,\mathrm{d}(\mathcal{P}_i(M))=\int_{\mathbb{D}}\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i\,\mathrm{d}M. \tag{18}$$

Combining this with (17), we find

$$V'_L(M-L)=\#\mathbb{I}U'_{p_L}(p_M-p_L)=\#\mathbb{I}\int_{\mathbb{W}} u_L\,\mathrm{d}(p_M-p_L)$$
$$=\int_{\mathbb{D}}\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i\,\mathrm{d}(M-L)$$

so V is integrally Gâteaux differentiable at L, with $\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i\in\nabla V_L$. By Lemma 5.4, V is locally expectational at L. This establishes the left-to-right direction in part (i) of the theorem, and indeed establishes the more specific claim of part (ii).

For part (iii), suppose that u_L is a local utility function for U at p_L. By Lemma 5.4, this means $u_L\in\nabla U_{p_L}$ and $U(p_L)=\int_{\mathbb{W}}u_L\,\mathrm{d}p_L$. We then have $V(L)=\#\mathbb{I}U(p_L)=\#\mathbb{I}\int_{\mathbb{W}}u_L\,\mathrm{d}p_L=\int_{\mathbb{D}}\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i\,\mathrm{d}L$, using (18) at the last step. Using Lemma 5.4 again, we find that $\sum_{i\in\mathbb{I}}u_L\circ\mathcal{W}_i$ is a local utility function for V at L. □

The next result is used in the proofs of Lemma 5.6 and Theorem 5.7. Recall the notation $P_\alpha:=\alpha P+(1-\alpha)1_\Omega$ for any $P\in\mathbb{P}^{\mathrm{v}}$ and $\alpha\in[0,1]$.

Lemma A.7. *Suppose $U^{\mathrm{v}}:\mathbb{P}^{\mathrm{v}}\to\mathbb{R}$ is Omega-linear. Fix $P,Q\in\mathbb{P}^{\mathrm{v}}$ and suppose that there is a $\mathbb{P}^{\mathrm{v}}$-integrable function u^{v} such that*

$$(U^{\mathrm{v}})'_P(Q-P)=\int_{\mathbb{W}^{\mathrm{v}}}u^{\mathrm{v}}\,\mathrm{d}(Q-P)\quad\textit{and}$$
$$(U^{\mathrm{v}})'_P(1_\Omega-P)=\int_{\mathbb{W}^{\mathrm{v}}}u^{\mathrm{v}}\,\mathrm{d}(1_\Omega-P). \tag{19}$$

Then, for any $\alpha\in(0,1]$ and $\beta\in[0,1]$, we have

$$(U^{\mathrm{v}})'_{P_\alpha}(Q_\beta-P_\alpha)=\int_{\mathbb{W}^{\mathrm{v}}}u^{\mathrm{v}}\,\mathrm{d}(Q_\beta-P_\alpha).$$

Proof. We first show that

$$(U^{\mathrm{v}})'_{P_\alpha}(Q_\beta-P_\alpha)=\beta(U^{\mathrm{v}})'_P(Q-P)+(\beta-\alpha)(U^{\mathrm{v}}(P)-U^{\mathrm{v}}(1_\Omega)) \tag{20}$$

given that, by hypothesis, the derivative on the right-hand side exists.

Suppose first that $\beta=0$. This reduces (20) to

$$(U^{\mathrm{v}})'_{P_\alpha}(1_\Omega-P_\alpha)=-\alpha(U^{\mathrm{v}}(P)-U^{\mathrm{v}}(1_\Omega)), \tag{21}$$

which follows from a direct calculation of the Gâteaux derivative (6) using Omega-linearity of U^{v}.

Suppose instead that $\beta>0$. Set $f(t):=U^{\mathrm{v}}(P_\alpha+t(Q_\beta-P_\alpha))$, for $t\in[0,1)$. Set $x(t)=\frac{\beta t}{\alpha+t(\beta-\alpha)}$ and $R(t):=P+x(t)(Q-P)$. Since $x(t)$ approaches 0 from above as t approaches 0 from above, $R(t)$ is in $\mathbb{P}^{\mathrm{v}}$ for all t small enough. Moreover, a straightforward calculation shows $P_\alpha+t(Q_\beta-P_\alpha)=R(t)_{\alpha+t(\beta-\alpha)}$.

Therefore, by Omega-linearity (8),

$$f(t)=U^{\mathrm{v}}(R(t)_{\alpha+t(\beta-\alpha)})=(\alpha+t(\beta-\alpha))U^{\mathrm{v}}(R(t))$$
$$+(1-(\alpha+t(\beta-\alpha)))U^{\mathrm{v}}(1_\Omega).$$

By definition, $(U^{\mathrm{v}})'_{P_\alpha}(Q_\beta-P_\alpha)$ is the partial derivative $\partial_+f(t)_{t=0}$, and by elementary calculus

$$(U^{\mathrm{v}})'_{P_\alpha}(Q_\beta-P_\alpha)=\partial_+f(t)|_{t=0}=\alpha\partial_+U^{\mathrm{v}}(R(t))|_{t=0}$$
$$+(\beta-\alpha)U^{\mathrm{v}}(R(0))-(\beta-\alpha)U^{\mathrm{v}}(1_\Omega).$$

Noting that $R(0)=P$, and comparing this with (20), it remains to establish

$$\alpha\partial_+U^{\mathrm{v}}(R(t))|_{t=0}=\beta(U^{\mathrm{v}})'_P(Q-P). \tag{22}$$

This is essentially just an application of the chain rule. To work it out in this unfamiliar setting,

$$\partial_+U^{\mathrm{v}}(R(t))|_{t=0}=\lim_{t\to0^+}\frac{U^{\mathrm{v}}(R(t))-U^{\mathrm{v}}(R(0))}{t}$$
$$=\lim_{t\to0^+}\frac{U^{\mathrm{v}}(R(t))-U^{\mathrm{v}}(R(0))}{x(t)}\cdot\frac{x(t)}{t}=(U^{\mathrm{v}})'_P(Q-P)\cdot x'(0).$$

The last equation follows from the definition of $(U^{\mathrm{v}})'_P$ and the fact that $x(t)\to0^+$ as $t\to0^+$. Since $x'(0)=\beta/\alpha$, we obtain (22). This concludes the proof of (20).

Now we calculate:

$$\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(Q_\beta-P_\alpha)=\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(Q_\beta-P_\beta)+\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(P_\beta-P_\alpha)$$
$$=\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(\beta Q-\beta P)$$
$$+\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}((\beta-\alpha)P-(\beta-\alpha)1_\Omega)$$
$$=\beta\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(Q-P)+(\beta-\alpha)\int_{\mathbb{W}^{\mathrm{v}}}u\,\mathrm{d}(P-1_\Omega)$$
$$=\beta(U^{\mathrm{v}})'_P(Q-P)-(\beta-\alpha)(U^{\mathrm{v}})'_P(1_\Omega-P)$$

applying the hypotheses of this lemma to obtain the last line. By (21) with $\alpha = 1$ we have $(U^{\vee})'_P(1_\Omega - P) = -(U^{\vee}(P) - U^{\vee}(1_\Omega))$. Therefore we have found

$$\int_{\mathbb{W}^{\vee}} u \, \mathrm{d}(Q_\beta - P_\alpha) = \beta(U^{\vee})'_P(Q-P) + (\beta-\alpha)(U^{\vee}(P) - U^{\vee}(1_\Omega)).$$

And, according to (20), this equals $(U^{\vee})'_{P_\alpha}(Q_\beta - P_\alpha)$, as desired. □

Proof of Lemma 5.6. For part (i), since $\mathbb{P}^{\vee}$ extends $\mathbb{P}$, every element of $\mathbb{P}^{\vee}$ is of the form $P_\alpha := \alpha P + (1-\alpha)1_\Omega$ for some $P \in \mathbb{P}$ and $\alpha \in [0,1]$. This presentation is unique except when $\alpha = 0$, so we may define $U^{\vee}(P_\alpha) = \alpha U(P) + (1-\alpha)c$. Then $U^{\vee}$ is the unique Omega-linear extension of U that satisfies $U^{\vee}(1_\Omega) = c$.

For part (ii), suppose that U is locally expectational on $\mathbb{P}$. We want to prove that $U^{\vee}$ as defined above is locally expectational at P_α, for each $P \in \mathbb{P}$ and $\alpha \in (0,1]$.

By Lemma 5.4, there is some u in ∇U_P with $\int_{\mathbb{W}} u \, \mathrm{d}P = U(P)$. Extend it to $u^{\vee}: \mathbb{W}^{\vee} \to \mathbb{R}$ by setting $u^{\vee}(\Omega) = c$. We first show that $u^{\vee}$ is $\mathbb{P}^{\vee}$-integrable. Since $\mathbb{P}^{\vee}$ extends $\mathbb{P}$, $\mathbb{W}$ is measurable in $\mathbb{W}^{\vee}$, with $Q(\mathbb{W}) = 1$ for any $Q \in \mathbb{P}$; therefore $u^{\vee}$ is Q-integrable, with $\int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}Q = \int_{\mathbb{W}} u^{\vee} \, \mathrm{d}Q = \int_{\mathbb{W}} u \, \mathrm{d}Q$. Similarly, $\{\Omega\}$ is measurable and $\int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}1_\Omega = u^{\vee}(\Omega) = c$. Together this shows that $u^{\vee}$ is $\mathbb{P}^{\vee}$-integrable, and specifically that for any $Q \in \mathbb{P}$ and $\beta \in [0,1]$, $\int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}Q_\beta = \beta \int_{\mathbb{W}} u \, \mathrm{d}Q + (1-\beta)c$.

We now fix $Q \in \mathbb{P}$ and verify the hypotheses (19) of Lemma A.7. Since $U^{\vee}$ extends U, and since u is a local utility function for U at P, we have

$$(U^{\vee})'_P(Q-P) = U'_P(Q-P) = \int_{\mathbb{W}} u \, \mathrm{d}(Q-P) = \int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}(Q-P).$$

Next, the definition (6) of the Gâteaux derivative and Omega-linearity yield

$$\begin{aligned}(U^{\vee})'_P(1_\Omega - P) &= \lim_{t\to 0^+} \frac{U^{\vee}((1-t)P + t1_\Omega) - U^{\vee}(P)}{t} \\ &= U^{\vee}(1_\Omega) - U^{\vee}(P).\end{aligned}$$

Given that $U^{\vee}(1_\Omega) = c = \int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}1_\Omega$ and $U^{\vee}(P) = U(P) = \int_{\mathbb{W}} u \, \mathrm{d}P = \int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}P$, we conclude

$$(U^{\vee})'_P(1_\Omega - P) = \int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}(1_\Omega - P).$$

Applying Lemma A.7 we find that, for any $\beta \in [0,1]$,

$$(U^{\vee})'_{P_\alpha}(Q_\beta - P_\alpha) = \int_{\mathbb{W}^{\vee}} u^{\vee} \, \mathrm{d}(Q_\beta - P_\alpha).$$

This shows that $U^{\vee}$ is integrally Gateaux differentiable at P_α. By Lemma 5.4 it is locally expectational at P_α, as desired. □

Proof of Theorem 5.7. We first prove the right-to-left direction of part (i) of the theorem. Suppose we are given an d_Ω-linear function $V^{\vee}: \mathbb{L}^{\vee} \to \mathbb{R}$ that is locally expectational on $\mathbb{L}^{\vee} \setminus \{1_{d_\Omega}\}$. Fix finite, nonempty $\mathbb{I} \subset \mathbb{I}^\infty$. We can now follow the proof of the right to left direction of Theorem 5.5(i), with variable population objects replacing constant population ones. That is, essentially the same argument shows that $U^{\vee} = V^{\vee} \circ \mathcal{L}^{\vee}_{\mathbb{I}}$ is a representation of $\succsim_{\mathbb{P}^{\vee}}$; that, for $P, Q \in \mathbb{P}^{\vee}$ with $P \neq 1_\Omega$, we have $(U^{\vee})'_P(Q-P) = (V^{\vee})'_{\mathcal{L}^{\vee}_{\mathbb{I}}(P)}(\mathcal{L}^{\vee}_{\mathbb{I}}(Q) - \mathcal{L}^{\vee}_{\mathbb{I}}(P))$ in analogy with (16); and that, for any $v^{\vee}_P \in \nabla V^{\vee}_{\mathcal{L}^{\vee}_{\mathbb{I}}(P)}$, we have

$$(U^{\vee})'_P(Q-P) = \int_{\mathbb{W}^{\vee}} v^{\vee}_P \circ \mathcal{D}^{\vee}_{\mathbb{I}} \, \mathrm{d}(Q-P).$$

Therefore $U^{\vee}$ is integrally Gâteaux differentiable on $\mathbb{P}^{\vee} \setminus \{1_\Omega\}$, and so by Lemma 5.4 it is locally expectational there, as desired.

To complete the right-to-left direction of part (i), it remains to show that $U^{\vee}$ is Omega-linear. Since the map $P \mapsto \mathcal{L}^{\vee}_{\mathbb{I}}(P)$ is mixture preserving, we find that for any $P \in \mathbb{P}^{\vee}$ and $\alpha \in [0,1]$, $U^{\vee}(\alpha P + (1-\alpha)P) = V^{\vee}(\alpha \mathcal{L}^{\vee}_{\mathbb{I}}(P) + (1-\alpha)\mathcal{L}^{\vee}_{\mathbb{I}}(1_\Omega))$. Since $\mathcal{L}^{\vee}_{\mathbb{I}}(1_\Omega) = 1_{d_\Omega}$, and $V^{\vee}$ is d_Ω-linear, this equals $\alpha V^{\vee}(\mathcal{L}^{\vee}_{\mathbb{I}}(P)) + (1-\alpha)V^{\vee}(\mathcal{L}^{\vee}_{\mathbb{I}}(1_\Omega))$. By definition of $U^{\vee}$, this equals $\alpha U^{\vee}(P) + (1-\alpha)U^{\vee}(1_\Omega)$, so $U^{\vee}$ is Omega-linear, as desired.

Conversely, for the left-to-right direction of part (i), suppose $U^{\vee}: \mathbb{P}^{\vee} \to \mathbb{R}$ is an Omega-linear representation of $\succsim_{\mathbb{P}^{\vee}}$ that is locally expectational on $\mathbb{P}^{\vee} \setminus \{1_\Omega\}$. For $L \in \mathbb{L}^{\vee}_{\mathbb{I}}$, define

$$V^{\vee}(L) := \#\mathbb{I}U^{\vee}(p^{\mathbb{I}}_L) - \#\mathbb{I}U^{\vee}(1_\Omega)$$

as in part (ii) of the theorem. We first show that $V^{\vee}$ represents $\succsim^{\vee}$. Since $\succsim_{\mathbb{P}^{\vee}}$ generates $\succsim^{\vee}$, and $U^{\vee}$ is a representation of $\succsim_{\mathbb{P}^{\vee}}$, we have, for $L, L' \in \mathbb{L}^{\vee}_{\mathbb{I}}$,

$$\begin{aligned} L \succsim^{\vee} L' &\iff U^{\vee}(p^{\mathbb{I}}_L) \geq U^{\vee}(p^{\mathbb{I}}_{L'}) \\ &\iff \#\mathbb{I}U^{\vee}(p^{\mathbb{I}}_L) - \#\mathbb{I}U(1_\Omega) \geq \#\mathbb{I}U^{\vee}(p^{\mathbb{I}}_{L'}) - \#\mathbb{I}U^{\vee}(1_\Omega)\end{aligned}$$

as desired.

Next we show that $V^{\vee}$ is d_Ω-linear. We have to show that, for $L \in \mathbb{L}^{\vee}_{\mathbb{I}}$ and $\alpha \in [0,1]$, $V^{\vee}(\alpha L + (1-\alpha)1_{d_\Omega}) = \alpha V^{\vee}(L) + (1-\alpha)V^{\vee}(1_{d_\Omega})$. Note first that $V^{\vee}(1_{d_\Omega}) = 0$, so we need to show that $V^{\vee}(\alpha L + (1-\alpha)1_{d_\Omega}) = \alpha V^{\vee}(L)$. But

$$\begin{aligned} &V^{\vee}(\alpha L + (1-\alpha)1_{d_\Omega}) \\ &= \#\mathbb{I}U^{\vee}(\alpha p^{\mathbb{I}}_L + (1-\alpha)p^{\mathbb{I}}_{1_{d_\Omega}}) - \#\mathbb{I}U^{\vee}(1_\Omega) \\ &= \alpha\#\mathbb{I}U^{\vee}(p^{\mathbb{I}}_L) + (1-\alpha)\#\mathbb{I}U^{\vee}(1_\Omega) - \#\mathbb{I}U^{\vee}(1_\Omega) = \alpha V^{\vee}(L).\end{aligned}$$

The first step uses the definition of $V^{\vee}$ and the fact that the map $L \mapsto p^{\mathbb{I}}_L$ is mixture preserving; the second step uses the fact that $p^{\mathbb{I}}_{1_{d_\Omega}} = 1_\Omega$ and the Omega-linearity of $U^{\vee}$.

To complete the proof of the left-to-right direction of part (i) of the theorem, as well as proving part (ii), we need to show that $V^{\vee}$ is locally expectational at each $L \in \mathbb{L}^{\vee} \setminus \{1_{d_\Omega}\}$. Fix such an L for the remainder of the proof, and $\mathbb{I} \subset \mathbb{I}^\infty$ such that $L \in \mathbb{L}^{\vee}_{\mathbb{I}}$. Note that, by Lemma 5.4, $U^{\vee}$ is integrally Gâteaux differentiable on $\mathbb{P}^{\vee} \setminus \{1_\Omega\}$, and it suffices to show that $V^{\vee}$ is integrally Gâteaux differentiable at L.

We first show that $V^{\vee}$ is Gâteaux differentiable at L; that is, for any $M \in \mathbb{L}^{\vee}$, the Gâteaux derivative $(V^{\vee})'_L(M-L)$ exists. We can find $\mathbb{J} \supset \mathbb{I}$ such that both L and M are in $\mathbb{L}^{\vee}_{\mathbb{J}}$. Note that since L is in $\mathbb{L}^{\vee} \setminus \{1_{d_\Omega}\}$, $p^{\mathbb{J}}_L$ is in $\mathbb{P}^{\vee} \setminus \{1_\Omega\}$, and therefore $U^{\vee}$ is integrally Gâteaux-differentiable at $p^{\mathbb{J}}_L$. Now, we have $V^{\vee}(L + t(M-L)) = \#\mathbb{J}U^{\vee}(p^{\mathbb{J}}_L + t(p^{\mathbb{J}}_M - p^{\mathbb{J}}_L)) - \#\mathbb{J}U^{\vee}(1_\Omega)$. Applying the definition of the Gâteaux derivative, we find that $V^{\vee}$ is Gâteaux differentiable at L, and in particular,

$$(V^{\vee})'_L(M-L) = \#\mathbb{J}(U^{\vee})'_{p^{\mathbb{J}}_L}(p^{\mathbb{J}}_M - p^{\mathbb{J}}_L). \tag{23}$$

We now show that $V^{\vee}$ is *integrally* Gâteaux differentiable at L. Since $U^{\vee}$ is integrally Gâteaux-differentiable at $p^{\mathbb{I}}_L$, we may pick $u_L \in \nabla U^{\vee}_{p^{\mathbb{I}}_L}$. By the variable population domain assumption **(D$^{\vee}$)** in Section 3.1,

$$f := \sum_{i\in\mathbb{I}^\infty}(u_L \circ \mathcal{W}^{\vee}_i - u_L(\Omega))$$

is a well-defined function on $\mathbb{D}^{\vee}$. To show that $V^{\vee}$ is integrally Gâteaux differentiable at L, we show specifically that $f \in \nabla V^{\vee}_L$.

As a preliminary, let us show that f is $\mathbb{L}^{\vee}$-integrable, i.e. integrable against an arbitrary $M \in \mathbb{L}^{\vee}$. Choose nonempty, finite $\mathbb{J} \subset \mathbb{I}^\infty$ with $M \in \mathbb{L}^{\vee}_{\mathbb{J}}$. Enlarging $\mathbb{J}$ if necessary, we can assume that $\mathcal{P}^{\vee}_i(M) = 1_\Omega$ for some $i \in \mathbb{J}$ (using Lemma 3.3(i)). For any finite $\mathbb{K} \subset \mathbb{I}^\infty$, define $f_{\mathbb{K}} := \sum_{i\in\mathbb{K}}(u_L \circ \mathcal{W}^{\vee}_i - u_L(\Omega))$. In parallel to the derivation of (18) in the proof of Theorem 5.5, u_L is integrable with respect to $\mathcal{P}^{\vee}_i(M) = M \circ (\mathcal{W}^{\vee}_i)^{-1}$. Using Lemma A.2, $u_L \circ \mathcal{W}^{\vee}_i$ is integrable with respect to M, and moreover

$$\int_{\mathbb{D}^{\vee}} u_L \circ \mathcal{W}^{\vee}_i \, \mathrm{d}M = \int_{\mathbb{W}^{\vee}} u_L \, \mathrm{d}\mathcal{P}^{\vee}_i(M).$$

This implies that $f_{\mathbb{J}}$ is integrable with respect to M, and specifically

$$\int_{\mathbb{D}^{\vee}} f_{\mathbb{J}}\, \mathrm{d}M = \#\mathbb{J} \int_{\mathbb{W}^{\vee}} (u_L - u_L(\Omega))\, \mathrm{d}p_M^{\mathbb{J}}.$$

We claim that f coincides with $f_{\mathbb{J}}$ on a set of M-measure 1, and is therefore M-integrable with the same integral.

Since u_L is $p_M^{\mathbb{J}}$-integrable, there is a measurable function $\bar{u}_L$ on $\mathbb{W}^{\vee}$ that equals u_L on a set $A \subset \mathbb{W}^{\vee}$ of measure 1 with respect to $p_M^{\mathbb{J}}$ (see note 48). By the definition of $p_M^{\mathbb{J}}$, this A must have measure 1 with respect to each $\mathcal{P}_i^{\vee}(M)$, $i \in \mathbb{J}$. In particular, A has measure 1 with respect to 1_{Ω}, so we have $\Omega \in A$.

Define $\bar{f}, \bar{f}_{\mathbb{K}}$ by the same formulas as $f, f_{\mathbb{K}}$, but using $\bar{u}_L$ instead of u_L. Then $\bar{f}_{\mathbb{K}}$, the sum of measurable functions, is itself measurable. Note that $\bar{f}|_{\mathbb{D}_{\mathbb{K}}^{\vee}} = \bar{f}_{\mathbb{K}}|_{\mathbb{D}_{\mathbb{K}}^{\vee}}$. Therefore, for any measurable $C \subset \mathbb{R}$, $(\bar{f})^{-1}(C) \cap \mathbb{D}_{\mathbb{K}}^{\vee} = (\bar{f}_{\mathbb{K}})^{-1}(C) \cap \mathbb{D}_{\mathbb{K}}^{\vee}$, showing that $(\bar{f})^{-1}(C) \cap \mathbb{D}_{\mathbb{K}}^{\vee}$ is measurable in $\mathbb{D}_{\mathbb{K}}^{\vee}$. Since this works for every $\mathbb{K}$, we conclude from coherence that $(\bar{f})^{-1}(C)$ is measurable. Therefore $\bar{f}$ is a measurable function. Since they are both measurable functions, the set $B_1 \subset \mathbb{D}^{\vee}$ on which $\bar{f}$ and $\bar{f}_{\mathbb{J}}$ coincide is measurable. B_1 clearly includes $\mathbb{D}_{\mathbb{J}}^{\vee}$, on which M is supported, so B_1 has M-measure 1.

Now consider the set $B_2 = \bigcap_{i \in \mathbb{I}^{\infty}} (\mathcal{W}_i^{\vee})^{-1}(A)$. Using that fact that $\Omega \in A$, we see that, for each $\mathbb{K} \subset \mathbb{I}^{\infty}$, we have $B_2 \cap \mathbb{D}_{\mathbb{K}}^{\vee} = \bigcap_{i \in \mathbb{K}} (\mathcal{W}_i^{\vee})^{-1}(A) \cap \mathbb{D}_{\mathbb{K}}^{\vee}$. This is the intersection of $\mathbb{D}_{\mathbb{K}}^{\vee}$ with a measurable set. Therefore $B_2 \cap \mathbb{D}_{\mathbb{K}}^{\vee}$ is measurable in $\mathbb{D}_{\mathbb{K}}^{\vee}$, for any $\mathbb{K}$. So, by coherence, B_2 is measurable. Since M is supported on $\mathbb{D}_{\mathbb{J}}^{\vee}$, Lemma A.3 also gives us $M(B_2) = M(\bigcap_{i \in \mathbb{J}} (\mathcal{W}_i^{\vee})^{-1}(A))$. Moreover, $M((\mathcal{W}_i^{\vee})^{-1}(A)) = \mathcal{P}_i^{\vee}(M)(A) = 1$ for $i \in \mathbb{J}$; therefore $M(B_2) = 1$. Finally, since $\bar{u}_L|_A = u_L|_A$, we have $\bar{f}|_{B_2} = f|_{B_2}$ and $\bar{f}_{\mathbb{J}}|_{B_2} = f_{\mathbb{J}}|_{B_2}$. Combining these equalities with the fact that $\bar{f}|_{B_1} = \bar{f}_{\mathbb{J}}|_{B_1}$, we find that $f|_{B_1 \cap B_2} = \bar{f}|_{B_1 \cap B_2} = \bar{f}_{\mathbb{J}}|_{B_1 \cap B_2} = f_{\mathbb{J}}|_{B_1 \cap B_2}$. That is, as claimed, f coincides with $f_{\mathbb{J}}$ on $B_1 \cap B_2$, a measurable set of M-measure 1. In summary, f is $\mathbb{L}^{\vee}$-integrable, with

$$\int_{\mathbb{D}^{\vee}} f\, \mathrm{d}M = \#\mathbb{J} \int_{\mathbb{W}^{\vee}} (u_L - u_L(\Omega))\, \mathrm{d}p_M^{\mathbb{J}}. \tag{24}$$

Now, given arbitrary $M \in \mathbb{L}^{\vee}$, we can again choose $\mathbb{J} \supset \mathbb{I}$ with $L, M \in \mathbb{L}_{\mathbb{J}}^{\vee}$, and $U^{\vee}$ is integrally Gâteaux differentiable at $p_L^{\mathbb{J}}$. Note that $p_L^{\mathbb{J}}$ is a mixture of $p_L^{\mathbb{I}}$ and 1_{Ω}. We now apply Lemma A.7, with $P := p_L^{\mathbb{I}}$, $P_{\alpha} := p_L^{\mathbb{J}}$, any $Q \in \mathbb{P}^{\vee}$, $\beta := 1$, and $u^{\vee} := u_L$. The hypotheses (19) hold because u_L was chosen from $\nabla U^{\vee}_{p_L^{\mathbb{I}}}$, and the conclusion is that this same u_L is also in $\nabla U^{\vee}_{p_L^{\mathbb{J}}}$.

Since $\nabla U^{\vee}_{p_L^{\mathbb{J}}}$ is closed under the addition of constant functions, $u_L - u_L(\Omega) \in \nabla U^{\vee}_{p_L^{\mathbb{J}}}$. Combining this fact with Eqs. (23) and (24), we find

$$\begin{aligned}(V^{\vee})'_L(M - L) &= \#\mathbb{J}(U^{\vee})'_{p_L^{\mathbb{J}}}(p_M^{\mathbb{J}} - p_L^{\mathbb{J}}) \\ &= \#\mathbb{J} \int_{\mathbb{W}^{\vee}} (u_L - u_L(\Omega))\, \mathrm{d}(p_M^{\mathbb{J}} - p_L^{\mathbb{J}}) = \int_{\mathbb{D}^{\vee}} f\, \mathrm{d}(M - L).\end{aligned}$$

This shows $f \in \nabla V_L^{\vee}$, establishing part (ii) of the theorem, and the left to right direction of part (i).

For part (iii) of the theorem, suppose that u_L is a local utility function for $U^{\vee}$ at $p_L^{\mathbb{I}}$, and that $U^{\vee}(1_{\Omega}) = 0$. We first verify that $u_L(\Omega) = 0$. By definition of integral Gâteaux differentiability,

$$\begin{aligned}\lim_{t \to 0^+} \frac{U^{\vee}(p_L^{\mathbb{I}} + t(1_{\Omega} - p_L^{\mathbb{I}})) - U^{\vee}(p_L^{\mathbb{I}})}{t} &= (U^{\vee})'_{p_L^{\mathbb{I}}}(1_{\Omega} - p_L^{\mathbb{I}}) \\ &= \int_{\mathbb{W}^{\vee}} u_L \mathrm{d}(1_{\Omega} - p_L^{\mathbb{I}}).\end{aligned}$$

Since $U^{\vee}$ is Omega-linear, the left-hand side simplifies to $-U^{\vee}(p_L^{\mathbb{I}})$, whereas, using Lemma 5.4, the right-hand side simplifies to $u_L(\Omega) - U^{\vee}(p_L^{\mathbb{I}})$. Hence $u_L(\Omega) = 0$.

Taking this into account, the final claim of the theorem is that our f is a local utility function for $V^{\vee}$ at L. Since we have already shown $f \in \nabla V_L^{\vee}$, it is enough by Lemma 5.4 to prove that $\int_{\mathbb{D}^{\vee}} f\, dL = V^{\vee}(L)$. Moreover, since by hypothesis $U^{\vee}(1_{\Omega}) = 0$, the definition of $V^{\vee}$ reduces to $V^{\vee}(L) = \#\mathbb{I}U^{\vee}(p_L^{\mathbb{I}})$.

By Eq. (24), putting $M := L$, we have $\int_{\mathbb{D}^{\vee}} f\, \mathrm{d}L = \#\mathbb{J} \int_{\mathbb{W}^{\vee}} u_L\, \mathrm{d}p_L^{\mathbb{J}}$. Here we cannot simply replace $\mathbb{J}$ by $\mathbb{I}$, because the derivation of (24) was premised on a large enough choice of $\mathbb{J}$. However, $p_L^{\mathbb{J}} = \frac{\#\mathbb{I}}{\#\mathbb{J}} p_L^{\mathbb{I}} + \frac{\#\mathbb{J} - \#\mathbb{I}}{\#\mathbb{J}} 1_{\Omega}$, so we find $\int_{\mathbb{D}^{\vee}} f\, \mathrm{d}L = \#\mathbb{I} \int_{\mathbb{W}^{\vee}} u_L\, \mathrm{d}p_L^{\mathbb{I}} + (\#\mathbb{J} - \#\mathbb{I}) \int_{\mathbb{W}^{\vee}} u_L\, \mathrm{d}1_{\Omega}$. Since $u_L(\Omega) = 0$, the last term vanishes, whereas Lemma 5.4 shows that $\int_{\mathbb{W}^{\vee}} u_L\, \mathrm{d}p_L^{\mathbb{I}} = U^{\vee}(p_L^{\mathbb{I}})$. Therefore $\int_{\mathbb{D}^{\vee}} f\, \mathrm{d}L = \#\mathbb{I}U^{\vee}(p_L^{\mathbb{I}})$ as desired. □

Proof of Proposition 6.1. The proof of (i) is an easy version of the proof of (ii), so we present only the latter.

Suppose that $\succsim_0^{\vee}$ is consistent with quasi utilitarianism, and specifically corresponds to an individual preorder $\succsim_{\mathbb{P}^{\vee}}$. For any finite, non-empty $\mathbb{I} \subset \mathbb{I}^{\infty}$ and $d \in \mathbb{D}_{\mathbb{I}}^{\vee}$, define $p_d^{\mathbb{I}} := \frac{1}{\#\mathbb{I}} \sum_{i \in \mathbb{I}} 1_{\mathcal{W}_i^{\vee}(d)}$. Thus for $d, d' \in \mathbb{D}_{\mathbb{I}}^{\vee}$, we have $d \succsim_0^{\vee} d'$ iff $p_d^{\mathbb{I}} \succsim_{\mathbb{P}^{\vee}} p_{d'}^{\mathbb{I}}$. Suppose that $c \in \mathbb{D}^{\vee}$ is an m-scaling of $d \in \mathbb{D}_{\mathbb{I}}^{\vee}$, and that s is a corresponding m-to-1 map. Then it is easy to see that $p_d^{\mathbb{I}} = p_c^{s^{-1}(\mathbb{I})}$. Now, given $d, d' \in \mathbb{D}_{\mathbb{I}}^{\vee}$, their m-scalings c, c', and corresponding m-to-1 maps s, s', we can, by applying a permutation to c, ensure that $s^{-1}(\mathbb{I}) = (s')^{-1}(\mathbb{I}) =: \mathbb{J}$. Since then c and c' are in $\mathbb{D}_{\mathbb{J}}^{\vee}$, we have

$$c \succsim_0^{\vee} c' \iff p_c^{\mathbb{J}} \succsim_{\mathbb{P}^{\vee}} p_{c'}^{\mathbb{J}} \iff p_d^{\mathbb{I}} \succsim_{\mathbb{P}^{\vee}} p_{d'}^{\mathbb{I}} \iff d \succsim_0^{\vee} d'.$$

Therefore $\succsim_0^{\vee}$ satisfies Scale Invariance.

Conversely, suppose that $\succsim_0^{\vee}$ satisfies Scale Invariance; we need to define a corresponding individual preorder. We first show that $\mathbb{P}^{\vee}$ contains the set $\mathbb{P}_0^{\vee}$ of convex combinations of delta-measures on $\mathbb{W}^{\vee}$ with rational coefficients. For any $w \in \mathbb{W}^{\vee}$ and finite, nonempty $\mathbb{I} \subset \mathbb{I}^{\infty}$, $\mathbb{L}^{\vee}$ contains $1_{\mathcal{D}_{\mathbb{I}}^{\vee}(w)}$: for by variable population domain condition **(B$^{\vee}$)** we have $\mathcal{D}_{\mathbb{I}}^{\vee}(w) \in \mathbb{D}^{\vee}$, and by hypothesis in Section 6.1, $\mathbb{L}^{\vee}$ contains 1_d for every $d \in \mathbb{D}^{\vee}$. So by the domain condition **(A$^{\vee}$)**, $\mathbb{P}^{\vee}$ contains 1_w; since $\mathbb{P}^{\vee}$ is convex, it contains $\mathbb{P}_0^{\vee}$.

For any w, $w' \in \mathbb{W}^{\vee}$, the sigma algebra on $\mathbb{D}^{\vee}$ separates $\mathcal{D}_{\mathbb{I}}^{\vee}(w)$ and $\mathcal{D}_{\mathbb{I}}^{\vee}(w')$ by assumption, and since $\mathcal{D}_{\mathbb{I}}^{\vee}$ is measurable, the sigma algebra on $\mathbb{W}^{\vee}$ separates w and w'. By Lemma A.4, the representation of members of $\mathbb{P}_0^{\vee}$ by convex combinations of delta-measures is essentially unique: any $p \in \mathbb{P}_0^{\vee}$ is the sum of a unique finite set of delta-measures with non-zero coefficients, and these (rational) coefficients are uniquely determined. We will use this to first define a preorder on $\mathbb{P}_0^{\vee}$ and then extend it to a preorder on $\mathbb{P}^{\vee}$.

Choose a sequence of populations $\mathbb{I}_1 \subset \mathbb{I}_2 \subset \ldots$ such that $\#\mathbb{I}_n = n$. For any $p \in \mathbb{P}_0^{\vee}$, there is some $n > 0$ and $d \in \mathbb{D}_{\mathbb{I}_n}^{\vee}$ such that $p = p_d^{\mathbb{I}_n}$. In this case say that d is a realization of p at n. More specifically, for any $p \in \mathbb{P}_0^{\vee}$, let $N(p)$ be the least common denominator of the rational coefficients appearing in p. Then p has a realization at n if and only if n is a multiple of $N(p)$. Moreover, any realization of p at, say, $mN(p)$ is an m-scaling of any realization of p at $N(p)$.

For any pair $p, p' \in \mathbb{P}_0^{\vee}$, let $N(p, p')$ be the least common multiple of $N(p)$ and $N(p')$: p and p' both have realizations at n if and only if n is a multiple of $N(p, p')$. Let $I(p, p')$ be the set of all such multiples. The scale-invariance of $\succsim_0^{\vee}$ yields the following observation. If d, d' are realizations of p, p' at $m \in I(p, p')$, and c, c' are realizations of p, p' at $n \in I(p, p')$, then $d \succsim_0^{\vee} d'$ if and only if $c \succsim_0^{\vee} c'$.

This allows us to define $\succsim_{\mathbb{P}_0^{\vee}}$ on $\mathbb{P}_0^{\vee}$ as follows:

$p \succsim_{\mathbb{P}_0^{\vee}} p' \iff$ for some (therefore any) $n \in I(p, p')$, there are realizations d, d' of p, p' at n with $d \succsim_0^{\vee} d'$.

This is a preorder. In particular it is transitive, since, given $p, p', p'' \in \mathbb{P}^{\mathrm{v}}_0$, we can consider realizations d, d', d'' of p, p', p'' at some common n. If $p \succsim_{\mathbb{P}^{\mathrm{v}}_0} p' \succsim_{\mathbb{P}^{\mathrm{v}}_0} p''$ then we must have $d \succsim^{\mathrm{v}}_0 d' \succsim^{\mathrm{v}}_0 d''$. Since $\succsim_0$ is transitive, $d \succsim_0 d''$, and therefore $p \succsim_{\mathbb{P}^{\mathrm{v}}_0} p''$.

Let us also check that $\succsim_{\mathbb{P}^{\mathrm{v}}_0}$ satisfies Omega Independence. Suppose given $p, p' \in \mathbb{P}^{\mathrm{v}}_0$, and $m/n =: \alpha \in (0, 1) \cap \mathbb{Q}$. Then realizations of p, p' at $N(p, p')m$ are elements of $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}_{N(p,p')m}}$; considered as elements of the larger set $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}_{N(p,p')n}}$, they are also realizations of $\alpha p + (1-\alpha)1_\Omega$ and $\alpha p' + (1-\alpha)1_\Omega$ at $N(p, p')n$. It follows that $p \succsim_{\mathbb{P}^{\mathrm{v}}_0} p'$ if and only if $\alpha p + (1-\alpha)1_\Omega \succsim_{\mathbb{P}^{\mathrm{v}}_0} \alpha p' + (1-\alpha)1_\Omega$, as desired.

We now extend $\succsim_{\mathbb{P}^{\mathrm{v}}_0}$ to a preorder $\succsim_{\mathbb{P}^{\mathrm{v}}}$ on $\mathbb{P}^{\mathrm{v}}$. Here is a construction that works in general (of course, in any given case there may be more natural ways to do it).

$$p \succsim_{\mathbb{P}^{\mathrm{v}}} p' \iff \begin{cases} p, p' \in \mathbb{P}^{\mathrm{v}}_0 \text{ and } p \succsim_{\mathbb{P}^{\mathrm{v}}_0} p', \text{ or} \\ p = p'. \end{cases}$$

Then $\succsim_{\mathbb{P}^{\mathrm{v}}}$ is a preorder on $\mathbb{P}^{\mathrm{v}}$ which satisfies Omega Independence. (Here we rely on the fact that Omega Independence only quantifies over rational values of α.) Let $\succsim^{\mathrm{v}}$ be the social preorder on $\mathbb{L}^{\mathrm{v}}$ it generates. Then, for any finite non-empty set $\mathbb{I} \subset \mathbb{I}^\infty$ such that d and d' are in $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$, $d \succsim^{\mathrm{v}}_0 d' \iff p^{\mathbb{I}}_d \succsim_{\mathbb{P}^{\mathrm{v}}} p^{\mathbb{I}}_{d'} \iff p^{\mathbb{I}}_{1_d} \succsim_{\mathbb{P}^{\mathrm{v}}} p^{\mathbb{I}}_{1_{d'}} \iff 1_d \succsim^{\mathrm{v}} 1_{d'}$. This shows that $\succsim^{\mathrm{v}}_0$ is consistent with the quasi utilitarian preorder $\succsim^{\mathrm{v}}$. □

Appendix B. Index of global notation

We recall the convention (see Section 3.1) that a superscript 'v' is used to distinguish variable population objects, while a superscript '∗' is used in Section 4.3 to indicate neutrality between a variable population model and a family of constant population models (see Section 4.3 for discussion).

Sets

- $\mathbb{D}, \mathbb{D}_{\mathbb{I}}$ The set of welfare distributions in a constant population model with population $\mathbb{I}$, Sections 2.2 and 4.3.
- $\mathbb{D}^{\mathrm{v}}$ The set of welfare distributions in a variable population model, Section 3.1.
- $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$ The subset of $\mathbb{D}^{\mathrm{v}}$ consisting of distributions in which only individuals in $\mathbb{I}$ exist, Section 3.1.
- $\mathbb{I}$ A finite set of individuals; the population in the basic constant population model, Section 2.2.
- $\mathbb{I}^\infty$ The infinite set of possible individuals in the basic variable population model (see Section 3.1), or in a family of constant population models, Section 4.3.
- $\mathbb{L}, \mathbb{L}_{\mathbb{I}}$ The set of lotteries in a constant population model with population $\mathbb{I}$, Sections 2.2 and 4.3.
- $\mathbb{L}^{\mathrm{v}}$ The set of lotteries in a variable population model, Section 3.1.
- $\mathbb{L}^{\mathrm{v}}_{\mathbb{I}}$ The subset of $\mathbb{L}^{\mathrm{v}}$ consisting of lotteries supported on $\mathbb{D}^{\mathrm{v}}_{\mathbb{I}}$, Section 3.1.
- $\mathcal{P}(Y)$ A generic convex set of probability measures on a measurable space Y, Section 4.2.
- $\mathbb{P}$ The set of prospects in a constant population model, Section 2.2.
- $\mathbb{P}_\Omega$ $\mathbb{P} \cup \{1_\Omega\}$, Section 3.4.
- $\mathbb{P}^{\mathrm{v}}$ The set of prospects in a variable population model, Section 3.1.
- $\mathbb{R}$ The set of real numbers.
- Σ The group of permutations of $\mathbb{I}$; it acts on $\mathbb{D}$ and $\mathbb{L}$, Section 2.2.
- Σ^∞ The group of permutations of $\mathbb{I}^\infty$; it acts on $\mathbb{D}^{\mathrm{v}}$ and $\mathbb{L}^{\mathrm{v}}$, Section 3.1.
- $\mathbb{W}$ The set of welfare states, excluding non-existence, Section 2.2.
- $\mathbb{W}^{\mathrm{v}}$ The set of welfare states, including Ω, Section 3.1.

Functions

- $\mathcal{D}$ $\mathcal{D}(w)$ is the constant population distribution in which all individuals in $\mathbb{I}$ have welfare state w, Section 2.2.
- $\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}$ $\mathcal{D}^{\mathrm{v}}_{\mathbb{I}}(w)$ is the variable population distribution in which all individuals in $\mathbb{I}$ have welfare state w, and no one else exists, Section 3.1.
- $\mathcal{L}$ $\mathcal{L}(P)$ is the constant population lottery in which all individuals in $\mathbb{I}$ face prospect P, with perfect correlation, Section 2.2.
- $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}$ $\mathcal{L}^{\mathrm{v}}_{\mathbb{I}}(P)$ is the variable population lottery in which all individuals in $\mathbb{I}$ face prospect P, with perfect correlation, and no one else exists, Section 3.1.
- $\mathcal{P}_i$ $\mathcal{P}_i(L)$ is the prospect faced by individual i in constant population lottery L, Section 2.2.
- $\mathcal{P}^{\mathrm{v}}_i$ $\mathcal{P}^{\mathrm{v}}_i(L)$ is the prospect faced by individual i in variable population lottery L, Section 3.1.
- $\mathcal{W}_i$ $\mathcal{W}_i(d)$ is the welfare state of individual i in constant population distribution d, Section 2.2.
- $\mathcal{W}^{\mathrm{v}}_i$ $\mathcal{W}^{\mathrm{v}}_i(d)$ is the welfare state of individual i in variable population distribution d, Section 3.1.

Preorders

- $\succsim_X$ A generic preorder on a (typically convex) set X, Section 1.
- $\sim_X$ For any preorder $\succsim_X$, $x \sim_X y$ if and only if both $x \succsim_X y$ and $y \succsim_X x$, Section 2.3.
- $\curlywedge_X$ For any preorder $\succsim_X$, $x \curlywedge_X y$ if and only if neither $x \succsim_X y$ nor $y \succsim_X x$, Section 2.3.
- $\succ_X$ For any preorder $\succsim_X$, $x \succ_X y$ if and only if $x \succsim_X y$ but not $y \succsim_X x$, Section 2.3.
- $\succsim_{\mathbb{P}}$ The constant population individual preorder on $\mathbb{P}$, Section 2.3.
- $\succsim, \succsim_{\mathbb{I}}$ The social preorder on $\mathbb{L}$ in a constant population model with population $\mathbb{I}$, Sections 2.3 and 4.3.
- $\succsim_0$ A social preorder on $\mathbb{D}$, Section 6.1.
- $\succsim_{\mathbb{P}^{\mathrm{v}}}$ The individual preorder on $\mathbb{P}^{\mathrm{v}}$ in a variable population model, Section 3.2.
- $\succsim^{\mathrm{v}}$ The social preorder on $\mathbb{L}^{\mathrm{v}}$ in a variable population model, Section 3.2.
- $\succsim^{\mathrm{v}}_0$ A social preorder on $\mathbb{D}^{\mathrm{v}}$, Section 6.1.
- $\approx^{\mathbb{J}}_{\mathbb{P}*}, \rhd^{\mathbb{J}}_{\mathbb{P}*}, \bowtie^{\mathbb{J}}_{\mathbb{P}*}$ Relations used to define Pareto axioms, Section 4.3.
- $\succsim^{SD}_X$ The stochastic dominance preordering on a set X of probability measures, Section 5.1.

Miscellaneous

- 1_y The delta measure supported at a point y, Section 2.3.
- $\int$ The weak integral, Section 4.2.
- Ω Non-existence, considered formally as a welfare state, Section 3.1.
- d_Ω The empty distribution, in which no individuals exist, Section 3.1.
- $\mathbb{F}$ A family of constant population models, Section 4.3.
- $L|_{\mathbb{J}}$ The restriction of a lottery L to a sub-population $\mathbb{J}$, Section 4.3.

$\mathbb{M}$ A constant population model, Section 4.3.

$\mathbb{M}^{\mathrm{v}}$ A variable population model, Section 4.3.

p_L The prospect derived from the constant population lottery L by averaging over individuals: $p_L = \frac{1}{\#\mathbb{I}} \sum_{i\in\mathbb{I}} \mathcal{P}_i(L)$, Section 2.4.

$p_L^{\mathbb{I}}$ The prospect derived from the variable population lottery L by averaging over individuals from $\mathbb{I}$: $p_L^{\mathbb{I}} = \frac{1}{\#\mathbb{I}} \sum_{i\in\mathbb{I}} \mathcal{P}_i^{\mathrm{v}}(L)$, Section 3.3.

U_p' The Gâteaux derivative of U at p, Section 5.2.

∇U_p The set of integral kernels for the Gâteaux derivative of U at p, Section 5.2.

References

Adler, M., Fleurbaey, M. (Eds.), 2016. The Oxford Handbook of Well-Being and Public Policy. Oxford University Press.
Allais, M., 1979. The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American school. In: Allais, M., Hagen, O. (Eds.), Expected Utility Hypothesis and the Allais Paradox. Reidel, Dordrecht, pp. 27–145.
Alon, S., Gayer, G., 2016. Utilitarian preferences with multiple priors. Econometrica 84, 1181–1201.
Anscombe, F., Aumann, R., 1963. A definition of subjective probability. Ann. Math. Stat. 34, 199–205.
Arrhenius, G., Rabinowicz, W., 2015. Value superiority. In: Hirose, I., Olson, J. (Eds.), The Oxford Handbook of Value Theory. Oxford University Press, pp. 225–248.
Barry, B., 1977. Rawls on average and total utility: a comment. Philos. Stud. 31, 317–325.
Ben-Porath, E., Gilboa, I., Schmeidler, D., 1977. On the measurement of inequality under uncertainty. J. Econom. Theory 75, 194–204.
Bewley, T., 2002. Knightian decision theory: part I. Decis. Econ. Finance 25, 79–110.
Billot, A., Vergopoulos, V., 2016. Aggregation of paretian preferences for independent individual uncertainties. Soc. Choice Welf. 47, 973–984.
Blackorby, C., Bossert, W., Donaldson, D., 1995. Intertemporal population ethics: critical-level utilitarian principles. Econometrica 65, 1303–1320.
Blackorby, C., Bossert, W., Donaldson, D., 1998. Uncertainty and critical-level population principles. J. Popul. Econ. 11, 1–20.
Blackorby, C., Bossert, W., Donaldson, D., 2005. Population Issues in Social Choice Theory, Welfare Economics, and Ethics. Cambridge University Press, New York.
Blackorby, C., Bossert, W., Donaldson, D., 2007. Variable-population extensions of social aggregation theorems. Soc. Choice Welf. 28, 567–589.
Blackorby, C., Donaldson, D., 1980. A theoretical treatment of indices of absolute equality. Internat. Econom. Rev. 21, 107–136.
Bogachev, V., 2007. Measure Theory. Springer.
Border, K., 1981. Notes on Von Neumann–Morgenstern Social Welfare Functions. California Institute of Technology, Department of Economics, Unpublished manuscript.
Borie, D., 2016. Lexicographic expected utility without completeness. Theory and Decision 81, 167–176.
Bostrom, N., 2011. Infinite ethics. Anal. Metaphys. 10, 9–59.
Broome, J., 1989. What's the good of equality? In: Hey, J.D. (Ed.), Current Issues in Microeconomics. Macmillan, pp. 236–262.
Broome, J., 1990. Bolker–Jeffrey expected utility theory and axiomatic utilitarianism. Rev. Econom. Stud. 57, 477–502.
Broome, J., 1991. Weighing Goods. Blackwell, Oxford.
Broome, J., 1993. A cause of preference is not an object of preference. Soc. Choice Welf. 10, 57–68.
Broome, J., 1999. Ethics Out of Economics. Cambridge University Press.
Broome, J., 2004. Weighing Lives. Oxford University Press, New York.
Brunnermeier, M., Simsek, A., Xiong, W., 2014. A welfare criterion for models with distorted beliefs. Q. J. Econ. 129 (4), 1753–1797.
Buchak, L., 2013. Risk and Rationality. Oxford University Press.
Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., 2017. Stochastic dominance analysis without the independence axiom. Manage. Sci. 63, 1097–1109.
Chambers, C., Hayashi, T., 2006. Preference aggregation under uncertainty: Savage vs. Pareto. Games Econom. Behav. 54, 430–440.
Chambers, C., Hayashi, T., 2014. Preference aggregation with incomplete information. Econometrica 82, 589–599.
Chateauneuf, A., Cohen, M., 1994. Risk-seeking with diminishing marginal utility in a non-expected utility model. J. Risk Uncertain. 9, 77–91.
Chew, S., Karni, E., Safra, Z., 1987. Risk aversion in the theory of expected utility with rank dependent probabilities. J. Econom. Theory 42, 370–381.
Chew, S., Mao, M., 1995. A Schur concave characterization of risk-aversion for non-expected utility preferences. J. Econom. Theory 67, 402–435.
Conrad, P., 1953. Embedding theorems for abelian groups with valuations. Amer. J. Math. 75, 1–29.
Coulhon, T., Mongin, P., 1989. Social choice theory in the case of von Neumann Morgenstern utilities. Soc. Choice Welf. 6, 175–187.
Crès, A., Gilboa, I., Vieille, N., 2011. Aggregation of multiple prior opinions. J. Econom. Theory 146, 2563–2582.
Danan, E., Gajdos, T., Hill, B., Tallon, J.-M., 2016. Robust social decisions. Amer. Econ. Rev. 106, 2407–2425.
Danan, E., Gajdos, T., Tallon, J.-M., 2015. Harsanyi's aggregation theorem with incomplete preferences. Am. Econ. J. Microecon. 7, 61–69.
d'Aspremont, C., Gevers, L., 1977. Equity and the informational basis of collective choice. Rev. Econom. Stud. 44, 199–209.
d'Aspremont, C., Gevers, L., 2002. Social welfare functionals and interpersonal comparability. In: Arrow, K., Sen, A., Suzumara, K. (Eds.), Handbook of Social Choice and Welfare, Vol. 1. pp. 459–541.
De Meyer, B., Mongin, P., 1995. A note on affine aggregation. Econom. Lett. 47, 177–183.
Diamond, P., 1967. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility: comment. J. Polit. Econ. 75, 765–766.
Dietrich, F., List, C., 2016. Probabilistic opinion pooling. In: Hàjek, A., Hitchcock, C. (Eds.), The Oxford Handbook of Philosophy and Probability. Oxford University Press, pp. 519–541.
Domotor, Z., 1979. Ordered sum and tensor product of linear utility structures. Theory and Decision 11, 375–399.
Dravecký, J., 1975. Spaces with measurable diagonal. Mat. Cas. 25, 3–9.
Dubra, J., 2011. Continuity and completeness under risk. Math. Social Sci. 61, 80–81.
Dubra, J., Maccheroni, F., Ok, E., 2004. Expected utility theory without the completeness axiom. J. Econom. Theory 115, 118–133.
Dudley, R., 2002. Real Analysis and Probability, second ed. Cambridge University Press, New York.
Evren, Ö, 2008. On the existence of expected multi-utility representations. Econom. Theory 35, 575–592.
Evren, Ö, 2014. Scalarization methods and expected multi-utility representations. J. Econom. Theory 151, 30–63.
Evren, Ö, Ok, E., 2011. On the multi-utility representation of preference relations. J. Math. Econom. 47, 554–563.
Fishburn, P., 1970. Utility Theory for Decision Making. Wiley, New York.
Fishburn, P., 1988. Nonlinear Preference and Utility Theory. The John Hopkins University Press, Baltimore, London.
Fleurbaey, M., 2003. On the informational basis of social choice. Soc. Choice Welf. 21, 347–384.
Fleurbaey, M., 2009. Two variants of Harsanyi's aggregation theorem. Econom. Lett. 105, 300–302.
Fleurbaey, M., 2010. Assessing risky social situations. J. Polit. Econ. 118, 649–680.
Fleurbaey, M., 2018. Welfare economics, risk and uncertainty. Can. J. Econ. 51.
Fleurbaey, M., Voorhoeve, A., 2016. Priority or equality for possible people? Ethics 126, 929–954.
Fleurbaey, M., Zuber, S., 2017. Fair management of social risk. J. Econom. Theory 169, 666–706.
Gajdos, T., Tallon, J.-M., Vergnaud, J.-C., 2008. Representation and aggregation of preferences under uncertainty. J. Econom. Theory 141, 68–99.
Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2003. A subjective spin on roulette wheels. Econometrica 71, 1897–1908.
Gilboa, I., Samet, D., Schmeidler, D., 2004. Utilitarian aggregation of beliefs and tastes. J. Polit. Econ. 112, 932–938.
Gilboa, I., Samuelson, L., Schmeidler, D., 2014. No-betting Pareto dominance. Econometrica 82, 1405–1442.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. J. Math. Econom. 18, 141–153.
Gorno, L., 2017. A strict expected multi-utility theorem. J. Math. Econom. 71, 92–95.
Grant, S., Kajii, A., Polak, B., Safra, Z., 2010. Generalized utilitarianism and Harsanyi's impartial observer theorem. Econometrica 78, 1939–1971.
Grant, S., Kajii, A., Polak, B., Safra, Z., 2012a. A generalized representation theorem for Harsanyi's ('impartial') observer. Soc. Choice Welf. 39, 833–846.
Grant, S., Kajii, A., Polak, B., Safra, Z., 2012b. Equally-distributed equivalent utility, ex post egalitarianism and utilitarianism. J. Econom. Theory 147, 1545–1571.
Greaves, H., 2017. A reconsideration of the Harsanyi–Sen–Weymark debate on utilitarianism. Utilitas 29, 175–213.
Gustafsson, J., 2019. Population axiology and the possibility of a fourth category of absolute value. Econ. Philos. http://dx.doi.org/10.1017/S0266267119000087 published online: 14 May, 2019.
Halpern, J., 2003. Reasoning About Uncertainty. M.I.T. Press.
Hammond, P., 1988. Consequentialist demographic norms and parenting rights. Soc. Choice Welf. 25, 127–145.
Hammond, P., 1991. Interpersonal comparisons of utility: why and how they are and should be made. In: Elster, J., Roemer, J. (Eds.), Interpersonal Comparisons of Well-Being. Cambridge University Press, pp. 200–254.

Hammond, P., 1998. Objective expected utility: a consequentialist perspective. In: Barberà, S., Seidel, C., Hammond, P. (Eds.), Handbook of Utility Theory, Vol. 1. Kluwer, Dordrecht, pp. 145–211.
Hara, K., Ok, E., Riella, G., 2019. Coalitional expected multi-utility theory. Econometrica 87 (3), 933–980.
Harsanyi, J., 1953. Cardinal utility in welfare economics and in the theory of risk-taking. J. Polit. Econ. 61, 434–435.
Harsanyi, J., 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. J. Polit. Econ. 63, 309–321.
Harsanyi, J., 1977a. Morality and the theory of rational behavior. Soc. Res. 44, 623–656.
Harsanyi, J., 1977b. Rational Behaviour and Bargaining Equilibrium in Games and Social Situations. Cambridge University Press, Cambridge.
Harsanyi, J., 1978. Bayesian decision theory and utilitarian ethics. Amer. Econ. Rev. 68, 223–228.
Harsanyi, J., 1979. Bayesian decision theory, rule utilitarianism, and Arrow's impossibility result. Theory and Decision 11, 289–317.
Hausner, M., Wendel, J., 1952. Ordered vector spaces. Proc. Amer. Math. Soc. 3, 977–982.
Herstein, I., Milnor, J., 1953. An axiomatic approach to measurable utility. Econometrica 21, 291–297.
Karni, E., Zhou, N., 2019. Weighted utility theory with incomplete preferences. https://www.econ2.jhu.edu/people/karni/WeightedUtilitywithIncompletePreferences.pdf.
Kavka, G., 1975. Rawls on average and total utility. Philos. Stud. 27, 237–253.
Kelley, J., Namioka, I., 1963. Linear Topological Spaces. Springer.
Kreps, D., 1988. Notes on the Theory of Choice. Westview Press, Boulder, London.
Luce, R., Raiffa, H., 1957. Games and Decisions. John Wiley and Sons, New York.
Machina, M., 1982. Expected utility' analysis without the independence axiom. Econometrica 50, 277–323.
McCarthy, D., 2015. Distributive equality. Mind 124, 1045–1109.
McCarthy, D., 2017. The priority view. Econ. Philos. 33, 215–257.
McCarthy, D., Mikkola, K., 2018. Continuity and completeness of strongly independent preorders. Math. Social Sci. 93, 141–145.
McCarthy, D., Mikkola, K., Thomas, T., 2017a. Representation of strongly independent preorders by sets of scalar-valued functions. MPRA Paper No. 79284. https://mpra.ub.uni-muenchen.de/79284/.
McCarthy, D., Mikkola, K., Thomas, T., 2017b. Representation of strongly independent preorders by vector-valued functions. MPRA Paper No. 80806. https://mpra.ub.uni-muenchen.de/80806/.
McCarthy, D., Mikkola, K., Thomas, T., 2019. Aggregation for potentially infinite populations without continuity or completeness. arXiv:1911.00872 [econ.TH].
Mongin, P., 1994. Harsanyi's aggregation theorem: multi-profile version and unsettled questions. Soc. Choice Welf. 11, 331–354.
Mongin, P., 1995. Consistent Bayesian aggregation. J. Econom. Theory 66, 313–351.
Mongin, P., 1998. The paradox of the Bayesian experts and state-dependent utility theory. J. Math. Econom. 29, 331–361.
Mongin, P., 2001a. The impartial observer theorem of social ethics. Econ. Philos. 17, 147–179.
Mongin, P., 2001b. A note on mixture sets in decision theory. Decis. Econ. Finance 24, 59–69.
Mongin, P., 2016. Spurious unanimity and the Pareto principle. Econ. Philos. 32, 511–532.
Mongin, P., d'Aspremont, C., 1998. Utility theory and ethics. In: Barberà, S., Hammond, P., Seidl, C. (Eds.), Handbook of Utility Theory, Vol. 1. Kluwer, Dordrecht, pp. 233–289.
Mongin, P., Pivato, M., 2015. Ranking multidimensional alternatives and uncertain prospects. J. Econom. Theory 157, 146–171.
Mongin, P., Pivato, M., 2016. Social choice under risk and uncertainty. In: Adler, M., Fleurbaey, M. (Eds.), In: The Oxford Handbook of Well-Being and Public Policy, Oxford University Press, pp. 711–745.
Mongin, P., Pivato, M., 2019. Social preference under twofold uncertainty. Econom. Theory http://dx.doi.org/10.1007/s00199-019-01237-0 published online: 15 November, 2019.
Myerson, R., 1981. Utilitarianism, egalitarianism, and the timing effect in social choice problems. Econometrica 49, 883–897.
Ng, Y., 1983. Some broader issues of social choice. In: Pattanaik, P., Salles, M. (Eds.), Social Choice and Welfare. North-Holland, Amsterdam, pp. 151–173.
Ok, E., 2002. Utility representation of an incomplete preference relation. J. Econom. Theory 104, 429–449.
Parfit, D., 1986. Reasons and Persons. Oxford University Press, New York.
Pivato, M., 2013. Risky social choice with incomplete or noisy interpersonal comparisons of well-being. Soc. Choice Welf. 40, 123–139.
Pivato, M., 2014. Additive representation of separable preferences over infinite products. Theory and Decision 77, 31–83.
Qu, X., 2017. Separate aggregation of beliefs and values under ambiguity. Econom. Theory 63, 1–17.
Quiggin, J., 1982. A theory of anticipated utility. J. Econ. Behav. Organ. 3, 323–343.
Quiggin, J., 1993. Generalized Expected Utility Theory: The Rank-Dependent Model. Kluwer, Dordrecht.
Rabinowicz, W., 2009. Broome and the intuition of neutrality. Philos. Issues 19, 389–411.
Rawls, J., 1971. A Theory of Justice. Belknap Press.
Richter, M., 1971. Rational choice. In: Chipman, J., Hurwicz, L., Richter, M., Sonnenschein, H. (Eds.), Preferences, Utility, and Demand: A Minnesota Symposium. Harcourt Brace Jovanovich, New York, pp. 29–58.
Rudin, W., 1991. Functional Analysis, second ed. McGraw-Hill.
Saito, K., 2013. Social preferences under risk: equality of opportunity versus equality of outcome. Amer. Econ. Rev. 103, 3084–3101.
Scanlon, T., 1982. Contractualism and utilitarianism. In: Sen, A., Williams, B. (Eds.), Utilitarianism and Beyond. Cambridge University Press, Cambridge, pp. 103–128.
Schmidt, U., 2004. Alternatives to expected utility: some formal theories. In: Barberà, S., Seidel, C., Hammond, P. (Eds.), Handbook of Utility Theory, Vol. 2. Kluwer, Dordrecht, pp. 757–838.
Sen, A., 1970. Collective Choice and Social Welfare. Holden Day, San Fransisco.
Sen, A., 1973. On Economic Inequality. Clarendon Press, Oxford.
Sen, A., 1976. Welfare inequalities and Rawlsian axiomatics. Theory and Decision 7, 243–262.
Sen, A., 1977. Non-linear social welfare functions: a reply to professor Harsanyi. In: Butts, R., Hintikka, J. (Eds.), Foundational Problems in the Special Sciences. Reidel, Dordrecht, pp. 297–302.
Sen, A., 1981. Plural utility. Proc. Aristot. Soc. 81, 193–216.
Sen, A., 1985. Well-being, agency and freedom. J. Philos. 82, 169–221.
Sen, A., 1986. Social choice theory. In: Arrow, K., Intriligator, M. (Eds.), Handbook of Mathematical Economics Amsterdam, Vol. 3. North Holland, pp. 1073–1181.
Starmer, C., 2000. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. J. Econ. Lit. 38, 332–382.
Thomas, T., 2018. Some possibilities in population axiology. Mind 127, 807–832.
Tungodden, B., 2003. The value of equality. Econ. Philos. 19, 1–44.
Vickrey, W., 1945. Measuring marginal utility by reactions to risk. Econometrica 13, 319–333.
Wakker, P., 1994. Separating marginal utility and probabilistic risk aversion. Theory and Decision 36, 1–44.
Wakker, P., 2010. Prospect Theory. Cambridge University Press, Cambridge.
Weymark, J., 1981. Generalized gini inequality indices. Math. Social Sci. 1, 409–430.
Weymark, J., 1991. A reconsideration of the Harsanyi–Sen debate on utilitarianism. In: Elster, J., Roemer, J. (Eds.), Interpersonal Comparisons of Well-Being. Cambridge University Press, pp. 255–320.
Weymark, J., 1993. Harsanyi's social aggregation theorem and the weak Pareto principle. Soc. Choice Welf. 10, 209–221.
Weymark, J., 1995. Further remarks on Harsanyi's social aggregation theorem and the weak Pareto principle. Soc. Choice Welf. 12, 87–92.
Yaari, M., 1987. The dual theory of choice under risk. Econometrica 55, 95–115.
Zhou, L., 1997. Harsanyi's utilitarianism theorems: general societies. J. Econom. Theory 72, 198–207.
Zuber, S., 2016. Harsanyi's theorem without the sure-thing principle: On the consistent aggregation of Monotonic Bernoullian and Archimedean preferences. J. Math. Econom. 63, 78–83.