Population Wellbeing Initiative at the UT-Austin PRC

Working Paper w2402

What Forms for Incomplete Population Ethics Does Social Expected Multi-Utility Permit?

Abstract: Famous problems in variable-population welfare economics have led some to suggest that social welfare comparisons over such populations may be incomplete. In the theory of rational choice with incomplete preferences, attention has recently centered on the Expected Multi-Utility framework, which permits incompleteness but preserves vNM independence and can be derived from weak, attractive axioms. Here, we apply this framework to variable-population welfare economics. We show that Expected Multi-Utility for social preferences, combined with a stochastic ex-ante-Pareto-type axiom, characterizes Expected Critical-Set Generalized Utilitarianism, in the presence of basic axioms. The further addition of Negative Dominance, an axiom recently introduced to the philosophy literature, yields a characterization of Expected Critical-Level Generalized Utilitarianism.

Authors

Harvey Lederman

Christian Tarsney

Dean Spears

The University of Texas at Austin
Population Wellbeing Initiative

What Forms for Incomplete Population Ethics Does Social Expected Multi-Utility Permit?

Harvey Lederman Christian Tarsney Dean Spears*

May 9, 2024

Abstract

Famous problems in variable-population welfare economics have led some to suggest that social welfare comparisons over such populations may be incomplete. In the theory of rational choice with incomplete preferences, attention has recently centered on the Expected Multi-Utility framework, which permits incompleteness but preserves vNM independence and can be derived from weak, attractive axioms. Here, we apply this framework to variable-population welfare economics. We show that Expected Multi-Utility for social preferences, combined with a stochastic ex-ante-Pareto-type axiom, characterizes Expected Critical-Set Generalized Utilitarianism, in the presence of basic axioms. The further addition of Negative Dominance, an axiom recently introduced to the philosophy literature, yields a characterization of Expected Critical-*Level* Generalized Utilitarianism.

Keywords: welfare economics, population economics, population ethics, incompleteness, expected multi-utility, utilitarianism, axiomatic characterization
JEL Classification numbers: D63, D81, I31, J10

*dspears@utexas.edu. Department of Economics and Population Research Center, University of Texas at Austin.

1 Introduction

The trajectory of the world population is changing, with important economic and policy implications. Birth rates are declining in richer countries (Kearney et al., 2022; Doepke et al., 2023) and in poorer countries (Jayachandran, 2017), so the size of the world population is projected to peak and then decline (Spears et al., 2024). The effects of on long-term economic growth (Jones, 2022) and on the climate (Scovronick et al., 2017) are two of many reasons why population change has emerged as a core policy question—highlighted, for example, with a chapter in the most recent Economic Report of the President (Council of Economic Advisors, 2024).

And yet, the assessment of social welfare when population size is endogenous remains a longstanding theoretical problem (Dasgupta, 1995; Golosov et al., 2007; Cowen, 2018). In the case of a fixed population, for instance, average welfare increases if and only if total welfare does. But if populations can vary in size, then these criteria no longer coincide. Tools such as social welfare functions and cost-benefit analysis are the engines of policy economics (Sunstein, 2024), but it is unclear how to extend them to questions of policy change.

Recently, a common approach has been to suggest that social preferences should be *incomplete* with respect to certain different-sized populations, not ranking them with respect to whether they are better, worse, or exactly as good. In this paper, we study such approaches, first, characterizing the form of the social welfare function it implies, and, second, providing a limitative result about the consequences of such incompleteness.

Our starting point is the paradigm for constraining the form of the social welfare function introduced by Harsanyi (1955). Harsanyi showed that, given standard axioms of rationality governing individual utility functions, similar axioms of rationality for the social planner, together with a mild form of "ex ante" Pareto axiom, tightly constrain the functional form of the social welfare function. In this paper, we develop an extension of the Harsanyi approach to variable-population comparisons, where comparisons across these different-sized populations may be incomplete.

In more detail: In his fixed population setting, Harsanyi assumed that the social planner has preferences satisfying the standard von Neumann-Morgenstern axioms: Completeness, Transitivity, Independence, and Continuity. Because our interest is in exploring possible incompleteness for variable populations, we here impose Completeness only over populations of the same size, dropping it between populations which vary in size. As is well-known, given Completeness, the standard Archimedean property of von Neumann-Morgenstern is derivable from Continuity. But without Completeness it is no longer derivable, so we impose it directly. A result of McCarthy et al. (2021) shows that if preferences satisfy these axioms (Transitivity, Independence, Continuity, and the Archimedean property) they are represented by an "expected multi-utility", that is:

- There is a set Φ of functions ϕ mapping outcomes to real numbers, each of which completely and transitively orders the final outcomes, and
- One lottery is at least as good as another if and only if its expectation of ϕ is at least as great for all $\phi \in \Phi$, and strictly better if and only if its expectation is at least as great for all ϕ and strictly greater for some ϕ.[1]

Our first main result shows that these constraints on social preferences, together with a Pareto-like axiom that social preferences respect first-order stochastic dominance for individuals in lotteries (which we call "Personal Good"), and minimal further background assumptions, allow us to characterize the form of the social-welfare function available for proponents of this kind of incompleteness. In particular, we show that, if social preferences satisfy Expected Multi-Utility and Personal Good, then social preferences take the form of "Expected Critical Band Generalized Utilitarianism." This already tightly constrains our approach to choices under uncertainty, and rules out several prominent treatments in the philosophy literature.[2]

Critical Band Generalized Utilitarianism allows for incompleteness in ranking populations of variable size. Our second result shows that a plausible dominance axiom, recently introduced into the philosophy literature, rules out even this incompleteness. Given this "Negative Dominance" axiom, which can be thought of as a further rationality constraint on social preferences in the spirit of Harsanyi, we can further move from Expected Critical Band Generalized Utilitarianism to a complete social welfare ordering, induced by Expected Critical Level Generalized Utilitarianism. In fact, Negative Dominance, Personal Good, and our background assumptions, on their own rule out incompleteness even in same-number cases. (The completeness of the social welfare ranking for these cases is an assumption of our first result.) The argument based on these

[1]Versions and close conceptual variants of this approach can be found in a wide array of works, including Seidenfeld et al. (1995); Shapley and Baucells (1986); Dubra et al. (2004); Nau (2006); Evren (2008); Evren and Ok (2011); Ok et al. (2012); Galaabaatar and Karni (2012, 2013); Riella (2015); Gorno (2017); Hara et al. (2019); McCarthy et al. (2021); Ok and Weaver (2023); Borie (2023).

[2]Danan et al. (2015) provide representations in the spirit of Harsanyi when both individual and social preferences are incomplete. We do not assume incompleteness over outcomes in individual preferences, we make stronger assumptions about the background space of possible outcomes for individuals, enforce completeness in same-number cases, and impose further axioms to guarantee the existence of an expected multi-utility.

two principles thus strengthens the case for completeness in social welfare, in both the same-number and variable-number cases.

Much of the paper is spent providing the characterization without Negative Dominance. But since Negative Dominance is a new axiom, and since its implications for this style of approach are striking, we provide the following Motivating Example to illustrate how Personal Good and Negative Dominance combine to force completeness. The example concerns two lotteries, each over two populations: in the first, p^*, there is a 0.5 chance of a population with three people (i, j, k) at welfare level u, and a 0.5 chance of a population with one person (i) at level u. (In the example, columns represent populations, and rows represent people.) In the second lottery, q^*, there is a 0.5 chance of a population with two people (i, j) at level u, and a 0.5 chance of a different two people (i, k) at level u.

$$\overbrace{\begin{bmatrix} u & u \\ u & \star \\ u & \star \end{bmatrix}}^{p^*} \sim \overbrace{\begin{bmatrix} u & u \\ u & \star \\ \star & u \end{bmatrix}}^{q^*} \qquad \text{(Motivating Example)}$$

These lotteries concern different sized populations. One might think that these populations should be unranked with respect to one another. But the following informal description of our axioms will show the key idea for how such incompleteness is ruled out.

Axiom 1* Personal Good (informal). *Consider two lotteries over variable-population social outcomes. If each potential person has the same probability of existence in both lotteries then:*

- *If each person faces the same individual prospect for lifetime utility in the two lotteries, then the two lotteries are equally good.*
- *If some people face individual prospects in the one lottery that stochastically dominate their individual prospects in the other, and everyone else faces the same individual prospects in the two lotteries, then the first lottery is better.*

Axiom 2* Negative Dominance (informal). *Consider two lotteries over variable-population social outcomes. If no outcome in the support of the first lottery is at least as good as any outcome in the support of the second lottery, then the first lottery is not strictly better than the second.*

This axiom was introduced in Lederman (2023) as a plausible dominance principle for incomplete betterness. One standard dominance principles states that, if every outcome in one lottery is strictly better than every outcome in another, then the first lottery is strictly better than the second. Negative Dominance states a similar idea for a lack of betterness between outcomes: there can be no strict betterness for one lottery by comparison to another without at least weak betterness between one outcome of the first and one outcome of the second.

These axioms already rule out particular forms of incompleteness in the above examples. In particular, Personal Good requires that the two lotteries are equivalent in the social welfare order. Any small improvement of any one u to $u + \varepsilon$ would make one of p^* or q^* better than the other, according to Personal Good. Negative Dominance then requires that some of the outcomes be comparable (that is, socially at least as good).

Multi-Utility is a natural framework for incomplete population ethics if the source of incompleteness is indeterminacy of the critical level of lifetime wellbeing for adding a new life. In this case, each ϕ could have a different critical level. For example, each ϕ could have the form of Critical-Level Total Utilitarianism, such that $\phi(\mathbf{u}) = \sum_i (u_i - c_\phi)$ for each outcome $\mathbf{u}$ with individual utilities u_i, but the critical utility level for adding a new life c_ϕ would differ across different ϕ in Φ. In other words, an economic analyst might be attracted to a multi-utility framework if she is comfortable with a standard, complete ordering of fixed-population, risk-free social distributions, but she also believes the critical level of utility to compare existence with non-existence is vague.

In our Motivating Example, as Proposition 1 below will demonate, Personal Good and an Expected Multi-Utility framework imply that for each ϕ,

$$\phi\left(\text{2 people with } u\text{-lives}\right) = \frac{1}{2}\phi\left(\text{1 person with a } u\text{-life}\right) + \frac{1}{2}\phi\left(\text{3 people with } u\text{-lives}\right).$$

This linear[3] relationship imports the additive structure of Expected Multi-Utility to create an additive structure for variable-population social welfare. Then, Negative Dominance restricts the incompleteness of final outcomes, so all ϕ agree. With these facts, we apply classic results from the population economics literature to characterize a complete and fully-additively-separable family of social welfare functions called Expected Critical-Level Generalized Utilitarianism. In so doing, we extensively build upon the definitions and results of Blackorby, Bossert, & Donaldson (2005), which hereafter we abbreviate BBD.

[3]In Jensen's functional equation, if $\frac{1}{2}\varphi(x) + \frac{1}{2}\varphi(y) = \varphi\left(\frac{1}{2}(x+y)\right)$, then φ is linear.

2 Setting and notation

Let $\mathbb{N}$ denote the set of positive integers, $\mathcal{N}$ the set of non-empty finite subsets of $\mathbb{N}$, $\mathbb{R}$ the set of real numbers, and $\mathbb{R}_{++}$ the set of positive real numbers. For a set D and any $n \in \mathbb{N}$, D^n is the n-fold Cartesian product of D. Also, for any two sets D and E, D^E denotes the set of mappings from E into D.

The set of *potential* individuals who may or may not exist is countable, and will be represented by the natural numbers $\mathbb{N}$. Only a finite non-empty subset of individuals exist in any realized outcome. That is, in any outcome, a population N exists: $N \in \mathcal{N}$.

We consider a welfarist framework where the only information necessary for social decisions is the utility levels of people alive in a certain state of affairs. An outcome's welfare information is given by $\mathbf{u} = (u_i)_{i \in N} \in \mathbb{R}^N$, where N is the population, and $u_i \in \mathbb{R}$ is, for each existing individual i, the lifetime utility experienced by i. Although we make use of the standard topology on $\mathbb{R}$ to assume continuity in utilities, beyond this we assume only that lifetime utilities are ordered and have the cardinality of the continuum.

We let $U = \bigcup_{N \in \mathcal{N}} \mathbb{R}^N$ denote the set of outcomes (which we will also call "populations") in which at least one individual exists. For each $\mathbf{u} \in U$, we denote $N(\mathbf{u})$ the set of individuals alive in $\mathbf{u}$, and $n(\mathbf{u})$ the number of individuals alive in $\mathbf{u}$. For two any outcomes $\mathbf{u}$ and $\mathbf{v}$ such that $N(\mathbf{u}) \cap N(\mathbf{v}) = \emptyset$ (that is, $\mathbf{u}$ and $\mathbf{v}$ are distributions of utility for two disjoint populations), we denote $\mathbf{uv}$ the outcome where the two populations are combined. Formally, it is the outcome $\mathbf{w}$ such that $N(\mathbf{w}) = N(\mathbf{u}) \cup N(\mathbf{v})$, $w_i = u_i$ for all $i \in N(\mathbf{u})$, and $w_j = v_j$ for all $j \in N(\mathbf{v})$. Abusing notation where the context allows, we will write $\mathbf{u}x$ for the population that adds one person at utility x to $\mathbf{u}$ and $x_i x_j$ for the two-person population consisting of one person at utility x_i and one at x_j. In the case where $N(\mathbf{u})$ and $N(\mathbf{v})$ have non-empty intersection, we let $\mathbf{uv}$ denote the combination of $\mathbf{u}$ with a permutation π of $\mathbf{v}$ such that $N(\mathbf{u})$ and $N(\pi(\mathbf{v}))$ have empty intersection. Similarly, $\overbrace{\mathbf{uu}\ldots\mathbf{u}}^{k \text{ times}}$ or $k\mathbf{u}$ will denote the combination of k disjoint permutations of $\mathbf{u}$.

We assume that it is not always known for sure what the final utility vector will be nor what set of individuals will exist. Our mixture space of interest is $\mathcal{L}$, the space of all finite lotteries on U. A lottery $p \in \mathcal{L}$ is a function $p : U \to [0, 1]$ such that $p(\mathbf{u}) > 0$ for only a finite number of $\mathbf{u} \in U$ and $\sum_{\mathbf{u} \in U} p(\mathbf{u}) = 1$. Where it is clear in context, we will sometimes write $\mathbf{u}$ for the risk-free lottery $p(\mathbf{u}) = 1$. For any finite set of outcomes $\mathbf{u}, \mathbf{v}, \ldots \mathbf{w} \in U$, we will write $\langle \mathbf{u}, \mathbf{v}, \ldots, \mathbf{w} \rangle$ for the lottery which assigns an equal probability to each outcome.

The goal of this paper is to characterize a social welfare preorder $\succsim$ on $\mathcal{L}$, with the normative interpretation that $p \succsim q$ means that p is at least as good as q and $p \succ q$ means that p is better than q. As a preorder, $\succsim$ is assumed to be reflexive and transitive, but we are specifically interested in the case where $\succsim$ may be incomplete.

Throughout we assume the following axioms for sure outcomes. Although these may be contentious in the philosophical population ethics literature, we expect that they will be noncontroversial in the population economics literature. We will refer to them below as the *Basic Axioms*.

Axiom 3 Same-people Pareto for sure outcomes. *For any* $\mathbf{u}, \mathbf{v} \in U$*, if* $N(\mathbf{u}) = N(\mathbf{v})$ *and if* $\mathbf{u} \geq \mathbf{v}$ *then* $\mathbf{u} \succsim \mathbf{v}$*; if, moreover,* $u_i > v_i$ *for some* $i \in N(\mathbf{u})$*, then* $\mathbf{u} \succ \mathbf{v}$*.*

Axiom 4 Same-people continuity-in-utilities for sure outcomes. *For any* $\mathbf{u} \in U$*, the sets* $\{\mathbf{v} \in U | N(\mathbf{v}) = N(\mathbf{u})$ *and* $\mathbf{v} \succsim \mathbf{u}\}$ *and* $\{\mathbf{v} \in U | N(\mathbf{v}) = N(\mathbf{u})$ *and* $\mathbf{v} \precsim \mathbf{u}\}$ *are closed in* $\mathbb{R}^{N(\mathbf{u})}$*.*

Axiom 5 Same-number anonymity for sure outcomes. *For any* $\mathbf{u}, \mathbf{v} \in U$*, if* $n(\mathbf{u}) = n(\mathbf{v})$ *and there exists a bijection* $\sigma : N(\mathbf{u}) \to N(\mathbf{v})$ *such that* $v_{\sigma(i)} = u_i$ *for all* i*, then* $\mathbf{u} \sim \mathbf{v}$*.*

3 Background: Incomplete population ethics and expected multi-utility

3.1 Variable-population social preorders

Definition 1 Critical-Set Generalized Utilitarianism. *There exists a continuous and increasing function* $f : \mathbb{R} \to \mathbb{R}$ *and a nonempty critical set for lifetime utility* $S \subseteq \mathbb{R}$*, such that, for any* $\mathbf{u}, \mathbf{v} \in U$*,*

$$\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \forall c \in S, \sum_{i \in N(\mathbf{u})} (f(u_i) - c) \geq \sum_{i \in N(\mathbf{v})} (f(v_i) - c) .$$

Because $\succsim$ satisfies Pareto and transitivity, S is convex; the convexity of S is established by BBD's Theorem 7.11. Because S is convex, if S is bounded, S is either one point or is an interval. If the critical set is a bounded, non-point interval, then $\succsim$ is Critical-Band Generalized Utilitarianism (BBD, p. 250). If the critical set is a point, then $\succsim$ is complete and it is Critical-Level Generalized Utilitarianism (CLGU), which was first characterized by Blackorby and Donaldson (1984).

Definition 2 Critical-Level Generalized Utilitarianism (BBD p. 166). *There exists a continuous and increasing function* $f : \mathbb{R} \to \mathbb{R}$ *and a critical-level for lifetime utility* $c \in \mathbb{R}$ *such that for any* $\mathbf{u}, \mathbf{v} \in U$*,*

$$\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \sum_{i \in N(\mathbf{u})} (f(u_i) - c) \geq \sum_{i \in N(\mathbf{v})} (f(v_i) - c) .$$

If no claim is made about the size of the critical set, we have Same-Number Generalized Utilitarianism:

Definition 3 Same-Number Generalized Utilitarianism (BBD p. 185). *There exists a continuous and increasing function $f : \mathbb{R} \to \mathbb{R}$ such that, and any $\mathbf{u}, \mathbf{v} \in U$ such that $n(\mathbf{u}) = n(\mathbf{v})$,*

$$\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \sum_{i \in N(\mathbf{u})} f(u_i) \geq \sum_{i \in N(\mathbf{v})} f(v_i).$$

In each case, we can extend the representation to lotteries by defining an Expected version.

Definition 4 Expected Critical-Set Generalized Utilitarianism. *There exists a continuous and increasing function $f : \mathbb{R} \to \mathbb{R}$ and a critical set for lifetime utility $S \subseteq \mathbb{R}$, such that, for any $p, q \in \mathcal{L}$,*

$$p \succsim q \Leftrightarrow \forall c \in S, \sum_{\mathbf{u} \in U} p(\mathbf{u}) \sum_{i \in N(\mathbf{u})} \left(f(u_i) - c \right) \geq \sum_{\mathbf{u} \in U} q(\mathbf{u}) \sum_{i \in N(\mathbf{u})} \left(f(u_i) - c \right).$$

Definition 5 Expected Critical-Level Generalized Utilitarianism. *There exists a continuous and increasing function $f : \mathbb{R} \to \mathbb{R}$ and a critical level for lifetime utility $c \in \mathbb{R}$, such that, for any $p, q \in \mathcal{L}$,*

$$p \succsim q \Leftrightarrow \sum_{\mathbf{u} \in U} p(\mathbf{u}) \sum_{i \in N(\mathbf{u})} \left(f(u_i) - c \right) \geq \sum_{\mathbf{u} \in U} q(\mathbf{u}) \sum_{i \in N(\mathbf{u})} \left(f(u_i) - c \right).$$

Gustafsson (2020) informally described a version of Expected Critical-Set Generalized Utilitarianism (p. 94), based on a proposal for "undistinguishedness" as a category of absolute value, but without an axiomatic characterization. Notice that there's a version of Expected Critical-Set Generalized Utilitarianism and Expected Same-Number Generalized Utilitarianism where the set of critical levels is the set of all reals (or, if f is bounded, is a superset of the image of f), so it's incomplete for any different-number comparison. Restricted to comparisons of sure outcomes, this resembles Bader (2022).

3.2 Expected Multi-Utility

Definition 6 Multi-Utility. *There exists a set Φ of functions $\phi : U \to \mathbb{R}$ such that, for any $\mathbf{u}, \mathbf{v} \in U$, $\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \forall \phi \in \Phi, \phi(\mathbf{u}) \geq \phi(\mathbf{v})$.*

Notice that $\succsim$ satisfies Critical-Set Generalized Utilitarianism if and only if $\succsim$ has a Multi-Utility representation in which each ϕ is Critical-Level Generalized Utilitarian. Moreover, all $\phi \in \Phi$ share the same function f and differ only in their critical levels c; $|\Phi| = |S|$, and for each $c \in S$ there is a $\phi \in \Phi$ with critical level c.

Definition 7 Expected Multi-Utility. *There exists a set Φ of functions $\phi : U \to \mathbb{R}$ such that, for any $p, q \in \mathcal{L}$,*

$$p \succsim q \Leftrightarrow \forall \phi \in \Phi, \sum_{\mathbf{u} \in U} p(\mathbf{u})\phi(\mathbf{u}) \geq \sum_{\mathbf{u} \in U} q(\mathbf{u})\phi(\mathbf{u}).$$

Each ϕ defines a complete and transitive ordering of U, which we will write as $\succsim_\phi$, such that, for all $\mathbf{u}, \mathbf{v} \in U, \mathbf{u} \succsim_\phi \mathbf{v} \Leftrightarrow \phi(\mathbf{u}) \geq \phi(\mathbf{v})$.

As noted earlier, Expected Multi-Utility can be derived from standard axiomatic constraints on preferences (dropping completeness). Consider the following three standard axioms:

Strong Independence $p \succsim q$ if and only if $\alpha p + (1-\alpha)r \succsim \alpha q + (1-\alpha)r$ for $\alpha \in (0,1)$.

Mixture Continuity If for all $\alpha \in (0,1]$, $\alpha p + (1-\alpha)r \succ q$, then $p \succsim q$.

Archimedeanness If $p \succ q \succ r$ then there are $\alpha, \beta \in (0,1)$ such that $\alpha p + (1-\alpha)r \succ q$ and $q \succ \beta p + (1-\beta)r$.

By McCarthy et al. (2021) Theorem 2.4, these imply the existence of an Expected Multi-Utility representation (see their Theorem 2.5 for alternative axiomatization, and also Dubra et al. (2004)).

4 Critical-Set Generalized Utilitarianism: A Characterization without Negative Dominance

The main result of this section is that Expected Multi-Utility, Personal Good, and the Basic Axioms characterize Critical-Set Generalized Utilitarianism. Two basic building blocks are Existence Independence and Extended Replication Invariance, which are important properties of population ethics for risk-free variable-population social outcomes. Here we show that Personal Good is sufficient for each $\succsim_\phi$ to have these two properties, if $\succsim$ satisfies the Basic Axioms and if it admits an Expected Multi-Utility Representation.

To state Personal Good requires notation for individual prospects, conditional on existence. First, for any $p \in \mathcal{L}$, let $N(p)$ denote the set of individuals with a positive probability of existence in p, i.e., the set of i such that $i \in N(\mathbf{u})$ for some $\mathbf{u}$ in the support of p. Then, for any $i \in N(p)$, let $p_i : \mathbb{R} \to [0,1]$ exist and be defined as $p_i(x) = \frac{\sum_{\{\mathbf{u} \in U : u_i = x\}} p(\mathbf{u})}{\sum_{\{\mathbf{u} \in U : i \in N(\mathbf{u})\}} p(\mathbf{u})}$. This represents the probability, on lottery p, that individual i receives welfare level x, conditional on i existing.

Axiom 1 Personal Good. *Suppose $p, q \in \mathcal{L}$ have the property that each possible person has the same probability of existing in either lottery—that is, for all $i \in \mathbb{N}$, $\sum_{\{\mathbf{u} \in U : i \in N(\mathbf{u})\}} p(\mathbf{u}) = \sum_{\{\mathbf{u} \in U : i \in N(\mathbf{u})\}} q(\mathbf{u})$. Then:*

- *if $p_i = q_i$ for all $i \in N(p)$ (or equivalently, all $i \in N(q)$), then $p \sim q$;*
- *if there exists a non-empty set of possible people $M \subseteq N(p)$ such that, for all $i \in M$, p_i first-order stochastically dominates q_i, and for all remaining $j \in N(p)$, $p_j = q_j$, then $p \succ q$.*[4]

Personal Good allows us to show that each $\succsim_\phi$ satisfies Existence Independence.[5] We do this by an argument similar to those in Thomas (2022) and Gustafsson et al. (2023).

Definition 8 Existence Independence (BBD p. 159). *For any three outcomes $\mathbf{u}, \mathbf{v}, \mathbf{w} \in U$ with $N(\mathbf{w})$ disjoint from $N(\mathbf{u})$ and $N(\mathbf{v})$, $\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \mathbf{uw} \succsim \mathbf{vw}$.*

Lemma 1. *If $\succsim$ has an* Expected Multi-Utility *representation and satisfies* Personal Good *and the* Basic Axioms *then:*

- *for all $\phi \in \Phi$, $\succsim_\phi$ satisfies* Existence Independence*, and therefore,*
- *$\succsim$ satisfies* Existence Independence*.*

Proof. Choose any ϕ. Then $\succsim_\phi$ is a complete ordering of U. Assume for indirect proof that there exist $\mathbf{u}, \mathbf{v}, \mathbf{w} \in U$ such that Existence Independence fails, and without loss of generality that $\mathbf{u} \succsim_\phi \mathbf{v}$. Then, by adding the real-valued ϕ-terms, it must be that $\phi(\mathbf{u}) + \phi(\mathbf{vw}) > \phi(\mathbf{v}) + \phi(\mathbf{uw})$. If so, then by Expected Multi-Utility, it cannot be that $\langle \mathbf{uw}, \mathbf{v} \rangle \succsim \langle \mathbf{u}, \mathbf{vw} \rangle$, but this is required by Personal Good (which holds that $\langle \mathbf{uw}, \mathbf{v} \rangle \sim \langle \mathbf{u}, \mathbf{vw} \rangle$), so there is a contradiction. The reverse direction (from $\mathbf{wu} \succsim \mathbf{wv}$ to $w \succsim v$) follows by a similar argument.

It then follows that $\succsim$ satisfies *Existence Independence*. If $\mathbf{u} \succsim \mathbf{v}$ then, for each ϕ, $\mathbf{u} \succsim_\phi \mathbf{v}$; so, for each ϕ, $\mathbf{uw} \succsim_\phi \mathbf{vw}$; so $\mathbf{uw} \succsim \mathbf{vw}$. These are equivalences so the reverse direction holds as well. □

With Existence Independence, we can show that each $\succsim_\phi$ satisfies Extended Replication Invariance.

[4]The equal-prospect part of Personal Good is what McCarthy (2017), McCarthy et al. (2020), and Thomas (2022) call Anteriority.

[5]This principle is often called "Separability", especially in the philosophical literature.

Definition 9 Extended Replication Invariance (BBD p. 165). *For any two outcomes* $\mathbf{u}, \mathbf{v} \in U$ *and any* $k \in \mathbb{N}$, $\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k\ times} \succsim \overbrace{\mathbf{vv} \ldots \mathbf{v}}^{k\ times}$.

Lemma 2. *If* $\succsim$ *has an* Expected Multi-Utility *representation and satisfies* Personal Good *and the* Basic Axioms *then*

- *each* $\succsim_\phi$ *satisfies* Extended Replication Invariance, *and therefore,*
- $\succsim$ *satisfies* Extended Replication Invariance.

Proof. Choose any ϕ and any $\mathbf{u}, \mathbf{v} \in U$ such that $\mathbf{u} \succsim_\phi \mathbf{v}$ and any $k \in \mathbb{N}$. For use in the proof, choose any $\mathbf{w} \in U$. By Existence Independence, $\mathbf{u} \succsim_\phi \mathbf{v} \Leftrightarrow \mathbf{uw} \succsim_\phi \mathbf{vw}$ and $\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \mathbf{uw} \succsim \mathbf{vw}$. Because $\succsim$ satisfies Personal Good, each $\succsim_\phi$ must be consistent with it. So by Personal Good, Anonymity, and Expected Multi-Utility,

$$\langle \overbrace{\mathbf{uw}, \mathbf{uw}, \ldots, \mathbf{uw}}^{k \text{ times}} \rangle \sim \langle \mathbf{w} \overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k \text{ times}}, \overbrace{\mathbf{w}, \mathbf{w}, \ldots, \mathbf{w}}^{k-1 \text{ times}} \rangle, \text{ and so}$$

$$k\phi(\mathbf{uw}) = \phi\left(\mathbf{w} \overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k \text{ times}}\right) + (k-1)\phi(\mathbf{w}).$$

and similarly for $\mathbf{v}$. $\succsim_\phi$ is a complete and transitive order, so we now have that

$$k\phi(\mathbf{uw}) \geq k\phi(\mathbf{vw}) \Leftrightarrow \phi\left(\mathbf{w} \overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k \text{ times}}\right) + (k-1)\phi(\mathbf{w}) \geq \phi\left(\mathbf{w} \overbrace{\mathbf{vv} \ldots \mathbf{v}}^{k \text{ times}}\right) + (k-1)\phi(\mathbf{w}),$$

where the latter inequality is equivalent to

$$\phi\left(\mathbf{w} \overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k \text{ times}}\right) \geq \phi\left(\mathbf{w} \overbrace{\mathbf{vv} \ldots \mathbf{v}}^{k \text{ times}}\right).$$

And finally by Existence Independence, that is equivalent to $\overbrace{\mathbf{uu} \ldots \mathbf{u}}^{k \text{ times}} \succsim_\phi \overbrace{\mathbf{vv} \ldots \mathbf{v}}^{k \text{ times}}$. Because all steps are equivalences, the proof holds in both directions. And because Extended Replication Invariance holds for all ϕ, it holds for $\succsim$. □

Although an Expected Multi-Utility approach would seem to be permissive, these results show that Personal Good considerably narrows the possibilities for members of Φ. That is because Existence Independence and Extended Replication Invariance are known to rule out many families of social welfare functions that have been named and investigated in the population ethics literature. Any type of CLGU (including Total

Utilitarianism) is consistent with both. But Average Utilitarianism, although consistent with Extended Replication Invariance, is not consistent with Existence Independence. In "Variable-Value" social welfare functions, adding lives at a given welfare level has diminishing marginal impact on social welfare (Hurka, 1983; Bossert et al., 2023). These typically reject both Existence Independence and Extended Replication Invariance (Blackorby et al., 2005). Both Number-Dampened Generalized Utilitarianism (Ng, 1989; Spears and Stefánsson, 2024) and Rank-Discounted Generalized Utilitarianism (Asheim and Zuber, 2014; Pivato, 2020; Spears and Stefánsson, 2021) reject both of these principles. So, even in a Multi-Utility framework, Personal Good rules out that any of the $\succsim_\phi$ can have an Average or Variable-Value form.

We use one final assumption to pin down a representation for $\succsim$. So far, we have characterized $\succsim_\phi$ without making any assumptions about the completeness of $\succsim$. But we now assume that fixed-population comparisons are complete. In doing this, our goal is to set aside the possibility of incompleteness arising for reasons *other than population ethics*. If the functions f mapping lifetime utility into social value differ across ϕ (perhaps because different ϕ have different inequality aversion or risk aversion), then $\succsim$ could be incomplete even for fixed-population comparisons. But this would be a distraction from the population ethics focus of this paper, so we introduce an axiom to eliminate this possibility.

Axiom 7 Same-Population Completeness. *For any* $\mathbf{u}, \mathbf{v} \in U$ *such that* $N(\mathbf{u}) = N(\mathbf{v})$, *either* $\mathbf{u} \succsim \mathbf{v}$, *or* $\mathbf{u} \precsim \mathbf{v}$, *or both.*

Lemma 3. *If* $\succsim$ *has an* Expected Multi-Utility *representation and satisfies* Personal Good, Same-Population Completeness *and the* Basic Axioms, *then* $\succsim$ *restricted to certain outcomes is* Same-Number Generalized Utilitarian. *Moreover, each* $\succsim_\phi$ *is* Same-Number Generalized Utilitarian.

Proof. This follows immediately from BBD's Theorem 6.2 because all of the conditions are met. In our notation, it says: $\succsim$ [which they elsewhere assume to be a complete, anonymous social order] satisfies continuity, strong Pareto, same-number independence [which is implied by Existence Independence], and replication invariance [which is implied by Extended Replication Invariance] if and only if $\succsim$ is Same-Number Generalized Utilitarian. □

The fact that $\succsim$ restricted to certain outcomes is Same-Number Generalized Utilitarianism has an important implication. Even if all ϕ were additively separable, they could have had different sub-functions f. Personal Good rules this disagreement out and collapses all ϕ down to one shared f. With this, we are ready for our main result:

Proposition 1. *The following statements are equivalent:*

- $\succsim$ *has an* Expected Multi-Utility *representation and satisfies* Personal Good, Same-Population Completeness *and the* Basic Axioms*.*
- $\succsim$ *is* Expected Critical-Set Generalized Utilitarianism*.*

If, moreover, there exists $\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{x} \in U$ *and* $a, b, c, d \in \mathbb{R}$ *such that* $\mathbf{u}a \succsim \mathbf{u}$*,* $\mathbf{v}b \precsim \mathbf{v}$*,* $c \neq d$*,* $\mathbf{w}c$ *is not ranked against* $\mathbf{w}$*, and* $\mathbf{x}d$ *is not ranked against* $\mathbf{x}$*, then* $\succsim$ *is* Expected Critical-Band Generalized Utilitarianism*.*

Proof. That the representation implies the axioms is trivial, so we will prove that the axioms imply the representation. We will show that each $\succsim_\phi$ is CLGU with the same f but potentially with different critical levels. Because Expected Multi-Utility holds that, for any $\mathbf{u}, \mathbf{v}$ $\mathbf{u} \succsim \mathbf{v}$ implies $\mathbf{u} \succsim_\phi \mathbf{v}$, and in the context of assuming Same-Population Completeness, each $\succsim_\phi$ inherits the Basic Axioms.

First we will show that each ϕ is additive CLGU on the set of outcomes U. Choose any $\mathbf{u}, \mathbf{v} \in U$. Without loss of generality, let $n(\mathbf{v}) \geq n(\mathbf{u})$. Fix any $x \in \mathbb{R}$. By Existence Independence, $\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \mathbf{u}x \succsim \mathbf{v}x$. Because each ϕ is expectational, this is true if and only if

$$\langle \mathbf{u}x, \overbrace{x, \ldots, x}^{n(\mathbf{v})-1 \text{ times}} \rangle \succsim \langle \mathbf{v}x, \overbrace{x, \ldots, x}^{n(\mathbf{v})-1 \text{ times}} \rangle$$

By Personal Good,

$$\mathbf{u} \succsim \mathbf{v} \Leftrightarrow \langle xu_1, xu_2, \ldots, xu_{n(\mathbf{u})}, \overbrace{x, \ldots, x}^{n(\mathbf{v})-n(\mathbf{u}) \text{ times}} \rangle \succsim \langle xv_1, xv_2, \ldots xv_{n(\mathbf{v})} \rangle.$$

So by Expected Multi-Utility, $\mathbf{u} \succsim \mathbf{v}$ if and only if for all ϕ,

$$\sum_{i=1}^{n(\mathbf{u})} \phi(xu_i) + (n(\mathbf{v}) - n(\mathbf{u}))\,\phi(x) \geq \sum_{i=1}^{n(\mathbf{v})} \phi(xv_i),$$

or, rearranging,

$$\sum_{i=1}^{n(\mathbf{u})} \left(\phi(xu_i) - \phi(x)\right) \geq \sum_{i=1}^{n(\mathbf{v})} \left(\phi(xv_i) - \phi(x)\right).$$

Now notice that, for any $u_i, x \in \mathbb{R}$, Personal Good and anonymity entail that $\langle xx, u_i \rangle \sim \langle x, xu_i \rangle$, so $\phi(xx) + \phi(u_i) = \phi(x) + \phi(xu_i)$, or, rearranging, $\phi(xu_i) = \phi(u_i) + \phi(xx) - \phi(x)$.

Then, substituting into the representation:

$$\sum_{i=1}^{n(\mathbf{u})} \left(\phi(u_i) + \phi(xx) - 2\phi(x)\right) \geq \sum_{i=1}^{n(\mathbf{v})} \left(\phi(v_i) + \phi(xx) - 2\phi(x)\right).$$

By Same-Number Generalized Utilitarianism, ϕ must rank in accord with f on single-person populations, so we can let $\phi(u_i) = f(u_i)$ and still represent the order. As a verification that the change in ϕ from adding any life at utility x is a constant, independent of population size, notice that

$$\forall n \in \mathbb{N}, \forall x \in \mathbb{R}, \forall \mathbf{u} \in U, \langle \mathbf{u}\overbrace{xx\ldots x}^{n+1 \text{ times}}, \mathbf{u}\rangle \sim \langle \mathbf{u}\overbrace{xx\ldots x}^{n \text{ times}}, \mathbf{u}x\rangle,$$

so

$$\phi\left(\mathbf{u}\overbrace{xx\ldots x}^{n+1 \text{ times}}\right) - \phi\left(\mathbf{u}\overbrace{xx\ldots x}^{n \text{ times}}\right) = \phi(\mathbf{u}x) - \phi(\mathbf{u}) = f(x) + \kappa.$$

for some κ. Therefore the remaining term can be absorbed into a constant:

$$\sum_{i \in N(\mathbf{u})} \left(f(u_i) - c_\phi\right) \geq \sum_{i \in N(\mathbf{v})} \left(f(v_i) - c_\phi\right)$$

So, each ϕ is CLGU. Step 1 has established a Critical-Set Generalized Utilitarian representation on U with a single continuous and increasing f.

What about lotteries? By the assumption of Expected Multi-Utility, we have that $\succsim$ over $\mathcal{L}$ is represented by

$$p \succsim q \Leftrightarrow \forall \phi \in \Phi, \sum_{\mathbf{u} \in U} p(\mathbf{u})\phi(\mathbf{u}) \geq \sum_{\mathbf{u} \in U} q(\mathbf{u})\phi(\mathbf{u}).$$

We next verify that if any ϕ is a transformed function of the representation on U, then that transformation is affine. That is, if $\phi(\mathbf{u}) = g\left(\sum_{i=1}^{n(\mathbf{u})} \left(f(u_i) - c_\phi\right)\right)$ for some increasing $g : \mathbb{R} \to \mathbb{R}$, then g is affine. To see this, choose any $\mathbf{u}, \mathbf{v}, \mathbf{w} \in U$. Notice that $\langle \mathbf{uv}, \mathbf{w}\rangle \sim \langle \mathbf{u}, \mathbf{wv}\rangle$. Therefore

$$g\left(\phi(\mathbf{u}) + \phi(\mathbf{v})\right) + g\left(\phi(\mathbf{w})\right) = g\left(\phi(\mathbf{u})\right) + g\left(\phi(\mathbf{w}) + \phi(\mathbf{v})\right),$$

which rules out curvature of g.

We have therefore characterized Expected Critical-Set Generalized Utilitarianism, with $S = \{c_\phi | \phi \in \Phi\}$. The restriction to Expected Critical Band-Utilitarianism further requires S to be bounded. This final step uses BBD's Theorem 7.11. Using a, b, c, d in the statement of the Proposition, a is an upper bound on S, b is a lower bound on S, and S is convex, establishing the characterization of Expected Critical-Band Generalized Utilitarianism.

□

As intuition for the characterization, apply the logic of Gustafsson et al. (2023), in which each lottery with only rational probabilities is shown to be equivalent to a large, certain population. Each $\succsim_\phi$ is CLGU. So p and q would each be as good, according to ϕ, as fixed-population lotteries created by replacing non-existence in a state with existence at $f^{-1}(c_\phi)$ for any person who exists in any state of either lottery. These fixed-population versions of p and q would also be as good as ones where the number of supported outcomes is expanded to a large number of equiprobable states. And by $\succsim_\phi$ respecting Personal Good, these are as good as all existences happening in one outcome. For example, where $a, b, d \in \mathbb{R}$ are arbitrary lifetime utilities, as before rows are potential people and columns are outcomes, and the numbers at the top are probabilities of an outcome, writing c_ϕ^{-1} for $f^{-1}(c_\phi)$:

$$\overbrace{\begin{bmatrix} \frac{1}{3} & \frac{2}{3} \\ \hline a & \star \\ \star & b \end{bmatrix}}^{p} \succsim \overbrace{\begin{bmatrix} 1 \\ \hline d \\ d \end{bmatrix}}^{q} \Leftrightarrow \forall \phi, \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \hline a & c_\phi^{-1} & c_\phi^{-1} \\ c_\phi^{-1} & b & c_\phi^{-1} \\ c_\phi^{-1} & c_\phi^{-1} & b \\ c_\phi^{-1} & c_\phi^{-1} & c_\phi^{-1} \\ c_\phi^{-1} & c_\phi^{-1} & c_\phi^{-1} \\ c_\phi^{-1} & c_\phi^{-1} & c_\phi^{-1} \end{bmatrix} \succsim \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \hline d & c_\phi^{-1} & c_\phi^{-1} \\ d & c_\phi^{-1} & c_\phi^{-1} \\ c_\phi^{-1} & d & c_\phi^{-1} \\ c_\phi^{-1} & d & c_\phi^{-1} \\ c_\phi^{-1} & c_\phi^{-1} & d \\ c_\phi^{-1} & c_\phi^{-1} & d \end{bmatrix} \Leftrightarrow \forall \phi, \begin{bmatrix} 1 \\ \hline a \\ b \\ b \\ c_\phi^{-1} \\ c_\phi^{-1} \\ c_\phi^{-1} \end{bmatrix} \succsim_\phi \begin{bmatrix} 1 \\ \hline d \\ d \\ d \\ d \\ d \\ d \end{bmatrix}.$$

The last comparison is a comparison within U, on which $\succsim_\phi$ is represented by CLGU. All ϕ share the same f, so the only scope for incompleteness that remains is differences among c_ϕ.

Notice that none of this necessitates the existence of any different-population-size comparability, because it could be that $S = \mathbb{R}$ (cf. Bader, 2022). But consider this example, where rows are possible people:

$$\begin{bmatrix} \frac{1}{2} + \varepsilon & \frac{1}{2} - \varepsilon \\ \hline 2 & \star \\ 2 & 2 \end{bmatrix} \succ \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \hline 1 & \star \\ 1 & 1 \end{bmatrix}.$$

If $S = \mathbb{R}$, then the higher-average-utility population is better only if $\varepsilon = 0$, but not if $\varepsilon > 0$. This seems implausible, and if so would be an argument for bounding S.

5 Critical-Level Generalized Utilitarianism: A Characterization with Negative Dominance

Axiom 2 Negative Dominance. *For any lotteries p, q, if for all $\mathbf{u}$ such that $p(\mathbf{u}) > 0$, and all $\mathbf{v}$ such that $q(\mathbf{v}) > 0$ is not the case that $\mathbf{u} \succsim \mathbf{v}$, then it is not the case that $p \succ q$.*

If $\succsim$ is complete, then Negative Dominance is implied by the universally assumed:

Axiom 7 Min-Max Dominance. *If for every $\mathbf{u}$ in the support of p and every $\mathbf{v}$ in the support of q, $\mathbf{u} \succ \mathbf{v}$, then $p \succ q$.*

In the absence of completeness, the implication no longer holds. But both dominance principles are plausible, independently of whether completeness is assumed. In particular, if there is no outcome in the support of p which is even weakly better than any outcome in the support of q, then it is hard to explain why p would be better than q—at least, some very natural explanations are unavailable to us, e.g., that p might yield a better outcome than q or increases our chance of an outcome at least as good as some desired threshold. We discuss the connection between Negative Dominance and some other relevant axioms in Section 6.1.

In population ethics, an Expected Multi-Utility representation need not satisfy Negative Dominance. Consider an example where one ϕ is Average Utilitarianism and another ϕ is Total Utilitarianism. Let outcomes a, b, c now be three populations with average utility 1, 2, and 5, respectively and population sizes 50, 10, and 2, respectively. Then:

$$\phi_{\text{AU}}(a) = 1. \quad \phi_{\text{AU}}(b) = 2. \quad \phi_{\text{AU}}(c) = 5.$$
$$\phi_{\text{TU}}(a) = 50. \quad \phi_{\text{TU}}(b) = 20. \quad \phi_{\text{TU}}(c) = 10.$$

If lotteries are valued expectationally, then Negative Dominance would be violated.

Proposition 1 rules out this example, however, because Personal Good together with Expected Multi-Utility rule out average utilitarianism (Gustafsson and Spears, 2022; Gustafsson et al., 2023). What are the implications of Negative Dominance if Personal Good is assumed? Adding Negative Dominance to Proposition 1 collapses S to a single critical level and ensures completeness, even while dropping the Same-Population Completeness axiom (used in Proposition 1).

Proposition 2. *The following statements are equivalent:*

- $\succsim$ *has an* Expected Multi-Utility *representation and satisfies* Personal Good*, the* Basic Axioms*, and* Negative Dominance*.*
- $\succsim$ *is* Expected Critical-Level Generalized Utilitarianism*.*

Proof. First we show that if $\succsim$ has an *Expected Multi-Utility Representation*, and satisfies *Personal Good*, the *Basic Axioms* and *Negative Dominance*, then it satisfies *Same-Number Completeness*. Suppose for contradiction that for some $\mathbf{u}, \mathbf{v}$ with $N(\mathbf{u}) = N(\mathbf{v})$ neither $\mathbf{u} \succsim \mathbf{v}$ nor $\mathbf{v} \succsim \mathbf{u}$. By *Same People Continuity for Sure Outcomes*, there is a $\mathbf{u}^+$ with $N(\mathbf{u}) = N(\mathbf{u}^+)$, $u_i^+ > u_i$ for some i, $u_j^+ \geq u_j$ for all $j \neq i$ and such that neither $\mathbf{u}^+ \succsim \mathbf{v}$ nor $\mathbf{v} \succsim \mathbf{u}^+$. By Existence Independence (which, by Lemma 1, is implied by our assumptions), neither $\mathbf{u}^+\mathbf{u} \succsim \mathbf{vu}$, nor $\mathbf{vu} \succsim \mathbf{u}^+\mathbf{u}$. By Anonymity, the same holds for $\mathbf{uv}$. A similar argument holds for the relationship between $\mathbf{vu}$, $\mathbf{uv}$ and $\mathbf{vv}$. Now consider, in line with the motivating example:

$$\overbrace{\begin{bmatrix} \mathbf{u} & \mathbf{v} \\ \mathbf{v} & \mathbf{u} \end{bmatrix}}^{\hat{p}} \qquad \overbrace{\begin{bmatrix} \mathbf{u}^+ & \mathbf{v} \\ \mathbf{u} & \mathbf{v} \end{bmatrix}}^{\hat{q}}$$

Personal Good requires that $\hat{q} \succ \hat{p}$, but Negative Dominance requires that $\hat{q} \not\succ \hat{p}$, contradicting our assumption. So $\succsim$ satisfies *Same-Number Completeness*.

By Proposition 1, it follows that these assumptions imply Expected Critical Set Generalized Utilitarianism. Suppose now for contradiction there are $c, c' \in S$ with $c > c'$. Recall that, by BBD's Theorem 7.11 and the Basic Axioms, S is a convex set. Now choose $u, v \in (f^{-1}(c), f^{-1}(c'))$ with $u < v$ and $f(v) - f(u)$ less than both $c - f(v)$ and $f(v) - c'$.. Now consider, in line with the Motivating Example, the following pair of lotteries:

$$\overbrace{\begin{bmatrix} u & u \\ v & \star \\ v & \star \end{bmatrix}}^{p^{**}} \qquad \overbrace{\begin{bmatrix} u & u \\ u & \star \\ \star & u \end{bmatrix}}^{q^{**}}$$

Personal Good implies that $p^{**} \succ q^{**}$. Expected Critical Set Generalized Utilitarianism implies that the outcomes in the support of p^{**} are incomparable with the outcomes in the support of q^{**}. (Since $f(v) - f(u) < c - f(v)$, the critical level c makes uvv worse than uu. Since $f(v) - f(u) < f(v) - c'$, the critical level c' makes uvv better than uu. Thus, uvv and uu are incomparable. Since $u \in (f^{-1}(c), f^{-1}(c'))$, u is likewise incomparable with uu.)

So by Negative Dominance, neither lottery can be better than the other, contradicting our assumption that there are distinct $c, c' \in S$. □

It may be that there is no u such that $f(u) = c$, if f is bounded. A further axiom can rule this possibility out, in the context of the other assumptions of Proposition 2:

Axiom 8 Population non-absolutism for sure outcomes. *There exists* $\mathbf{u} \in U$ *and* $a, b \in \mathbb{R}$ *such that* $\mathbf{u}a \not\succ \mathbf{u}$ *and* $\mathbf{u} \not\succ \mathbf{u}b$*.*

Population non-absolutism for sure outcomes rules out that the c of CLGU is greater than, or less than, every element of the image of f.

6 Conclusion and Discussion

6.1 Why Negative Dominance?

Manzini and Mariotti (2008) introduce the following axiom:

Axiom 9 Vagueness Sure Thing (VST). *If* p *is incomparable with* q *and* r *is incomparable with* s*, then* $p\alpha r$ *is incomparable with* $q\alpha s$*.*

As they observe (p. 310), this axiom rules out Stochastic Dominance, provided there are outcomes a, a^+, b, b^+ such that a^+ is incomparable with b, a incomparable with b^+, but a^+ is strictly better than a and b^+ is strictly better than b.

Hare (2010, 2013) presents a weaker axiom, in a Savage-style framework, where actions are functions from states to outcomes:

Axiom 10 Statewise Negative Dominance. *For any actions* f, g*, if for every state* s*,* $f(s)$ *is not weakly better than* $g(s)$*, then* f *is not strictly better than* g*.*

Although this principle is intuitively weaker than the VST because of its restriction to comparisons in the same states, Hare provides an example that Bader (2018) shows presents a conflict between Statewise Negative Dominance and Stochastic Dominance, provided there are again outcomes a, a^+, b, b^+, as above.

The assumption that such quadruples exist is a very modest richness assumption, which is natural in applications like the ones we have considered. Stochastic Dominance is almost universally assumed. So the incompatibility of these earlier axioms with Stochastic Dominance is a strong mark against them. Negative Dominance is consistent with Stochastic Dominance in general (see Lederman (2023)). We take this to show that it is interestingly weaker than the earlier axioms. However, as we have shown, it is still strong enough to imply completeness with Personal Good in the background.

6.2 Related Literature Weakening Harsanyi's Approach to Welfare Economics Under Risk

Variable-population incompleteness has long been a prominent threat to finding a consensus in population ethics that could be applied to questions of economic and policy importance. Parfit's (1986) Depletion case (§123) offered a powerful early response, based on the substantive implausibility of the implication that different human futures with radically different wellbeing—but also different populations—could not be compared as better or worse. This paper offers a different response, arguing from axiomatic population ethics, rather than cases.

Our response joins an active recent literature, following Fleurbaey (2009), that has weakened Harsanyi's aggregation theorem for social choice under risk. These include Danan et al. (2015), McCarthy et al. (2020), Spears and Zuber (2022), Li et al. (2023), Gustafsson et al. (2023). A theme of this literature is that the assumptions of Harsanyi's project can be considerably weakened: Additively-separable social evaluation can result from many different sets of plausible decision-making principles applied to risky social outcomes.

One question is what are the most plausible options that remain for a theorist seeking to reject an additively-separable, generalized-totalist social welfare function.[6] One option is to reject Negative Dominance or Expected Multi-Utility as requirements for social decision-making under risk. But Negative Dominance is plausible, and each ϕ is additively-separable CLGU even without it. As long as S is bounded, dropping Negative Dominance would still result in Expected Critical-Band Generalized Utilitarianism. And Expected Multi-Utility can be derived from weak principles (McCarthy et al., 2021).

Another option would be to reject Personal Good, perhaps in the egalitarian spirit of Fleurbaey (2010). But Gustafsson et al. (2023) have recently offered an alternative axiomatzation of Expected CLGU that uses a weakened axiom of Personal Good which, considered in isolation, is compatible with egalitarian choice for fixed-population social risk. And while incompleteness is not as central to their contribution as it is in this paper, they use dominance axioms that, in isolation, do not assume away incompleteness. More radically, one might deny that it is possible to identify individuals across different possible outcomes (i.e., in philosophical parlance, that there are any true "transworld identities"), as Personal Good and all related ex ante Pareto axioms require. In this

[6] We explore a similar question in greater depth in Tarsney et al. (2024), which is a companion paper in the philosophy literature that considers the implications of Personal Good and Negative Dominance in the absence of Expected Multi-Utility and without the characterization results of this paper.

case, any principle in the spirit of Personal Good will be either incoherent (if such cross-outcome identifications are incoherent) or vacuous (if such identifications are coherent but always false).

A third option is to reject the whole project of learning about the evaluative ranking of populations by considering decisions under risk. Arguably the "better than" relation for outcomes or populations is more fundamental than the "better than" relation for lotteries, and one can't properly reason from the latter to the former.[7] In this case, we should conclude that *if* there is antecedent reason to think that some outcomes are incomparable, *then* either Negative Dominance, Personal Good, or one of our other assumptions about the ranking of lotteries must be false, even if we don't know which.

Ultimately, we see our results and the recent literature it joins as bolstering a version of Harsanyi's project of justifying additively-separable social welfare evaluation on the basis of decision-making under risk.[8] Even without addressing contentious issues about the shape of f, the measurement of u_i, or the possibilities for c, the results of this recent literature have substantive implications, such as:

- Non-separable approaches such as average utilitarianism are rejected. This is relevant to efforts in population ethics to avoid the "repugnant conclusion" (Zuber et al., 2021), because the pivotal "mere addition" axiom is closely related to existence independence/separability (see Ng, 1989; Blackorby et al., 2005).

- If adding any life at wellbeing $u^\star$ ever leaves any population at least as good, then adding any other life at any wellbeing greater than $u^\star$ always makes any population better. This is true however large the population, so variable-value population ethics is rejected.

- Risk and population size trade off against one another in an expectation-taking way. So do average wellbeing and population size.

Each of these conclusions would hold for Critical-Band Generalized Utilitarianism, even without Negative Dominance. And even without precise measurement of wellbeing,

[7]As an analogy: It would seem odd to reason from the claims that (1) "two sets are equinumerous iff there's a bijection between them" and (2) "every set is more numerous than its proper subsets" to the claim that "there are no infinite sets", because *what mathematical objects exist* seems like a prior question to *how to compare the sizes of mathematical objects*.

[8]Here we intentionally weaken Harsanyi's claim of justifying "utilitarianism." Sen and Weymark influentially criticized Harsanyi's interpretation of von Neumann–Morgenstern utilities (f in our notation) as utilitarian wellbeing, but see Greaves (2017). Additive separability is important in population ethics and in social choice under risk whether f is linear, concave, or something else.

sometimes these conclusions will have clear implications in practically-relevant cases (because they are conditional on other judgments). Collectively, the results of this literature are narrowing the alternative paths and highlighting that such alternatives have sharper theoretical costs than have been previously recognized.

Acknowledgments

We are grateful for comments from seminar participants at the Global Priorities Institute at Oxford University. Although this research received no specific funding, our collaborative research is supported by the NIH (P2CHD042849) and the Musk Foundation and Spears' population research is supported by the NIH (K01HD098313).

References

Asheim, Geir B and Stéphane Zuber. 2014. "Escaping the repugnant conclusion: Rank-discounted utilitarianism with variable population." *Theoretical Economics*, 9(3): 629–650.

Bader, Ralf M. 2018. "Stochastic dominance and opaque sweetening." *Australasian Journal of Philosophy*, 96(3): 498–507.

Bader, Ralf M.. 2022. "Person-Affecting Utilitarianism." in Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns eds. *The Oxford Handbook of Population Ethics*, New York: Oxford University Press: 251–270.

Blackorby, Charles and David Donaldson. 1984. "Social Criteria for Evaluating Population Change." *Journal of Public Economics*, 25(1): 13–33.

Blackorby, Charles, Walter Bossert, and David Donaldson. 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, Cambridge: Cambridge University Press.

Borie, Dino. 2023. "Expected utility in Savage's framework without the completeness axiom." *Economic Theory*, 76(2): 525–550.

Bossert, Walter, Susumu Cato, and Kohei Kamaga. 2023. "Revisiting variable-value population principles." *Economics & Philosophy*, 39(3): 468–484.

Council of Economic Advisors. 2024. *Economic Report of the President*: Executive Office of the President.

Cowen, Tyler. 2018. *Stubborn attachments: A vision for a society of free, prosperous, and responsible individuals*: Stripe Press.

Danan, Eric, Thibault Gajdos, and Jean-Marc Tallon. 2015. "Harsanyi's aggregation theorem with incomplete preferences." *American Economic Journal: Microeconomics*, 7(1): 61–69.

Dasgupta, Partha. 1995. "The population problem: theory and evidence." *Journal of economic literature*, 33(4): 1879–1902.

Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt. 2023. "The economics of fertility: A new era." in *Handbook of the Economics of the Family*, 1: Elsevier: 151–254.

Dubra, Juan, Fabio Maccheroni, and Efe A Ok. 2004. "Expected utility theory without the completeness axiom." *Journal of Economic Theory*, 115(1): 118–133.

Evren, Özgür. 2008. "On the Existence of Expected Multi-Utility Representations." *Economic Theory*, 35(3): 575–592.

Evren, Özgür and Efe Ok. 2011. "On the multi-utility representation of preference relations." *Journal of Mathematical Economics*, 47(4-5): 554–563.

Fleurbaey, Marc. 2009. "Two variants of Harsanyi's aggregation theorem." *Economics Letters*, 105(3): 300–302.

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

Galaabaatar, Tsogbadral and Edi Karni. 2012. "Expected multi-utility representations." *Mathematical Social Sciences*, 64(3): 242–246.

——— 2013. "Subjective expected utility with incomplete preferences." *Econometrica*, 81(1): 255–284.

Golosov, Mikhail, Larry E Jones, and Michele Tertilt. 2007. "Efficiency with endogenous population growth." *Econometrica*, 75(4): 1039–1071.

Gorno, Leandro. 2017. "A strict expected multi-utility theorem." *Journal of Mathematical Economics*, 71: 92–95.

Greaves, Hilary. 2017. "A reconsideration of the Harsanyi–Sen–Weymark debate on utilitarianism." *Utilitas*, 29(2): 175–213.

Gustafsson, Johan E.. 2020. "Population Axiology and the Possibility of a Fourth Category of Absolute Value." *Economics and Philosophy*, 36(1): 81–110.

Gustafsson, Johan E. and Dean Spears. 2022. "Ex-Post Average Utilitarianism Can Be Worse for All Affected." *PWI Working Paper, UT-Austin*.

Gustafsson, Johan E., Dean Spears, and Stéphane Zuber. 2023. "Utilitarianism is Implied by Social and Individual Dominance." *IZA Discussion Paper*.

Hara, Kazuhiro, Efe A Ok, and Gil Riella. 2019. "Coalitional Expected Multi-Utility Theory." *Econometrica*, 87(3): 933–980.

Hare, Caspar. 2010. "Take the sugar." *Analysis*, 70(2): 237–247.

——— 2013. *The limits of kindness*: Oxford University Press.

Harsanyi, John C.. 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *Journal of Political Economy*, 63(4): 309–321.

Hurka, Thomas. 1983. "Value and population size." *Ethics*, 93(3): 496–507.

Jayachandran, Seema. 2017. "Fertility decline and missing women." *American Economic Journal: Applied Economics*, 9(1): 118–139.

Jones, Charles I.. 2022. "The End of Economic Growth? Unintended Consequences of a Declining Population." *American Economic Review*, 112(11): 3489–3527, November.

Kearney, Melissa S., Phillip B. Levine, and Luke Pardue. 2022. "The Puzzle of Falling US Birth Rates since the Great Recession." *Journal of Economic Perspectives*, 36(1): 151–76, February.

Lederman, Harvey. 2023. "Incompleteness, Independence, and Negative Dominance." *arXiv preprint arXiv:2311.08471*.

Li, Chen, KI Rohde, and Peter P Wakker. 2023. "The deceptive beauty of monotonicity, and the million-dollar question: row-first or column-first aggregation." *Erasmus School of Economics working paper*.

Manzini, Paola and Marco Mariotti. 2008. "On the representation of incomplete preferences over risky alternatives." *Theory and Decision*, 65: 303–323.

McCarthy, David. 2017. "The priority view." *Economics & Philosophy*, 33(2): 215–257.

McCarthy, David, Kalle Mikkola, and Joaquin Teruji Thomas. 2020. "Utilitarianism with and Without Expected Utility." *Journal of Mathematical Economics*, 87: 77–113.

McCarthy, David, Kalle Mikkola, and Teruji Thomas. 2021. "Expected utility theory on mixture spaces without the completeness axiom." *Journal of Mathematical Economics*, 97: 102538.

Nau, Robert. 2006. "The shape of incomplete preferences." *The Annals of Statistics*, 34(5): 2430 – 2448.

Ng, Yew-Kwang. 1989. "What Should We Do About Future Generations?: Impossibility of Parfit's Theory X." *Economics & Philosophy*, 5(2): 235–253.

Ok, Efe A and Nik Weaver. 2023. "Lipschitz Bernoulli utility functions." *Mathematics of Operations Research*, 48(2): 728–747.

Ok, Efe A, Pietro Ortoleva, and Gil Riella. 2012. "Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes." *Econometrica*, 80(4): 1791–1808.

Parfit, Derek. 1986. *Reasons and Persons*: OUP Oxford.

Pivato, Marcus. 2020. "Rank-additive population ethics." *Economic Theory*, 69(4): 861–918.

Riella, Gil. 2015. "On the representation of incomplete preferences under uncertainty with indecisiveness in tastes and beliefs." *Economic Theory*, 58(3): 571–600.

Scovronick, Noah, Mark B Budolfson, Francis Dennig, Marc Fleurbaey, Asher Siebert, Robert H Socolow, Dean Spears, and Fabian Wagner. 2017. "Impact of population growth and population ethics on climate change mitigation policy." *Proceedings of the National Academy of Sciences*, 114(46): 12338–12343.

Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane. 1995. "A Representation of Partially Ordered Preferences." *Annals of Statistics*, 23: 2168–2217.

Shapley, Lloyd S. and Manel Baucells. 1986. "Multiperson Utility." Tech Rep. 779 UCLA Department of Economics.

Spears, Dean and H. Orri Stefánsson. 2021. "Additively-separable and rank-discounted variable-population social welfare functions: A characterization." *Economics letters*, 203: 109870.

Spears, Dean and H. Orri Stefánsson. 2024. "What calibrating variable-value population ethics suggests." *Economics & Philosophy*.

Spears, Dean and Stéphane Zuber. 2022. "Foundations of utilitarianism under risk and variable population." *Social Choice and Welfare*.

Spears, Dean, Sangita Vyas, Gage Weston, and Michael Geruso. 2024. "Long-term population projections: Scenarios of low or rebounding fertility." *Plos one*, 19(4): e0298190.

Sunstein, Cass R.. 2024. "The Economic Constitution of the United States." *Journal of Economic Perspectives*.

Tarsney, Christian, Harvey Lederman, and Dean Spears. 2024. "Share the Sugar."

Thomas, Teruji. 2022. "Separability and Population Ethics." in Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns eds. *The Oxford Handbook of Population Ethics*, New York: Oxford University Press: 271–295.

Zuber, Stéphane, Nikhil Venkatesh, Torbjörn Tännsjö, Christian Tarsney, H Orri Stefánsson, Katie Steele, Dean Spears, Jeff Sebo, Marcus Pivato, Toby Ord et al.. 2021. "What should we agree on about the repugnant conclusion?" *Utilitas*, 33(4): 379–383.