Expectational Total Utilitarianism Is Implied by Social and Individual Dominance

Johan E. Gustafsson,* Dean Spears,† and Stéphane Zuber‡

We provide a new axiomatization of expectational total utilitarianism, which is a core framework for evaluating policies and social welfare under variable population and social risk. Our innovation is a previously unrecognized combination of weak assumptions that yields expectational utilitarianism, in a fixed-population result, and expectational total utilitarianism, in a variable-population result. We show that two dimensions of weak dominance (over risk and individuals) characterize a social welfare function with two dimensions of additive separability. So social expected utility emerges merely from social statewise dominance (given other axioms). Moreover, additive utilitarianism arises merely from individual stochastic dominance, which is assumed only for lives certain to exist (so this axiom does not compare life against non-existence), without assuming individual preference orderings. A variable-population setting, we show, allows us to eliminate an assumption that a fixed-population setting requires. Our result provides an important foundation for evaluating climate change, growth, and global depopulation.

Keywords: Social risk, variable population, utilitarianism.

JEL Classification numbers: D63, D81.

*University of Texas at Austin, USA. University of York, UK. University of Gothenburg, Sweden. IFFS, Sweden. E-mail: *johan.eric.gustafsson@gmail.com*.

†University of Texas at Austin, USA. IZA, Germany. IFFS, Sweden. GPI, Oxford University, UK. r.i.c.e., India. E-mail: *dspears@utexas.edu*.

‡Paris School of Economics (CNRS) and Centre d'Économie de la Sorbonne, France. E-mail: *Stephane.Zuber@univ-paris1.fr*.

1. Introduction

Policies that change future mortality rates (like climate mitigation) or change future fertility rates (like public education) not only change the quality of lives in the future but also who will live in the future. Our global population has quadrupled in size over the past hundred years and is projected to peak then shrink within the lifetime of children born today, with uncertain consequences.[1] Hence evaluating economic policies requires assessing both social risk and variable population. A standard principle for economic policy evaluation is Expectational Total Utilitarianism, which maximizes the expected value of the sum of individuals' lifetime well-being. Despite the prominent use in public economics of both additive utilitarianism and expectation-taking under risk, these methods continue to be questioned in welfare economics, in part because existing axiomatic justifications make seemingly strong assumptions (Fleurbaey, 2010; Golosov et al., 2007).

We provide a new axiomatic path to Generalized Expectational Total Utilitarianism, which is a set of social welfare functions, introduced by Blackorby, Bossert, and Donaldson (1998), of which Expectational Total Utilitarianism is the most straightforward case. Our result builds on a new combination of weak assumptions that yields, first, in a fixed-population theorem, Generalized Expectational Fixed-Population Utilitarianism, and second, in a variable-population theorem, Generalized Expectational Total Utilitarianism.[2] The variable-population setting, we show, allows us to eliminate a contentious assumption that the fixed-population setting requires.

In particular, our result extends Harsanyi's aggregation theorem to variable populations but does not assume Expected Utility Theory either for society or for individuals. Instead, we show that two dimensions of weak dominance (one over risky states and the other over individuals) characterize a social welfare function with two dimensions of additive separability. So social expected utility emerges merely from social

[1] For two recent overviews of the economists' emerging understanding and open questions about low fertility, see Kearney et al. (2022) and Doepke et al. (2022). For possible macroeconomic implications of depopulation, see Jones (2022).

[2] An active current theoretical literature in welfare economics has been exploring additive separability over two dimensions, including Fleurbaey (2009), Mongin and Pivato (2015), McCarthy et al. (2020), and Spears and Zuber (2022). We discuss related prior literature in more depth in our conclusion section.

Table 1: Motivating example: The only affected person is sure to exist in both prospects and is stochastically better-off in g^*

states:	prospect f^* Ann	prospect f^* Bob	prospect g^* Ann	prospect g^* Bob
s_1	1	1	1	9
s_2	Ω	7	Ω	1

statewise dominance (in the context of our other axioms). Moreover, additively-separable utilitarianism arises merely from individual stochastic dominance, which is only assumed for lives certain to exist; this is crucial because it means that our core variable-population axiom does not compare life against non-existence. Further, without assuming complete individual preferences, we derive that the social order respects individual-level expected utility.

We introduce a new axiom, Stochastic Dominance for Sure Individuals, which is at the core of our new results. To understand the axiom, consider the example in Table 1, where columns are individuals, rows are equiprobable states, and for the purposes of this example, Ω represents an individual's non-existence in a state.[3]

The principle behind our result only makes comparisons for individuals who are sure to exist and, even for those individuals, only uses stochastic dominance. In this comparison of f^* and g^*, Ann is not sure to exist in either prospect. But she is altogether unaffected by the social choice of f^* or g^*: her utility conditional-on-existence does not differ between her prospects, nor her probability of existence, nor even the state in which she exists. Bob is the only person sure to exist. Given the equal probability of the two states, prospect g^* stochastically dominates prospect f^* for him. Our principle says that, in this situation, g^* is better than f^*. An egalitarian, however, might say that f^* is better than g^*, if stochastically "levelling down" Bob is a price worth paying to ensure *ex-post* equality

[3] This example was first introduced in the philosophy literature, as a counterexample against Expectational Average Utilitarianism, by Gustafsson and Spears (2022); their informal paper does not include any characterization results. Gustafsson and Spears emphasize that their counterexample uses only *positive* lifetime utilities, so unlike other classic objections to Average Utilitarianism and to other non-separable approaches to population ethics, their counter-example does not depend upon a meaningful zero for utility nor upon the plausible existence of lives not worth living.

(Myerson 1981, Parfit n.d., ch. 1; 1995, p. 17).

Our basic principle says that a prospect is better if it can be shown, without comparing existence to non-existence, to be better for somebody sure to exist and worse for nobody, where "better for" merely means in the sense of stochastic dominance, which is an incomplete ranking. And, in fact, our principle weakens stochastic dominance to only apply if states can be permuted the same way for all sure-to-exist individuals. Our new principle is consistent with egalitarian (or any other) choice in Myerson's counter-example to utilitarianism (which we introduce in Section 3 and discuss further with our principle in Section 4).

But notice that Expectational Average Utilitarianism, Expectational Maximin, and Expectational Rank-Discounted Utilitarianism[4] would each choose f^* over g^* in Table 1. And so would a non-expectation-taking rule that evaluates outcomes according to Total Utilitarianism but then uses maximin for social risk, choosing the prospect with the highest least-valuable state.[5] Our principle is violated by each of these alternative approaches to variable-population social choice under risk. So our principle offers a new foundation for Generalized Expectational Total Utilitarianism, grounded in a simple axiom that is in the spirit of respecting Pareto improvements.

Our paper begins, after introducing our framework, with a theorem for a fixed-population setting. This theorem requires an axiom (that we call Compensation) that can be dropped in our second theorem, which is in a variable-population setting. In this way, we apply the insight of Blackorby and Donaldson (1984), who showed that variable population provides a second axiomatic path to additive utilitarianism. More broadly, we demonstrate what can be achieved by combining Blackorby and Donaldson's approach with Harsanyi's result, which is a standard basis of utilitarian welfare economics. Our result shows that Harsanyi's assumptions can be considerably weakened in a variable-population setting — which

[4] Expectational Average Utilitarianism maximizes the expectation of the average of lifetime utility among people born. Expectational Maximin maximizes the expectation of the utility of the worst-off person born. See Blackorby et al. (2005) for an overview of alternative approaches to population ethics. Expectational Rank-Discounted Utilitarianism would maximize the expectation of $\sum_r \beta^r u_{[r]}$, where $u_{[r]}$ orders individuals by rank r from worst-off; see Asheim and Zuber (2016).

[5] That is, $\min_s \left(\sum_i u_{is}\right)$. There is literature in decision theory proposing to maximize the minimal expected utility using a set of probability, in particular Gilboa and Schmeidler (1989).

is also the setting that is pragmatically needed for economic policy assessments.

After presenting our main results, we show that they can be reinterpreted to apply to further economic settings and questions with two dimensions of value. Macroeconomists, for one example, canonically use time-separable social objective functions that add a value for each time period, which in turn is found by adding a period utility for each individual: $\sum_t \alpha_t \sum_i u_{it}$ for individuals i and times t. Although our main result considers risky states and individuals, our result can also be applied to time periods and individuals. So our results offer a new axiomatic justification for this practice in the macroeconomic literature. For example, we provide a microfoundation for Klenow, et al.'s (2022) recent growth accounting model that incorporates population growth while assuming a Total Utilitarian perspective: They conclude that, even though per-person living standards have improved radically, over the past decades population growth has accounted for even more of the vast improvements in aggregate well-being. Climate economics, such as Nordhaus' (2017) DICE and RICE models, also uses this functional form for optimizing macroeconomic climate policy. We provide a new microfoundation for this standard tool in normative macroeconomics.

In another application, we show that our formal result can be reinterpreted to propose that a prudent individual making risky decisions about a life of unknown length and per-period utility should maximize the expected sum of per-period utility over time, $\sum_s \alpha_s \sum_t u_{st}$ for risky states s and time periods t. This application would exclude, for example, risk aversion over the length of the person's life.

2. Framework

Let $\mathbb{N}$ denote the set of positive integers, 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 I. We will consider two cases: a fixed-population case where $I = \{1, \cdots, n\}$ for some finite n and individuals always exist; and a variable-population case where $I = \mathbb{N}$ and only a finite subset of individuals exist in any realized outcome. In the variable-population case, since $I = \mathbb{N}$, it holds that N

Don't need to be utilities, just ordinally ranked

is also the set of all possible realized populations of individuals, meaning the set of all subsets of I. That is, in any outcome, N is the population that exists. So $N \in \mathcal{N}$ in the variable-population case and $N = I$ in the fixed-population case.

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 $u = (u_i)_{i \in N} \in \mathbb{R}^N$, where N is the population, and $u_i \in \mathbb{R}$ is, for each existing individual i, the lifetime utility experienced by i. We denote U the set of outcomes; the exact definition will be different in the fixed-population and variable-population cases.

For each u, we denote $N(u)$ the set of individuals alive in u, and $n(u)$ the number of individuals alive in u. For two any outcomes u and v such that $N(u) \cap N(v) = \varnothing$ (that is, u and v are distributions of utility for two disjoint populations), we denote uv the outcome where the two populations are merged. Formally, it is the outcome w such that $N(w) = N(u) \cup N(v)$, $w_i = u_i$ for all $i \in N(u)$, and $w_j = u_j$ for all $j \in N(v)$.

We assume that it is not always known for sure what the final utility vector will be nor what set of individuals will exist. For now, we assume that there is a fixed, finite set of states of the world $S = \{1, \cdots, m\}$, with $m \geq 2$, where all states are equally probable so that each state has probability $\frac{1}{m}$. A "supplementary material" appendix discusses the extension to the more general cases. Note that we are in a framework where probabilities are given and/or individuals have the same beliefs (an "objective" probability framework).

A social prospect f is a function from S to U. For $s \in S$, $f(s)$ is therefore the outcome induced by the prospect f in state s. We let $F = U^S$ be the set of functions from S to U. For an outcome $u \in U$, we abuse notation and also denote u the sure social prospect f such that $f(s) = u$ in all $s \in S$.

For any outcome $u \in U$, whenever $i \in N(u)$, $u_i \in \mathbb{R}$ denotes the utility of individual i. For a prospect $f \in F$, whenever $i \in N(f(s))$, $f_i(s)$ denotes the utility of individual i in state of the world $s \in S$. For an individual $i \in I$ and a social prospect $f \in F$, we let $S_i(f) = \{s \in S | i \in N(f(s))\}$ be the set of states of the world where individual i exists. In the fixed-population case, we have $S_i(f) = S$ for all individuals and prospects. But this may not be case in the variable-population case: In that case a potential individual may not exist in some (or any) state of the world.

Consider an individual $i \in I$ and a social prospect $f \in F$ such that

$S_i(f) = S$, meaning that the individual exists for sure in all states of the world. Given objective probabilities of states, we can define the individual prospect for i that is associated with f. We write such an individual prospect using the associated cumulative distribution function $q_i^f : \mathbb{R} \to [0,1]$, with $q_i^f(z) = \frac{|\{s \in S | f_i(s) \leq z\}|}{m}$; $q_i^f(z)$ is the probability that individual i obtains a level of wellbeing level that is at most z. Notice that probabilities are formed as fractions of the set of equiprobable states.

The task of our paper is to characterize a social preorder $\succsim$ on F. That $\succsim$ is a preorder means that it is a reflexive and transitive binary relation. In particular, the preorder $\succsim$ is not directly assumed by our axioms to be complete on F, although both of our theorems derive completeness on F from the combination of our axioms. Throughout the paper we assume completeness only on sure prospects, as stated in our first axiom:

> *Completeness for Sure Prospects* For all $u, v \in U$, either $u \succsim v$, or $v \succsim u$, or both.

Completeness for Sure Prospects is not contentious within the literature for same-population cases. It is more contentious in the philosophical population-ethics literature. Completeness, in that case, would hold that populations with different sizes are always comparable. Some authors argue that variable-population completeness may not hold because we do not know the critical-level.[6] But approaches with incompleteness are typically subject to time-consistency problems or money pump arguments (Hammond, 1988; Gustafsson, 2022). Variable-population incompleteness would also have deeply unattractive practical and normative implications, such as that climate mitigation policy is not preferable to large global temperature increases, because different sets of people would exist.[7] Ordinary economic analysis and policy-making routinely (if implicitly) assume that outcomes with different populations can be compared; we follow that tradition.

3. Fixed-population results

We assume in this section that the population is a fixed set of n individuals, $I = \{1, \cdots, n\}$ with $n \geq 3$. The set U of outcomes is $U = \mathbb{R}^I$. In a slight

[6] See the literature on range critical level: Blackorby et al. 1996, Rabinowicz 2009, and Gustafsson 2020.

[7] This observation is an application of Parfit's (1984, p. 362) *Depletion case* in the philosophical population-ethics literature.

abuse of notation, it will sometimes be useful to consider subsets of I, which we will call N in this section, and to consider the utility distribution of the subpopulation within N that is an element of $\mathbb{R}^N$.

Our first results are based on two dominance principles, one for society and one for individuals. In our social dominance principle, the notation $f(s) \succsim g(s)$ means "if $\succsim$ faced a binary choice between the outcome of f in s occurring for sure (that is, in every state), or the outcome of g in s occurring for sure, then $\succsim$ would prefer the former to the latter."

> *Social Statewise Dominance* For all $f, g \in F$, if $f(s) \succsim g(s)$ for all $s \in S$, then $f \succsim g$. If in addition there exists $s \in S$ such that $f(s) \succ g(s)$, then $f \succ g$.

Social Statewise Dominance is a very weak rationality principle for social decision making. It means that if we are sure that a social prospect would be better than another under any state, then we should prefer it.

Our individual dominance principle uses our notation for individual prospects. For two social prospects $f, g \in F$ and an individual i such that $S_i(f) = S_i(g) = S$, q_i^f stochastically dominates q_i^g if and only if $q_i^f(z) \leq q_i^g(z)$ for all $z \in \mathbb{R}$ and $q_i^f(z') < q_i^g(z')$ for some $z' \in \mathbb{R}$. If $q_i^f = q_i^g$, then q_i^f and q_i^g correspond to the same individual prospect.

> *Individual Stochastic Dominance* For all $f, g \in F$, if $q_i^f(z) \leq q_i^g(z)$ for all $z \in \mathbb{R}$ and $i \in I$ then $f \succsim g$. If in addition there exists $j \in I$ and $z' \in \mathbb{R}$ such that $q_j^f(z') < q_j^g(z')$, then $f \succ g$.

Individual Stochastic Dominance can be interpreted as a weak *ex-ante* Pareto principle: If a prospect is better than another for all individuals (in the sense of first-order stochastic dominance), then it is also socially better. In that sense, it is in the lineage of Harsanyi's foundational result on social aggregation under risk (Harsanyi, 1955). Note, however, that Individual Stochastic Dominance is weaker than the usual *ex-ante* principles for two reasons: because it is compatible with non-expected utility assessments of individual prospects, and because it is only an incomplete ranking of individual prospects. An interpretation is that the social ranking needs not always respect individual preferences, but instead only respects the part of individual preferences, if any, which is compatible with first-order stochastic dominance.

Table 2: Myerson's objection to utilitarianism: Equal indivdiual prospects can yield unequal outcomes

	prospect f^{**}		prospect g^{**}	
states:	Ann	Bob	Ann	Bob
s_1	1	1	1	0
s_2	0	0	0	1

We recognize that the Individual Stochastic Dominance principle itself has important and controversial normative implications for *ex-post* fairness. To show this, let us reframe our introductory example into the fixed-population example in Table 2, which is originally due to Myerson (1981).[8] Continue to assume that the two states are equally probable. So both individuals face the same individual prospect in f^{**} and g^{**}. Accordingly, Individual Stochastic Dominance implies that those two prospects must be indifferent. But several authors have argued that society may prefer prospect f^{**} because it does not imply any inequality *ex post* (Fleurbaey, 2010). We acknowledge this implication for our fixed-population result, but note that our variable-population individual dominance axiom, introduced below, avoids it.

The conflict with the intuition behind Table 2 emerges, in fact, from the Anteriority axiom, which is weaker than Individual Stochastic Dominance and which says that the social preorder only depends on which prospect each individual faces, that is:

> *Anteriority* If $q_i^f(z) = q_i^g(z)$ for all $z \in \mathbb{R}$ and all $i \in I$ then $f \sim g$.

McCarthy et al. (2020) have argued that Anteriority expresses a weak sense in which the social preorder is *ex ante*. So, our characterization results can be seen as attractive to people endorsing a weak *ex-ante* view, or as additional arguments for people who resist that view.

Our first result is to show that those two dominance principles, together with Completeness for Sure Prospects, imply the following separability property for sure prospects:

> *Separability for Sure Prospects* For any $N \subset I$, for any $u, v \in \mathbb{R}^N$ and $w, \hat{w} \in \mathbb{R}^{I \setminus N}$, $uw \succsim vw$ if and only if $u\hat{w} \succsim v\hat{w}$.

[8] See also Broome 1991, p. 185.

Lemma 1. *If the social ordering $\succsim$ satisfies Completeness for Sure Prospects, Social Statewise Dominance and Individual Stochastic Dominance, then it satisfies Separability for Sure Prospects.*

Proof. The proof is by contradiction. Assume that $N \subset I$, $u, v \in \mathbb{R}^N$ and $w, \hat{w} \in \mathbb{R}^{I \setminus N}$ are such that $uw \succsim vw$ but $v\hat{w} \succ u\hat{w}$. Consider the three following prospects f, g, and h (where each row gives the vector of utilities in a specific state of the world):

state:	f	g	h
1	uw	vw	uw
2	$u\hat{w}$	$u\hat{w}$	$v\hat{w}$
3	uw	uw	uw
…	…	…	…
m	uw	uw	uw

By Social Statewise Dominance, given that $uw \succsim vw$, we must have $f \succsim g$. By Social Statewise Dominance, given that $v\hat{w} \succ u\hat{w}$, we must have $h \succ f$. Hence, by transitivity, we should have $h \succ g$. But it is the case that $q_i^g(z) = q_i^h(z)$ for all $z \in \mathbb{R}$ and $i \in I$, so that Individual Stochastic Dominance requires $g \sim h$, a contradiction. Completeness for Sure Prospects implies that, if we do not have $v\hat{w} \succ u\hat{w}$, we must have $u\hat{w} \succsim v\hat{w}$. ∎

Note that for this first result, we do not need the full force of Individual Stochastic Dominance, but only Anteriority.

Lemma 1 is already a big step towards additive separability, because we now have a strong separability condition. But to obtain our fixed-population main result, we need two additional technical properties.

> *Continuity* For all $u \in U$, the sets $\{v \in U | u \succsim v\}$ and $\{v \in U | v \succsim u\}$ are closed.

> *Compensation* For any $u \in U$ and $i \in I$, there exists $z \in \mathbb{R}$ such that, if $v \in U$ is defined by $v_i = z$ and $v_j = 0$ for all $j \neq i$, then $u \sim v$.

Compensation means that we can compensate losses or gains (from 0) of all but one individuals by adjusting the welfare level of the last individual. Such a property is sometimes named Solvability in the literature (see for instance Pivato and Tchouante, 2022). Although Compensation may intuitively appear utilitarian, it is consistent with views that are sensitive to distribution, such as equally-distributed-equivalent egalitarianism $\left(\phi^{-1}\left(\frac{1}{n}\sum_i \phi(u_i)\right)\right)$ and rank-discounted generalized utilitarianism $\left(\sum_{[r]} \beta^r \phi(u_r)\right.$, where $[r]$ indicates rank from worst-off$\left.\right)$, if ϕ is an unbounded positive transformation.

Theorem 1. *If the social ordering $\succsim$ satisfies Completeness for Sure Prospects, Social Statewise Dominance, Individual Stochastic Dominance, Continuity and Compensation then there exist continuous and increasing functions $\phi_i : \mathbb{R} \to \mathbb{R}$ such that:*

$$f \succsim g \iff \sum_{s\in S} \frac{1}{m} \sum_{i\in I} \phi_i\big(f_i(s)\big) \geq \sum_{s\in S} \frac{1}{m} \sum_{i\in I} \phi_i\big(g_i(s)\big).$$

An impartiality axiom would replace ϕ_i with a shared ϕ.

The proof is in the Appendix. It has two steps that we describe here informally. The first step is to derive additive separability within a state (or for sure prospects). It relies on Separability for Sure Prospects, using Lemma 1, and on the theorem by Debreu (1960) on additive representations. The second step is to construct the across-state additivity of social expected utility. Informally, this is done by combining the use of the additive formula within a state and Stochastic Dominance for Sure Individuals to move the consequences of other states all into one state. This is illustrated by the following two-by-two example (which disregards ϕ for illustration), where columns are individuals, rows are two equiprobable risky states (s_1 and s_2), and x, y, z, and w are real lifetime utilities:

$$\begin{array}{c} s_1 \\ s_2 \end{array} \begin{bmatrix} x & y \\ w & z \end{bmatrix} \sim \begin{bmatrix} x+y & 0 \\ 0 & w+z \end{bmatrix} \sim \begin{bmatrix} x+y & w+z \\ 0 & 0 \end{bmatrix} \sim \begin{bmatrix} x+y+w+z & 0 \\ 0 & 0 \end{bmatrix}.$$

The first equivalence uses Social Statewise Dominance and the additive structure within states. The second equivalence uses Stochastic Dominance for Sure Individuals. The third equivalence again uses Social Statewise Dominance and the additive structure within states.

Two remarks can be made on Theorem 1. First, we do not have a full equivalence result. This is because, to satisfy Compensation, the ϕ_i functions have to satisfy conditions that are hard to write but not very interesting. It is clear that if $\succsim$ is represented by the expected value of a generalized utilitarian function then it also satisfies all the principles except Compensation. Second, we can derive the result with even weaker principles. As explained before, Lemma 1 only requires Anteriority. Similarly, our full proof only requires Anteriority and a Pareto-like property on alternatives that is implied by Individual Stochastic Dominance (this is detailed in the proof). So Theorem 1 can be reformulated using Anteriority and this Pareto-like property in place of Individual Stochastic Dominance. However, we present Individual Stochastic Dominance because it foreshadows our variable-population theorem.

4. Variable-population results

Theorem 1 is a powerful weakening of the Harsanyi approach. But fixed-population utilitarianism leaves open the question of how to expand to variable-population questions — which are the real-world questions of much actual economic policy decision-making. Blackorby et al. (2005) detail many variable-population social welfare functions (such as Average Utilitarianism or Number-Dampened Utilitarianism) that simplify to fixed-population utilitarianism in fixed-population cases. This section shows how social and individual dominance further narrow down the possibilities for utilitarianism in a variable-population setting. We show that our axioms imply a specific family of generalized utilitarianisms for variable-population cases, namely, Generalized Expectational Total Utilitarianism.

Additionally, we can take advantage of the variable-population setting — which has its own "existence independence" route to additive separability, due to Blackorby and Donaldson (1984) — to weaken our assumptions. In particular, our individual dominance axiom in the variable-population setting is consistent with making the egalitarian choice that $f^{**} \succ g^{**}$ in Table 2. Moreover, in the fixed-population case, we use the Compensation principle, but this may not be obviously appealing from the ethical viewpoint. By moving to the variable-population case, we will be able to instead use the principles of Anonymity and Weak critical level, which are widely accepted in the variable-population welfare economics literature.

In this section, the set of *potential* individuals who may or may not exist is $I = \mathbb{N}$. In an alternative, only a non empty finite population $N \in \mathcal{N}$ exists. We thus define $U = \bigcup_{N \in \mathcal{N}} \mathbb{R}^N$ as the set of possible outcomes when at least one individual exists.

This section requires additional notation to refer to people who may or may not exist. For each $u \in U$, we denote $N(u)$ the set of individuals alive in u, and $n(u)$ the number of individuals alive in u. For each population $N \in \mathcal{N}$, $U_N = \{u \in U | N(u) = N\}$ is the set of outcomes such that the population is N.

We adopt six principles that are properties of the social preorder $\succsim$.

Completeness for Sure Prospects and Social Statewise Dominance are the same as in the previous sections. We first have three principles that we expect to be uncontroversial in the economics literature.

> *Anonymity for Sure Outcomes* For all $u, v \in U$, if $n(u) = n(v)$ and there exists a bijection $\pi : N(u) \to N(v)$ such that $u_i = v_{\pi(i)}$ for all $i \in N(u)$, then $u \sim v$.
>
> *Same-Population Continuity for Sure Outcomes* For all $N \in \mathcal{N}$, for all $u \in U_N$, the sets $\{v \in U_N | u \succsim v\}$ and $\{v \in U_N | v \succsim u\}$ are closed.
>
> *Minimal Critical Level* There exists $z \in \mathbb{R}_+$, $u \in U$, and $j \in (I \setminus N(u))$ such that, if v is defined by $N(v) = N(u) \cup \{j\}$, $v_i = u_i$ for all $i \in N(u)$ and $v_j = z$, then $u \sim v$.

Notice that Anonymity and Same-population continuity only apply to comparisons among sure outcomes. Minimal Critical Level only asserts the existence of one instance of variable-population comparability. Minimal Critical Level assumes a weak comparability between existence and non-existence, but we have already assumed this by assuming a complete social ordering of sure variable-population outcomes. In this context, minimal critical level merely rules out the implausible cases that adding a life is always worse or always better, no matter how good or bad.

The heart of our variable-population characterization is our stochastic dominance axiom for individuals: Stochastic Dominance for Sure Individuals. This axiom formalizes the principle behind our motivating example in Table 1. To adapt Individual Stochastic Dominance to the variable-population setting, we apply the principle only to individuals who are sure to exist — like Bob is in Table 1's motivating example. Additionally, this axiom only applies when states of the world where not-sure-to-exist individuals exist and their utilities conditional on existence are left unchanged.[9]

[9] We thank Marcus Pivato for suggesting this formulation of the Stochastic Dominance for Sure Individuals principle.

Stochastic Dominance for Sure Individuals For all $f, g \in F$, if:

(i) $S_i(f) = S_i(g)$ for all $i \in I$;

(ii) for all $j \in I$ such that $S_j(f) \notin \{\emptyset, S\}$, there exists $x_j \in \mathbb{R}$ such that $f_j(s) = g_j(s) = x_j$ for all $s \in S_j(f)$;

(iii) there exists a bijection $\sigma : S \to S$ such that for all $k \in I$ such that $S_k(f) = S$ and all $s \in S$, $f_k\big(\sigma(s)\big) \geq g_k(s)$;

then $f \succsim g$.
If in addition there exists $l \in I$ such that $S_l(f) = S$ and $s' \in S$ such that $f_l\big(\sigma(s')\big) > g_l(s')$ then $f \succ g$.

This axiom, Stochastic Dominance for Sure Individuals, has three important features:

- In condition (i), individuals exist in the same states of the world in the two social prospects f and g being compared, which equivalently means that in each state of the world the populations existing with f and g are the same. The principle does not speak to situations with different populations in a state of the world.

- In condition (ii), people who do not exist for sure either do not exist at all, or they do not bear any risk and exist with the same level of utility in f and g. People who do not necessarily exist, in the comparison between f and g, are altogether unaffected.

- For people who are sure to exist, the condition (iii) entails that the individual prospect they face in f stochastically dominates the one they face in g. That is, condition (iii) implies that $q_i^f(z) \leq q_i^g(z)$ for all $z \in \mathbb{R}$ for any individual i exiting for sure. But, in fact, condition (iii) is weaker than individual stochastic dominance because the same permutation of states σ is used for all individuals.

Notice, then, that Stochastic Dominance for Sure Individuals requires that $g^* \succ f^*$ in the example from Table 1 but it permits any ranking of f^{**} and g^{**} in Table 2, including the non-utilitarian judgment that $f^{**} \succ g^{**}$. We cannot conclude from Stochastic Dominance for Sure Individuals that f^{**} and g^{**} in Table 2 are socially equivalent, because to obtain dominance we need to use different permutations of states for Ann and Bob. Yet we can conclude for Table 1 that $g^* \succ f^*$ because only Bob exists for sure so we can permute the outcome for Bob in states 1 and 2. These examples, therefore, distinguish Stochastic Dominance for Sure Individuals from Anteriority, because Anteriority would immediately imply the utilitarian judgement that $f^{**} \sim g^{**}$.[10] Indeed, although the representation in Theorem 2 implies that $f^{**} \sim g^{**}$, no one axiom used in our variable-population theorem individually requires this.

Our first result is that the restricted social ordering of sure prospects must be a critical-level generalized utilitarian social ordering. Fundamentally, we achieve additive separability from our axioms because, in our variable population setting, Social Statewise Dominance and Stochastic Dominance for Sure Individuals are sufficient for Blackorby et al.'s (2005) axiom Existence independence.

Proposition 1. *If $\succsim$ satisfies Completeness for Sure Prospects, Anonymity for Sure Outcomes, Same-Population Continuity for Sure Outcomes, Minimal Critical Level, Social Statewise Dominance, and Stochastic Dominance for Sure Individuals, then there exists a continuous and increasing function* $\phi : \mathbb{R} \to \mathbb{R}$ *and a number* $c \in \mathbb{R}$ *such that for all* $u, v \in U$, $u \succsim v$ *if and only if* $\sum_{i \in N(u)} \left[\phi(u_i) - \phi(c)\right] \geq \sum_{i \in N(v)} \left[\phi(v_i) - \phi(c)\right]$.

Proof. Let us first show that $\succsim$ satisfies the following principle, Existence Independence for Sure Prospects:

Existence Independence for Sure Prospects For all $u, v, w \in U$, $uw \succsim vw$ if and only if $u \succsim v$.

The proof is by contradiction and is similar to that of Lemma 1. It is obtained by considering the three prospects:

[10] Anteriority, as written above, is not defined for variable-population cases. So consider, further, a variable population extension of Anteriority which holds that two prospects are equally good if each potential person faces the same individual distribution of the probability of non existence and the same distribution of utility levels conditional on existence (McCarthy et al., 2020). Such an Anteriority axiom would both hold that Table 1's $g^* \succ f^*$, like Stochastic Dominance for Sure Individuals, and that Table 2's $f^{**} \sim g^{**}$. Such an Anteriority axiom is thus stronger than Stochastic Dominance for Sure Individuals.

	f	g	h
1	uw	vw	uw
2	u	u	v
3	u	u	u
…	…	…	…
m	u	u	u

The next step is to recognize that, with Existence Independence, all requirements of Theorem 6.10 of (Blackorby et al., 2005, p. 191) for Critical-Level Generalized Utilitarianism (henceforth CLGU)[11] are met, so $\succsim$ restricted to U is CLGU – which is the form obtained in the Proposition. ∎

We can then state our main result for this section:

Theorem 2. *The following statements are equivalent:*

(1) The social preorder $\succsim$ satisfies Completeness for Sure Prospects, Anonymity for Sure Outcomes, Same-Population Continuity for Sure Outcomes, Minimal Critical Level, Social Statewise Dominance and Stochastic Dominance for Sure Individuals.

(2) $\succsim$ is a complete social order; there exists a continuous function $\phi : \mathbb{R} \to \mathbb{R}$ and a number $c \in \mathbb{R}$ such that for all $f, g \in F$, $f \succsim g$ if and only if

$$\sum_{s \in S} \tfrac{1}{m} \Big[\sum_{i \in N(f(s))} [\phi(f_i(s)) - \phi(c)] \Big] \geq \sum_{s \in S} \tfrac{1}{m} \Big[\sum_{i \in N(g(s))} [\phi(g_i(s)) - \phi(c)] \Big].$$

The basic approach is to use the within-state additivity of CLGU to construct the across-state additivity of social expected utility. Informally, this is done first by using CLGU to have a separate set of individuals with welfare different from c in each state of the world. Then we use Stochastic Dominance for Sure Individuals to move the consequences of other states all into one state. Then we can apply CLGU to get an additive formula. Consider the following example for intuition of the proof. There are four possible individuals (in columns) and two equiprobable states (in rows);

[11] In the background context of an anonymous social order, "[$\succsim$] satisfies continuity, strong Pareto, weak existence of critical levels, and existence independence if and only if [$\succsim$] is CLGU."

x, y, z, and w are real lifetime utilities:

$$\begin{bmatrix} x & y & \Omega & \Omega \\ \Omega & w & z & \Omega \end{bmatrix} \sim \begin{bmatrix} x & y & c & c \\ \Omega & \Omega & w & z \end{bmatrix} \sim \begin{bmatrix} x & y & w & z \\ \Omega & \Omega & c & c \end{bmatrix}.$$

The first equivalence uses CLGU from Proposition 1 in each state of the world (and then Social Statewise Dominance). The second equivalence uses Stochastic Dominance for Sure Individuals (the last two individuals). We can then use the additive formula of Generalized Expectational Total Utilitarianism applied to the first state of the last prospect. The full proof is presented in the Appendix.

Notice that Stochastic Dominance for Sure Individuals is independent of the other axioms of Theorem 2, because the other axioms are each consistent with Expectational Average Utilitarianism, but Stochastic Dominance for Sure Individuals is not. The next logical weakening of Stochastic Dominance for Sure Individuals would be to weaken shared-permutation stochastic dominance to statewise dominance, but this would not be sufficient for Theorem 2, which suggests that Stochastic Dominance for Sure Individuals may be the weakest axiom that can narrow variable-population utilitarianism to Generalized Expectational Total Utilitarianism.

5. Further applications

In this section, we note that our formal result can be usefully reinterpreted if the dimensions and utility-bearers are understood in different ways.[12] We give an example for macroeconomics and another for individual rational choice. Where our main setting uses risky states and individuals as the two dimensions, our applications below use, first, time periods and individuals and, second, time periods and risky states.

5.1 Macroeconomic welfare accounting with time separability: time periods and individuals

Macroeconomists typically use a social welfare function that is additively separable across time periods and sums individual time-period-specific

[12] Mongin and Pivato (2015) have made a similar observation, in surveying multiple applications of their own result about two-dimensional separability, although their result and applications are different.

utility within time periods. This practice has two important implications: that individual lifetime utility is also additively time-separable, and that the implied population ethics is totalist. For example, the Nobel-winning climate-economy model of Nordhaus (2017) (like other leading climate-policy models) maximizes a social objective function $\sum_t \alpha_t \sum_i u_{it}$, for individuals i and periods t experiencing flow utility u_{it}, or more precisely $\sum_t \alpha_t L_t \bar{u}_t$, where L_t is population size and $\bar{u}_t$ is average wellbeing in t. Particularly relevantly to our paper, Klenow et al. (2022) use this functional form (without time discount factors α_t) to conduct a growth accounting exercise that decomposes aggregate growth into population growth and improvements in per-person living standards.

These conventions invite the question: How can this social objective function be normatively justified? Our Theorem 1 provides a justification, if cells are reinterpreted as individual-by-time flows of utility, risky states are reinterpreted as discrete time periods (ignoring risk for this application), and potential individuals have lives composed of a variable number of time periods.

- Our Social Statewise Dominance axiom would become Social periodwise dominance, holding that an intertemporal allocation f is better than another g if each time period of f would be better, if made permanent, than the corresponding time period of g, if made permanent.
- Our Stochastic Dominance for Sure Individuals would become Temporal dominance for fixed-longevity individuals, holding that an intertemporal allocation f is better than another g if
  - every person who only lives for some (but not all) populated time periods is unaffected by a choice between f and g, and
  - every person who lives throughout the entire span of populated time has a lifelong distribution of period well-being in f that dominates that person's lifelong distribution of well-being in g.

These, combined with the technical axioms, would yield additivity across and within time periods.[13] So this result can justify Klenow, et al.'s (2022) Total Utilitarian growth accounting, with the same sort of weak axioms that justify our result.[14] To be sure, various intuitions (including a taste for "pattern goods" such as flat or increasing utility profiles over time) might lead an economist to reject Temporal dominance for fixed-longevity individuals, but such economists would already have rejected macroeconomists' entire time-separable project. This is formally analogous to, in our original social risk setting, a concern for egalitarianism that might bring about a rejection of utilitarianism and our axioms that characterize it.

5.2 Individual decision-making for a lifetime of risky length and per-period utility: time periods and risky states

Consider an individual's rational choice over a risky temporal distribution of state-specific period flows of utility, u_{st}, where s are risky states and t are time periods when the individual may or may not be alive and, if so, would experience a flow utility. Reinterpreting our model of social risky choice as a model of individual risky choice, with i in our model now becoming periods t in a life of unknown length, results in the decision criterion that maximizes the expectation of the sum of period-specific utility flows over a lifetime: $\sum_s \frac{1}{m} \sum_t u_{st}$.

- Our Social Statewise Dominance axiom would become Individual statewise dominance, but its interpretation would otherwise be similar to the interpretation of statewise dominance in our main setting, holding that a risky intertemporal allocation f is better than another g if each state of f would be better, if received for certain, than the corresponding state of g, if received for certain.

- Our Stochastic Dominance for Sure Individuals would become

[13] Pure social time preference could be accommodated by period weights which would be analogous to probabilities in our interpretation. Note that Blackorby et al. (1995), in an early contribution to population ethics, also derive additive separability from lives born at different times, but consider only lifetime utilities, not period-specific utility flows.

[14] In fact, because they compare time periods with other time periods, without integrating over time and without time discounting, our Proposition 1 is sufficient to justify their approach.

> Stochastic dominance for fixed-longevity outcomes, holding that a risky intertemporal allocation f is better than another g if
>
> – every time period in which the decision-maker is not certain to live is unaffected by a choice between f and g, and
>
> – every time period in which the decision-maker is certain to live has a period-specific lottery of well-being in f that dominates that period's lottery of well-being in g.

This would be a novel justification of individual-level expected utility and of evaluating lifetime utility as the sum of period utility flows. As in the macroeconomic interpretation, the axioms rule out certain pattern goods. So whether or not this application makes normative sense for a prudent decision-maker may depend upon your interpretation of personal identity over a lifetime and whether lifetime pattern goods make sense and are valuable.

6. Discussion and conclusion

6.1 Extension to more general probability distributions

Until now, we have assumed that we have a finite number m of states of the world, all of them having the same probability $1/m$ to occur. In this section, we show that the results very easily extend to a case with states of the world whose probability of occurrence is a rational number.

Assume that there exists an infinite set of states of the world S, with typical element $s \in S$. We denote with Σ a σ-algebra over S. So, (S, Σ) is a measurable space, and we assume that there is a probability measure P on this measurable space. We define social prospects as measurable functions from S to U and assume that we only look at simple prospects, that is prospects such that only finitely many outcomes in U have a positive probability to be obtained. More technical details are provided in Supplementary materials, but here we just sketch how the extension works.

We need to make the following richness assumption:

> Furthermore, for any $m \in \mathbb{N}$, there exists a partition of S into m measurable events, $(E^1, \cdots, E^m)$ such that $P(E^k) = 1/m$ for all $k \in \{1, \cdots, m\}$.

Social prospects that take values in the partition $(E^1, \cdots, E^m)$ are formally similar to social prospects on m equiprobable states of the world, which is the framework we have studied. So, all our results extend to that case, whatever the number $m \geq 2$. What we can prove next, is: first, that if probabilities are rational numbers, we can match the prospects under consideration to other prospects that take values on a partition with equiprobable events and which are socially indifferent to the initial prospects; second, when probabilities are not rational numbers, we can add a property of continuity based on convergence in probability to extend our results.

Remark that all our axiom straightforwardly extend to the more general framework with infinitely many states, except for Individual Stochastic Dominance and Stochastic Dominance for Sure Individuals. For Individual Stochastic Dominance, we only have to adjust the definition of the cumulative distribution functions q_i^f. Specifically, we define the individual cumulative distribution function associated to a prospect, $q_i^f : \mathbb{R} \to [0, 1]$, by $q_i^f(z) = P\big(\{s \in S | f_i(s) \leq z\}\big)$ for any $z \in \mathbb{R}$. Using this new definition in the formulation of Individual Stochastic Dominance is sufficient for our purpose.

The formulation of Stochastic Dominance for Sure Individuals is more complex. It must be adapted in the following way:

Stochastic Dominance for Sure Individuals For all $f, g \in F$, if:

(i) $S_i(f) = S_i(g)$ for all $i \in I$;

(ii) for all $j \in I$ such that $S_j(f) \notin \{\varnothing, S\}$, there exists $x_j \in \mathbb{R}$ such that $f_j(s) = g_j(s) = x_j$ for all $s \in S_j(f)$;

(iii) there exists $\ell \in \mathbb{N}$ and two partitions $\{E_1, \cdots, E_\ell\}$ and $\{\tilde{E}_1, \cdots, \tilde{E}_\ell\}$ such that $P(E_r) = P(\tilde{E}_r)$ for each $r \in \{1, \cdots, \ell\}$, and for all $k \in I$ such that $S_k(f) = S$ and $r \in \{1, \cdots, \ell\}$, $f_k(s) \geq g_k(s')$ whenever $s \in E_r$ and $s' \in \tilde{E}_r$;

then $f \succsim g$.
If in addition there exists $h \in I$ such that $S_h(f) = S$ and $r' \in \{1, \cdots, \ell\}$ such that $f_h(s) > g_h(s')$ for all $s \in E_{r'}$ and $s' \in \tilde{E}_{r'}$ then $f \succ g$.

The difficulty in the formulation is that we cannot just permute states, we need to make sure that we preserve the measure of each outcome.

Hence, we need to associate each event E_r to another event $\tilde{E}_r$ of same measure and make sure that each concerned individual i obtain a better outcome with f on E_r than with g on $\tilde{E}_r$.

6.2 RELATED LITERATURE

Our paper joins a recent literature that has characterized objective functions with two dimensions of value. A theme of this literature is that separability in one dimension creates pressure for separability in another. None of these papers connect axioms as weak as ours to a conclusion as strong as ours.

Harsanyi's (1955) aggregation theorem is recognized as a foundation of utilitarian welfare economics, which is widely used throughout macroeconomics and public economics. As Fleurbaey (2009) summarized, Harsanyi showed that "in the presence of risk, weighted utilitarianism is the only criterion that satisfies the ex-ante Pareto principle and can be written as the expected value of social welfare," where *ex-ante* Pareto, in Harsanyi's case, meant assuming complete individual expected utilities. Harsanyi's result has received much attention and has been weakened in several directions. Fleurbaey (2009), in a founding contribution to this recent literature, weakens Harsanyi's assumptions in a setting of fixed-population social risk. Fleurbaey uses a weak dominance axiom like ours for social risk, but maintains an assumption of expectation-taking for individual *ex-ante* Pareto. In an uncertainty framework à la Savage, without objective probabilities, Mongin and Pivato (2015) obtained the Harsanyi's result with assumptions akin to statewise dominance for the social ordering and *ex-ante* Pareto for individuals, without assuming that individuals maximize an expected utility. A similar result is obtained by Zuber (2016) in an uncertainty framework à la Anscombe–Aumann. Li et al. (2023) recall the generality of this result that applies also to the context of risk and time or time and individuals as explained above. One way that all of these axiomatizations are stronger than ours is in requiring an individual order, where our axiom for individuals requires only dominance; also we do not assume a complete social ordering of all prospects.

Another contribution is the paper by McCarthy et al. (2020). They consider a framework with objective probabilities and use the property of Anteriority, which is related to our properties of Individual Stochastic Dominance and Stochastic Dominance for Sure Individuals. They obtain a “quasi-utilitarian” result with axioms that are similar to ours. But it must be clarified that their result is not exactly utilitarian in the sense that we use here. What they get is that the society should evaluate social prospects as if one of the individuals in the society was facing an “average prospect”, in the sense that she faces the prospect of each individual with equal probability. To clarify the difference, assume that individuals assess prospects only on the basis of first order stochastic dominance (to be consistent with our axioms). Consider a society with two individuals and two prospects: in one prospect the two individuals get utility $1/2$ for sure, in the other prospect one individual gets utility 1 for sure and the other gets 0 for sure. McCarthy et al. (2020) require that we assess these prospects like an individual would do if she compared a sure outcome of $1/2$ with the lottery of having 0 or 1 with equal probability. Given that the individual uses first order stochastic dominance, these two prospects cannot be compared. On the contrary, our approach can compare them and will prefer the former to the latter if $\phi(1/2) > \frac{1}{2}\phi(0) + \frac{1}{2}\phi(1)$ — for instance when ϕ is concave. McCarthy et al. (2020) could obtain this result by further assuming that individuals maximize an expected utility — which we do not assume.

Harsanyi (1955) only considered a fixed-population case. We show that the axioms leading to Harsanyi’s result can be significantly weakened in a variable-population setting. There exist other extensions of Harsanyi’s to the variable population framework. A founding result is by Blackorby et al. (1998) but they assume social expected utility as well as some utility independence for unconcerned individuals (or individual-level expected utility). Other, more recent, papers do combine the logic of two dimensions with variable population. Spears and Zuber (2022), for example, extend Harsanyi’s result to variable population, but maintain an assumption of social expected utility throughout. McCarthy et al. (2020) that we mentioned above is a recent contribution with wide mathematical generality, including the variable-population case. Their variable-population results differ from ours in assuming an axiom that they call Omega Independence that contains a comparison of existence in a risky outcome to non-existence. As explained above, they also do not get a generalized utilitarian criterion in the sense that we use here. Finally, Thomas (2022) offers an overview of the relationship between separability and additivity for the philosophical population-ethics literature. Thomas makes use of the Anteriority axiom that we have discussed.

Any axiomatization of a social welfare function can be read as an argument for that approach or as a warning of what the approach entails, depending upon one's perspective. To a reader who shares our interpretation that our axioms are weak and normatively attractive, our result raises the theoretical cost of departing either from additively-separable utilitarianism or from social expectation-taking. We note again that none of the axioms in our variable-population theorem is individually inconsistent with the egalitarian choice that $f^{**} \succ g^{**}$ in Table 2.

Acknowledgements

We thank Marcus Pivato and Itai Sher for very thoughtful comments on an earlier version of the paper.

A. Proof of Theorem 1

Proof. The proof has two steps.

Step 1: an additive representation for $\succsim$ for sure prospects. By Completeness for sure prospects, we know that $\succsim$ is a complete pre-order for sure prospects. By Lemma 1, $\succsim$ satisfies Separability for Sure Prospects. By the definition of $\succsim$ and Continuity, it is also continuous. By definition of $\succsim$ and Individual stochastic dominance, it is easily shown that $\succsim$ satisfies the following *Pareto-like property* (mentioned in the text on page 10):

for any $u, v \in \mathbb{R}^I$ if $u \geq v$ and $u \neq v$, then $u \succ v$,

where $\geq$ means at least as good for each person and better for one person. Hence, by the well-known result of Debreu (1960), there exist continuous and increasing functions $\phi_i : \mathbb{R} \to \mathbb{R}$ such that, for all $u, v \in U$,

$$u \succsim v \Longleftrightarrow \sum_{i \in I} \phi_i(u_i) \geq \sum_{i \in I} \phi_i(v_i).$$

Without loss of generality, we can normalize the ϕ_i functions so that $\phi_i(0) = 0$ and $\sum_{i \in I} \phi_i(1) = 1$.

Step 2: an expected utility representation. Consider any $f, g \in F$. Let us first construct $\hat{f}^{(1)}, \hat{g}^{(1)} \in F$ in the following way, using Compensation: $\hat{f}_1^{(1)}(1) = z^{(1)}$ and $\hat{g}_1^{(1)}(1) = \bar{z}^{(1)}$ while $\hat{f}_i^{(1)}(1) = \hat{g}_i^{(1)}(1) = 0$ for all $i \neq 1$, where $z^{(1)}$ and $\bar{z}^{(1)} \in \mathbb{R}$ are such that $\hat{f}^{(1)}(s) \sim f(s)$ and $\hat{g}^{(1)}(s) \sim g(s)$

(we know that such $z^{(1)}$ and $\bar{z}^{(1)}$ exist by Compensation). For each $s > 1$, $\hat{f}^{(1)}(s) = f(s)$ and $\hat{g}^{(1)}(s) = g(s)$ so that $\hat{f}^{(1)} \sim f$ and $\hat{g}^{(1)} \sim g$ by Social statewise dominance. By Step 1, it is the case that $\phi_1(\hat{f}_1^{(1)}(1)) = \sum_{i \in I} \phi_i(f_i(1))$ and $\phi_1(\hat{g}_1^{(1)}(1)) = \sum_{i \in I} \phi_i(g_i(1))$.

The next move is to construct two sequences of prospects $\hat{f}^{(1)}, \cdots, \hat{f}^{(m)}$ and $\hat{g}^{(1)}, \cdots, \hat{g}^{(m)}$ with the following properties:

- $\phi_1\big(\hat{f}_1^{(k)}(1)\big) = \sum_{s=1}^{k} \sum_{i \in I} \phi_i(f_i(s))$, $\phi_1\big(\hat{g}_1^{(k)}(1)\big) = \sum_{s=1}^{k} \sum_{i \in I} \phi_i(g_i(s))$, and $\hat{f}_i^{(k)}(1) = \hat{g}_i^{(k)}(1) = 0$ for all $i > 1$;
- $\hat{f}_i^{(k)}(s) = \hat{g}_i^{(k)}(s) = 0$ for all $i \in I$ and $2 \le s \le k$;
- $\hat{f}^{(k)}(s) = f(s)$ and $\hat{g}^{(k)}(s) = g(s)$ for all $s > k$;
- $\hat{f}^{(k+1)} \sim \hat{f}^{(k)}$ and $\hat{g}^{(k+1)} \sim \hat{g}^{(k)}$ for all $k = 1, \cdots, m-1$.

Let us show that the construction is possible by recursion. Notice that all the properties (except the last) are already satisfied by $\hat{f}^{(1)}$ and $\hat{g}^{(1)}$. Let $k \in \{1, \cdots, m-1\}$ and assume that we have constructed $\hat{f}^{(k)}$. Let us show that we can construct $\hat{f}^{(k+1)}$ with the desired properties so that $\hat{f}^{(k+1)} \sim \hat{f}^{(k)}$ (the proof is similar for $\hat{g}^{(1)}, \cdots, \hat{g}^{(m-1)}$, and thus not repeated).

By Compensation, there exists a number $\tilde{z}^{(k+1)} \in \mathbb{R}$ such that, if we define $\tilde{u}^{(k+1)} \in U$ by $\tilde{u}_2^{(k+1)} = \tilde{z}^{(k+1)}$ and $\tilde{u}_j^{(k+1)} = 0$ for all $j \in I \setminus \{2\}$, it is the case that $\tilde{u}^{(k+1)} \sim \hat{f}^{(k)}(k+1)$. By construction and step 1, it is the case that

$$\phi_2\left(\tilde{u}_2^{(k+1)}\right) = \sum_{i \in I} \phi_i\left(\hat{f}_i^{(k)}(k+1)\right) = \sum_{i \in I} \phi_i\left(f_i(k+1)\right). \tag{1}$$

Define $\tilde{f}^{(k+1)}$ by $\tilde{f}^{(k+1)}(k+1) = \tilde{u}^{(k+1)}$ and $\tilde{f}^{(k+1)}(s) = \hat{f}^{(k)}(s)$ for all $s \neq k+1$. Social statewise dominance gives $\tilde{f}^{(k+1)} \sim \hat{f}^{(k)}$. Next construct $\bar{f}^{(k+1)}$ in the following way: $\bar{f}_i^{(k+1)}(s) = \tilde{f}_i^{(k+1)}(s)$ for all $s \in S$ and $i \neq 2$; $\bar{f}_2^{(k+1)}(1) = \tilde{f}_2^{(k+1)}(k+1)$, $\bar{f}_2^{(k+1)}(k+1) = 0$, while $\bar{f}_2^{(k+1)}(s) = \tilde{f}_2^{(k+1)}(s)$ for all $s \neq 1, k+1$. Individual 2 faces the same individual prospect in $\bar{f}^{(k+1)}$ and $\tilde{f}^{(k+1)}$, while all other individuals are not affected. By Individual stochastic dominance, $\bar{f}^{(k+1)} \sim \tilde{f}^{(k+1)}$, and by transitivity $\bar{f}^{(k+1)} \sim \hat{f}^{(k)}$.

The prospect $\bar{f}^{(k+1)}$ is such that $\bar{f}_1^{(k+1)}(1) = \hat{f}_1^{(k)}(1)$, $\bar{f}_2^{(k+1)}(1) = \tilde{f}_2^{(k+1)}(k+1) = \tilde{u}_2^{(k+1)}$ and $\bar{f}_i^{(k+1)}(1) = 0$ for all $i > 2$. By Compensation, there exists a number $z^{(k+1)} \in \mathbb{R}$ such that, if we define $\bar{u}^{(k+1)} \in U$ by $\bar{u}_1^{(k+1)} = z^{(k+1)}$ and $\bar{u}_j^{(k+1)} = 0$ for all $j \in I \setminus \{1\}$, it is the case that

$\bar u^{(k+1)} \sim \bar f^{(k+1)}(1)$. By construction and step 1, it is also the case that

$$\phi_1\left(\bar u_1^{(k+1)}\right) = \sum_{i\in I}\phi_i\left(\bar f^{(k+1)}(1)\right) = \phi_1\left(\hat f_1^{(k)}(1)\right)+\phi_2\left(\tilde u_2^{(k+1)}\right) = \sum_{s=1}^{k+1}\sum_{i\in I}\phi_i(f_i(s)).$$

(Recall that $\phi_1\left(\hat f_1^{(k)}(1)\right) = \sum_{s=1}^{k}\sum_{i\in I}\phi_i(f_i(s))$ and $\phi_2\left(\tilde u_2^{(k+1)}\right) = \sum_{i\in I}\phi_i\left(f_i(k+1)\right)$ – see Equation (1)). It suffices to define $\hat f^{(k+1)}$ by $\hat f^{(k+1)}(1) = \bar u^{(k+1)}$ and $\hat f^{(k+1)}(s) = \bar f^{(k)}(s)$ for all $s > 1$ to obtain $\bar f^{(k+1)} \sim \hat f^{(k+1)}$ by Social statewise dominance. By transitivity, $\hat f^{(k)} \sim \hat f^{(k+1)}$. It can be checked that $\hat f^{(k+1)}$ has all the aforementioned features. Figure 1 describes the step between $\hat f^{(k)}$ and $\hat f^{(k+1)}$.

By our construction and transitivity, we have $f \sim \hat f^{(m)}$ and $f \sim \hat g^{(m)}$. But $\hat f^{(m)}$ and $\hat g^{(m)}$ are such that $\hat f_i^{(m)}(s) = \hat g_i^{(m)}(s) = 0$ for all $i \in I$ and $s > 1$. By Social statewise dominance and Completeness for sure prospects, we know that $\hat f^{(m)} \succsim \hat g^{(m)} \iff \hat f^{(m)}(1) \succsim \hat g^{(m)}(1)$. By transitivity, we also have $f \succsim g \iff \hat f^{(m)}(1) \succsim \hat g^{(m)}(1)$.

Figure 1: Construction of prospect $\hat f^{(k+1)}$ for $k \geq 2$

	$\hat f^{(k)}$				$\tilde f^{(k+1)}$			
	1	2	3	$\cdots$	1	2	3	$\cdots$
1	$\hat f_1^{(k)}(1)$	0	0	$\cdots$	$\hat f_1^{(k)}(1)$	0	0	$\cdots$
2	0	0	0	$\cdots$	0	0	0	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
k	0	0	0	$\cdots$	0	0	0	$\cdots$
$k+1$	$f_1(k+1)$	$f_2(k+1)$	$f_3(k+1)$	$\cdots$	0	$\tilde f_2^{(k+1)}(k+1)$	0	$\cdots$
$k+2$	$f_1(k+2)$	$f_2(k+2)$	$f_3(k+2)$	$\cdots$	$f_1(k+2)$	$f_2(k+2)$	$f_3(k+2)$	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$

	$\bar f^{(k+1)}$				$\hat f^{(k+1)}$			
	1	2	3	$\cdots$	1	2	3	$\cdots$
1	$\hat f_1^{(k)}(1)$	$\tilde f_2^{(k+1)}(k+1)$	0	$\cdots$	$\hat f_1^{(k+1)}(1)$	0	0	$\cdots$
2	0	0	0	$\cdots$	0	0	0	$\cdots$
k	0	0	0	$\cdots$	0	0	0	$\cdots$
$k+1$	0	0	0	$\cdots$	0	0	0	$\cdots$
$k+2$	$f_1(k+2)$	$f_2(k+2)$	$f_3(k+2)$	$\cdots$	$f_1(k+2)$	$f_2(k+2)$	$f_3(k+2)$	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$

Using Step 1 and the definition of $\hat{f}^{(m)}$ and $\hat{g}^{(m)}$ we get:

$$
\begin{aligned}
f \succsim g &\iff \hat{f}^{(m)}(1) \succsim \hat{g}^{(m)}(1) \\
&\iff \sum_{i \in I} \phi_i\left(\hat{f}_i^{(m)}(1)\right) \geq \sum_{i \in I} \phi_i\left(\hat{g}_i^{(m)}(1)\right) \\
&\iff \sum_{s \in S} \sum_{i \in I} \phi_i\left(f_i(s)\right) \geq \sum_{s \in S} \sum_{i \in I} \phi_i\left(g_i(s)\right) \\
&\iff \sum_{s \in S} \frac{1}{m} \sum_{i \in I} \phi_i\left(f_i(s)\right) \geq \sum_{s \in S} \frac{1}{m} \sum_{i \in I} \phi_i\left(g_i(s)\right).
\end{aligned}
$$

∎

B. Proof of Theorem 2

Proof. It is straightforward to check that Generalized Expectational Total Utilitarianism satisfies all of our six principles.

Let us show that the six principles imply Generalized Expectational Total Utilitarianism. By Proposition 1, we know that there exists a continuous and increasing function $\phi : \mathbb{R} \to \mathbb{R}$ and a number $c \in \mathbb{R}$ such that for all $u, v \in U$, $u \succsim v$ if and only if $\sum_{i \in N(u)} \left[\phi(u_i) - \phi(c)\right] \geq \sum_{i \in N(v)} \left[\phi(v_i) - \phi(c)\right]$.

Consider any social prospect $f \in F$. Let us construct the social prospect $\tilde{f} \in F$ with the following properties:

- There exists a collection of state-indexed populations $N^1, \cdots, N^m$ such that: (i) $|N^s| = |N(f(s))|$ for all $s \in S$; (ii) $N^{s'} \cap N^s = \varnothing$ for all $s' \neq s$;
- $N\left(\tilde{f}(s)\right) = \bigcup_{s' \in S} N^{s'}$ for all $s \in S$;
- There exist bijections $\sigma^s : N^s \to N\big(f(s)\big)$ such that $\tilde{f}_i(1) = f_{\sigma^s(i)}(s)$ for all $i \in N^s$;
- $\tilde{f}_i(s) = c$ for all $i \in N\left(\tilde{f}(s)\right)$ when $s \in \{2, \cdots, m\}$.

Social prospect $\tilde{f}$ is a prospect where all utility levels of all states of the world have been moved to state 1 (by creating new people), and all individuals have level c or do not exist in other states of the world. We want to show that $f \sim \tilde{f}$. Notice that, by the definition of $\tilde{f}$:

$$\begin{aligned}\sum_{i \in N(\tilde{f}(1))} \left[\phi\left(\tilde{f}_i(1)\right) - \phi(c)\right] &= \sum_{s \in S} \sum_{j \in N^s} \left[\phi\left(\tilde{f}_j(1)\right) - \phi(c)\right] \\ &= \sum_{s \in S} \sum_{j \in N^s} \left[\phi(f_{\sigma^s(j)}(s)) - \phi(c)\right] \\ &= \sum_{s \in S} \sum_{k \in N(f(s))} \left[\phi\left(f_k(s)\right) - \phi(c)\right].\end{aligned}$$

To show that $f \sim \tilde{f}$, let us construct two sequences of social prospects $(\hat{f}^{(1)}, \cdots, \hat{f}^{(m)})$ and $(\tilde{f}^{(1)}, \cdots, \tilde{f}^{(m)})$ in the following way.

We have $\hat{f}^{(1)} = \tilde{f}^{(1)}$, defined as follows: $N\left(\tilde{f}^{(1)}\right)(1) = \bigcup_{s' \in S} N^{s'}$, $\tilde{f}_i^{(1)}(1) = f_{\sigma^1(i)}(1)$ for all $i \in N^1$, and $\tilde{f}_j^{(1)}(1) = c$ for all $j \in \left(N\left(\tilde{f}^{(1)}\right)(1) \setminus N^1\right)$; for all $s \geq 2$, $N\left(\tilde{f}^{(1)}\right)(s) = N^s$ and $\tilde{f}_i^{(1)}(s) = f_{\sigma^1(i)}(s)$ for all $i \in N^s$.

For any $k \in \{2, \cdots, m\}$:

- $N\left(\hat{f}^{(k)}(1)\right) = N\left(\tilde{f}^{(k)}(1)\right) = \bigcup_{s'=1}^{k} N^{s'}$; $\hat{f}_i^{(k)}(1) = \tilde{f}_i^{(k)}(1) = f_{\sigma^s(i)}(s)$ for all $i \in N^s$ and $s < k$; $\hat{f}_j^{(1)}(1) = c$ and $\tilde{f}_j^{(1)}(1) = f_{\sigma^k(j)}(k)$ for all $j \in N^k$;
- For all $1 < s < k$, $N\left(\hat{f}^{(k)}(s)\right) = N\left(\tilde{f}^{(k)}(s)\right) = N^k$, and $\hat{f}_i^{(k)}(s) = \tilde{f}_i^{(k)}(s) = c$ for all $i \in N^k$;
- $N\left(\hat{f}^{(k)}(k)\right) = N\left(\tilde{f}^{(k)}(k)\right) = N^k$, $\hat{f}_i^{(k)}(k) = f_{\sigma^k(i)}(k)$ and $\tilde{f}_i^{(k)}(k) = c$ for all $i \in N^k$;
- For all $s > k$, $N\left(\hat{f}^{(k)}(s)\right) = N\left(\tilde{f}^{(k)}(s)\right) = N^k \cup N^s$, $\hat{f}_i^{(k)}(s) = \tilde{f}_i^{(k)}(s) = c$ for all $i \in N^k$, and $\hat{f}_j^{(k)}(s) = \tilde{f}_j^{(k)}(s) = f_{\sigma^s(j)}(s)$ for all $j \in N^s$.

Figure 2 illustrates those social prospects.

By Statewise dominance, $\tilde{f}^{(k)} \sim \hat{f}^{(k+1)}$ for any $k \in \{1, \cdots, m-1\}$; indeed, $\tilde{f}^{(k)}$ and $\hat{f}^{(k+1)}$ differ only in each state of the world (except state 1 where they are identical) by the existence of people with utility level c. By Proposition 1 and CLGU with critical-level c, they are thus equivalent in each state of the world. On the other hand, $\tilde{f}^{(k)} \sim \hat{f}^{(k)}$ for any $k \in S$ by Stochastic dominance for sure individuals. Indeed, only individuals in N^k face different prospects, but they live for sure and their utility is permuted from state 1 to state k.

Figure 2: Construction of prospects $\hat{f}^{(k)}$ and $\tilde{f}^{(k)}$ for $k \geq 2$. Like in the main text, Ω denotes non-existence, here appplied to a group of persons.

$\hat{f}^{(k)}$

	N^1	$\cdots$	N^{k-1}	N^k	N^{k+1}	N^{k+2}	$\cdots$
1	$(f_i(1))_{i\in N^1}$	$\cdots$	$(f_i(k-1))_{i\in N^{k-1}}$	c	Ω	Ω	$\cdots$
2	Ω	$\cdots$	Ω	c	Ω	Ω	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
$k-1$	Ω	$\cdots$	Ω	c	Ω	Ω	$\cdots$
k	Ω	$\cdots$	Ω	$(f_i(k))_{i\in N^k}$	Ω	Ω	$\cdots$
$k+1$	Ω	$\cdots$	Ω	c	$(f_i(k+1))_{i\in N^{k+1}}$	Ω	$\cdots$
$k+2$	Ω	$\cdots$	Ω	c	Ω	$(f_i(k+2))_{i\in N^{k+2}}$	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$

$\tilde{f}^{(k)}$

	N^1	$\cdots$	N^{k-1}	N^k	N^{k+1}	N^{k+2}	$\cdots$
1	$(f_i(1))_{i\in N^1}$	$\cdots$	$(f_i(k-1))_{i\in N^{k-1}}$	$(f_i(k))_{i\in N^k}$	Ω	Ω	$\cdots$
2	Ω	$\cdots$	Ω	c	Ω	Ω	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
$k-1$	Ω	$\cdots$	Ω	c	Ω	Ω	$\cdots$
k	Ω	$\cdots$	Ω	c	Ω	Ω	$\cdots$
$k+1$	Ω	$\cdots$	Ω	c	$(f_i(k+1))_{i\in N^{k+1}}$	Ω	$\cdots$
$k+2$	Ω	$\cdots$	Ω	c	Ω	$(f_i(k+2))_{i\in N^{k+2}}$	$\cdots$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$

We thus obtain the chain of equivalences $\tilde{f}^1 \sim \hat{f}^{(2)} \sim \tilde{f}^{(2)} \sim \cdots \sim \tilde{f}^{(m-1)} \sim \hat{f}^{(m)}$. In addition, $f \sim \tilde{f}^{(1)}$ and $\hat{f}^{(m)} \sim \tilde{f}$ by Statewise dominance (they differ only in each state of the world by the set of people with certain utility levels, and/or a number of people at critical-level c). So, by transitivity $f \sim \tilde{f}$.

Consider any f and $g \in F$. By the arguments above, there exist $\tilde{f}$ and $\tilde{g}$ such that (where ϕ and c are given by Proposition 1):

- $f \sim \tilde{f}$ and $g \sim \tilde{g}$;
- $\sum_{i\in N(\tilde{f}(1))} \left[\phi\left(\tilde{f}_i(1)\right) - \phi(c)\right] = \sum_{s\in S} \sum_{j\in N(f(s))} \left[\phi\left(f_j(s)\right) - \phi(c)\right]$;
- $\sum_{i\in N(\tilde{g}(1))} \left[\phi\left(\tilde{g}_i(1)\right) - \phi(c)\right] = \sum_{s\in S} \sum_{j\in N(g(s))} \left[\phi\left(g_j(s)\right) - \phi(c)\right]$; and
- for all $s \in \{2, \cdots, m\}$, $\tilde{f}_i(s) = c$ for all $i \in N(\tilde{f}(s))$ and $\tilde{g}_j(s) = c$ for all $j \in N(\tilde{g}(s))$.

By Proposition 1, $\tilde{f}(s) \sim \tilde{g}(s)$ for all $s \in \{2, \cdots, m\}$, so that, by Statewise dominance and completeness for sure prospects $\tilde{f} \succsim \tilde{g} \iff \tilde{f}(1) \succsim$

$\tilde{g}(1)$. Gathering all the results, we obtain:

$$
\begin{aligned}
f \succsim g &\iff \tilde{f} \succsim \tilde{g} \\
&\iff \tilde{f}(1) \succsim \tilde{g}(1) \\
&\iff \sum_{i \in N(\tilde{f}(1))} \left[\phi\left(\tilde{f}_i(1)\right) - \phi(c)\right] \geq \sum_{j \in N(\tilde{g}(1))} \left[\phi\left(\tilde{g}_j(1)\right) - \phi(c)\right] \\
&\iff \sum_{s \in S} \sum_{i \in N(f(s))} \left[\phi\left(f_i(s)\right) - \phi(c)\right] \geq \sum_{s \in S} \sum_{j \in N(g(s))} \left[\phi\left(g_j(s)\right) - \phi(c)\right] \\
&\iff \sum_{s \in S} \frac{1}{m} \left[\sum_{i \in N(f(s))} \left[\phi\left(f_i(s)\right) - \phi(c)\right] \right] \geq \sum_{s \in S} \frac{1}{m} \left[\sum_{j \in N(g(s))} \left[\phi\left(g_j(s)\right) - \phi(c)\right] \right].
\end{aligned}
$$

∎

References

Asheim, G. B. and S. Zuber (2016) 'Evaluating Intergenerational Risks', *Journal of Mathematical Economics* 65:104–117.

Blackorby, C., W. Bossert, and D. Donaldson (1995) 'Intertemporal Population ethics: Critical-level utilitarian principles', *Econometrica* 63:1303–1320.

— (1996) 'Quasi-Orderings and Population Ethics', *Soc Choice Welfare* 13:129–151.

— (1998) 'Uncertainty and Critical-Level Population Principles', *Journal of Population Economics* 11:1–20.

— (2005) *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, Cambridge University Press.

Blackorby, C. and D. Donaldson (1984) 'Social Criteria for Evaluating Population Change', *Journal of Public Economics* 25:13–33.

Broome, J. (1991) *Weighing Goods: Equality, Uncertainty and Time*, Oxford: Blackwell.

Debreu, G. (1960) 'Topological Methods in Cardinal Utility Theory', in *Mathematical Methods in the Social Sciences 1959*, pp. 16–26, Stanford University Press.

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2022) 'The Economics of Fertility: A New Era', Working Paper w29948, National Bureau of Economic Research.

Fleurbaey, M. (2009) 'Two Variants of Harsanyi's Aggregation Theorem', *Economics Letters* 105:300–302.

— (2010) 'Assessing Risky Social Situations', *Journal of Political Economy* 118:649–680.

Gilboa, I. and D. Schmeidler (1989) 'Maxmin Expected Utility with Non-Unique Prior', *J. Math. Econ.* 18:141–153.

Golosov, M., L.E. Jones, and M. Tertilt (2007) 'Efficiency with Endogenous Population Growth', *Econometrica* 75 (4): 1039–1071.

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

— (2022) *Money-Pump Arguments*, Cambridge: Cambridge University Press.

Hammond, P. J. (1988) 'Consequentialist Foundations for Expected Utility', *Theory and Decision* 25 (1): 25–78.

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

Jones, C. I (2022) 'The end of economic growth? Unintended consequences of a declining population', *American Economic Review* 112 (11): 3489–3527.

Kearney, M.S., P.B. Levine, and L. Pardue (2022) 'The Puzzle of Falling US Birth Rates since the Great Recession', *Journal of Economic Perspectives* 36 (1): 151–76.

Klenow, P., C. Jones, M. Bils, and M. Adhami (2022) 'Population and Welfare: The Greatest Good for the Greatest Number', Stanford University working paper.

Li, C., K.I.M. Rhodes, and 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.

McCarthy, D., K. Mikkola, and T. Thomas (2020) 'Utilitarianism with and without Expected Utility', *Journal of Mathematical Economics* 87:77–113.

Mongin, P. and M. Pivato (2015) 'Ranking Multidimensional Alternatives and Uncertain Prospects', *Journal of Economic Theory* 157:146–171.

Myerson, R.B. (1981) 'Utilitarianism, Egalitarianism, and the Timing Effect in Social Choice Problems', *Econometrica* 49 (4): 883–897.

Nordhaus, W.D. (2017) 'Revisiting the Social Cost of Carbon', *Proceedings of the National Academy of Sciences* 114 (7): 1518–1523.

Parfit, D. (1984) *Reasons and Persons*, Oxford: Clarendon Press.

— (1995) *Equality or Priority?*, Lawrence: University of Kansas.
— (n.d.) 'On Giving Priority to the Worse Off', unpublished manuscript.
Pivato, M. and E. F. Tchouante (2022) 'Bayesian Social Aggregation with Non-Archimedean Utilities and Probabilities', mimeo.
Rabinowicz, W. (2009) 'Broome and the Intuition of Neutrality', *Philosophical Issues* 19 (1): 389–411.
Spears, D. and S. Zuber (2022) 'Foundations of Utilitarianism under Risk and Variable Population', *Social Choice and Welfare*.
Thomas, T. (2022) 'Separability in Population Ethics', in *Oxford Handbook of Population Ethics*.
Zuber, S. (2016) 'Harsanyi's Theorem without the Sure-Thing Principle: On the Consistent Aggregation of Monotonic Bernoullian and Archimedean Preferences', *Journal of Mathematical Economics* 63:78–83.