Belief-Averaging and Relative Utilitarianism: Savage Meets Arrow

Florian Brandl[†]

(click for most recent revision)

We consider social welfare functions when the preferences of individual agents and society maximize subjective expected utility in the tradition of Savage. A system of axioms is introduced whose unique solution is the social welfare function that averages the agents' beliefs and sums up their utility functions, normalized to have the same range. The first distinguishing axiom requires that an act about which beliefs agree becomes socially more preferred if it gains support among the agents. The second is a weakening of Arrow's independence of irrelevant alternatives that only applies to non-redundant acts.

Keywords: Uncertainty, subjective expected utility, Pareto optimality, relative utilitarianism.

Harsanyi (1955) studied how a society should rank risky alternatives. His aggregation theorem shows that if the agents as well as society are von Neumann-Morgenstern expected utility maximizers and the preferences of society satisfy the Pareto principle, then the utility function of society is a linear combination of the agents' utility functions. Among others, two streams of literature originated from here.

The first studies the implications of the Pareto principle when the social alternatives are not lotteries but uncertain acts. The approach of Savage (1954) is to view acts as maps from states of the world to outcomes. A preference relation over acts maximizes subjective expected utility if it ranks acts by their expected utility for a utility function over outcomes and an idiosyncratic belief that assigns probabilities to states. Mongin (1995) showed that if agents and society are *subjective* expected utility maximizers, the Pareto principle implies that the preferences of society have to coincide with those of one of the agents unless all agents hold the same belief (in which case we are back to Harsanyi's setting). Mongin (1995, 1997) noted, however, that the Pareto principle also applies to spurious unanimities where all agents prefer

[†]Department of Economics, Princeton University, brandl.ffx@gmail.com

one act to another because of differences in beliefs *and* differences in utility functions. Thus, the Pareto principle in its full force is not necessarily compelling when applied to subjective expected utility maximizers. Gilboa, Samet, and Schmeidler (2004) introduced a restricted Pareto condition, which avoids spurious unanimities. It applies only to acts that induce the same distribution over outcomes for all agents' beliefs. Such acts correspond to lotteries in Harsanyi's setting. Thus, by Harsanyi's aggregation theorem, the restricted Pareto condition implies that the utility function of society is a linear combination of the agents' utility functions. The surprising part is that it also implies that the belief of society is an affine combination of the agents' beliefs.

The second stream stays with decisions under risk and deals with a question left open by Harsanyi's work: if we follow Harsanyi's utilitarian doctrine, how do we assign weights to the agents' utility functions? The approach pursued by Dhillon and Mertens (1999) considers social welfare functions in the tradition of Arrow (1951), which map every preference profile to a preference relation for society. Dhillon and Mertens impose a set of axioms on social welfare functions that, roughly speaking, constitute weakenings of the conditions in Arrow's impossibility theorem. They show that only *relative utilitarianism* satisfies all of their axioms: normalize the agents' utility functions so that their range is the unit interval and sum them up.[1] Their result is multi-profile since the axioms relate the preferences of society for different preference profiles to each other. By contrast, the results of Harsanyi, Mongin, and Gilboa et al. discussed above are single-profile since Pareto conditions talk about a single preference profile considered in isolation.

This paper seeks to combine both streams. That is, it takes a multi-profile approach to decisions under uncertainty. We consider social welfare functions when the preferences of the agents as well as the society are subjective expected utility maximizing. The goal is to determine a social welfare function that can be justified on axiomatic grounds. To this end, two axioms are introduced. First, consider a society that is indifferent between two acts and an outside agent whose belief induces the same distribution over outcomes as the belief underlying the preferences of the society for either of the two acts. If the agent is added to the society, the preference of the augmented society over the two acts should coincide with that of the additional agent. This axiom, called *restricted monotonicity*, applies the restricted Pareto condition of Gilboa et al. (2004) to the preferences of a society and an outside agent. Second, recall that Arrow's (1951) independence of irrelevant alternatives requires that society's preference over any pair of alternatives—acts in our case—depends only on the agents' preferences over these two acts. In the present setting, it forces the social welfare function to ignore the agents' expected utilities for the two acts apart from how they are ordered. The second axiom weakens Arrow's condition so that society's preference may also depend on the agents expected utilities for the two acts. (In fact, it is even weaker.) Referring to the independence of redundant alternatives axiom

[1] Several other characterizations of relative utilitarianism, for example, by Karni (1998); Dhillon (1998); Segal (2000); Börgers and Choo (2017), exist. The discussion focuses on the characterization of Dhillon and Mertens (1999) since it is most relevant to the present paper.

	Risk	Uncertainty
Single-profile	linear aggregation of utilities (Harsanyi, 1955)	linear aggregation of beliefs and utilities (Gilboa et al., 2004)
Multi-profile	relative utilitarianism (Dhillon and Mertens, 1999)	belief-averaging and relative utilitarianism

Table 1: Placement of the present paper relative to previous works. The papers in the "Risk" ("Uncertainty") column assume that agents as well as society are (subjective) expected utility maximizers. The single-profile results use only the (restricted) Pareto condition. The multi-profile results rely on several axioms, some of which relate the preferences of society for different preference profiles to each other.

of Dhillon and Mertens (1999) by which it is inspired, it is called *independence of redundant acts*. Both axioms are discussed in detail in the next section.

The main result of this paper shows that these two conditions, together with four undiscriminating axioms, characterize *belief-averaging and relative utilitarianism*: average the agents' beliefs and sum up their utility functions, normalized to the unit interval. The resulting belief and utility function determine the preferences of society. This extends Gilboa et al.'s linear aggregation result to a multi-profile framework and Dhillon and Mertens's relative utilitarianism to subjective expected utility maximizers. Table 1 summarizes how the present paper relates to the described works.

Discussion of the axioms Restricted monotonicity is defined as follows. Suppose we apply a social welfare function to a preference profile on some subset of the agents and the thus-derived preferences of society rank two acts f and g as equally desirable. Restricted monotonicity applies when we add an outside agent to the profile. If (i) this agent prefers f to g and (ii) for either of f and g, the agent's belief induces the same distribution over outcomes as the belief underlying the preferences of the original society, then the augmented society should prefer f to g as well.[2] More precisely, applying the social welfare function to the preference profile on the original subset of agents plus the additional agent should give a preference relation that ranks f above g.

Part (ii) above makes restricted monotonicity apply only to acts that are risky alternatives for the belief of the original society and the additional agent. That is, acts for which differences in those two beliefs are irrelevant and the utility functions are the only source of preference heterogeneity. The rationale for part (ii) is the same as that of Gilboa et al. (2004) for restricting the Pareto condition to acts that induce the same distribution over outcomes for all agents' beliefs: it avoids accidental preference agreements through differences in beliefs and utility

[2] Two beliefs induce the same distribution over outcomes for an act f if the push-forward measure on the set of outcomes under f is the same for both beliefs. Equivalently, both beliefs agree on the sigma-algebra over the set of states induced by f.

	belief		utility function			expected utilities	
	ω_1	ω_2	a	b	c	f	g
Agent 1	90%	10%	1	0	0.9	0.9	0.9
Agent 2	10%	90%	0	1	0.8	0.9	0.8

Table 2: Numerical values for Example 1. ω_1 and ω_2 are the two states of the world and a, b, and c are possible outcomes. f is an act that yields a if the state is ω_1 and b if the state is ω_2; g yields c in both states. The last two rows give the beliefs, utility functions, and expected utilities of Agent 1 and Agent 2.

functions. To see why restricted monotonicity is less substantiated if differences in beliefs matter, consider the following example.

Example 1. There are two states of the world ω_1 and ω_2, and three possible outcomes a, b, and c. Call f the act that results in a in the state ω_1 and in b in the state ω_2 and g the act that gives outcome c for both states. We consider two agents, Agent 1 and Agent 2, whose beliefs and utility functions are given in Table 2.[3] The expected utilities show that Agent 1 is indifferent between f and g and Agent 2 prefers f to g. An unrestricted monotonicity condition (without part (ii)) would thus demand that the society composed of Agent 1 and Agent 2 prefers f to g. With the beliefs and utilities in Table 2, this conclusion is questionable. Agent 1 evaluates f as a lottery with probability 90% on a and 10% on b, whereas for Agent 2 the probabilities are reversed; for both agents, g equals a sure bet on c. Hence, their expected utilities for f and g found on different views on the outcomes of the two acts. In particular, there is no belief so that if both agents held that belief, they would both prefer f to g.

This is not to say that any monotonicity condition stronger than the one defined above is undesirable. The approach here is merely to be cautious about the assumptions that are imposed. Note however that when requiring that the preferences of every single-agent society are those of its sole member (which we call faithfulness), an unrestricted monotonicity condition would imply the (unrestricted) Pareto principle and lead to an impossibility by the result of Mongin (1997).

The restricted monotonicity axiom and faithfulness together imply the restricted Pareto condition of Gilboa et al.. Hence, the belief and utility function of the society have to be linear combinations of the agents' beliefs and utility functions. The restricted Pareto condition allows the weights in both linear combinations to be arbitrary functions of *all* agents' preferences. The additional strength of restricted monotonicity implies that the weight of an agent in either linear combination depends only on the agent's own preferences. The restricted monotonicity axiom connects society's preferences for different sets of agents. We will thus consider social

[3]The structure of this example is very similar to the example of Gilboa et al. (2004) of the two gentlemen who are contemplating having a duel.

welfare functions that take as input the preferences of an arbitrary finite set of agents.

Arrow's (1951) independence of irrelevant alternatives prescribes that the preferences of society over any two acts must depend only on the agents' preferences over the two acts, and not on their preferences over other acts. It precludes that society's ranking of two acts depends more flexibly on the agents' expected utilities for the two acts. For example, consider two agents whose expected utilities for three acts f, g, and h are 1, 0, and α (for the first agent) and 0, 1, and α (for the second agent). Independence of irrelevant alternatives asserts that the preferences of this two-agent society over the three acts are the same for every value of α strictly between 0 and 1. However, it does not seem unreasonable that h is society's least preferred act if α is close to 0 and the most preferred act if α is close to 1. This is ruled out by independence of irrelevant alternatives. Independence of redundant acts weakens independence of irrelevant alternatives so that it does not apply to profiles with different values for α.

Since the agents report preference relations that identify utility functions only up to scaling, we need to capture the above intuition on the level of preferences. Call a set of acts $\mathcal{A}'$ *co-redundant* for a profile if every act is unanimously indifferent to some act in $\mathcal{A}'$.[4] Every act outside a co-redundant set $\mathcal{A}'$ is redundant in the sense that $\mathcal{A}'$ contains an act that all agents consider equally desirable. Independence of redundant acts demands that if $\mathcal{A}'$ is co-redundant for two profiles and every agent has the same preferences over acts in $\mathcal{A}'$ in both profiles, then society's preferences over $\mathcal{A}'$ are the same for both profiles. Thus, society's preferences over the co-redundant acts must not depend on the agents' preferences over the redundant acts. This condition restricts independence of irrelevant alternatives to pairs of profiles that are equal except for preferences over redundant acts. Independence of redundant acts extends Dhillon and Mertens's (1999) independence of redundant alternatives to decisions under uncertainty.

It is convenient to consider the situation in utility space. Given a preference profile, each act maps to a vector of (subjective expected) utilities with one component for every agent. One can interpret the image of this map as the bargaining problem associated with the profile. A set of acts is co-redundant for a profile if and only if removing all other acts does not change the bargaining problem. Hence, two profiles that agree on a co-redundant set induce the same bargaining problem, and every act in the co-redundant set maps to the same point. In particular, independence of redundant acts applies only to sets of acts for which the agents' expected utilities are the same for both profiles. Figure 1 illustrates the preceding discussion.

On top of restricted monotonicity and independence of redundant acts, we make four assumptions about the social welfare function: the preferences of any single-agent society are those of its sole member (*faithfulness*), no agent can impose its belief on a society (*no belief imposition*), the preferences of society depend continuously on the preferences of its members (*continuity*), and all agents are treated symmetrically (*anonymity*). Together these six conditions characterize belief-averaging and relative utilitarianism.

The proof is modular and yields two intermediary results that are interesting in their own

[4] In the formal definition in Section 2, we will additionally require that each agent's preferences over $\mathcal{A}'$ are subjective expected utility-maximizing for a *non-atomic* belief.

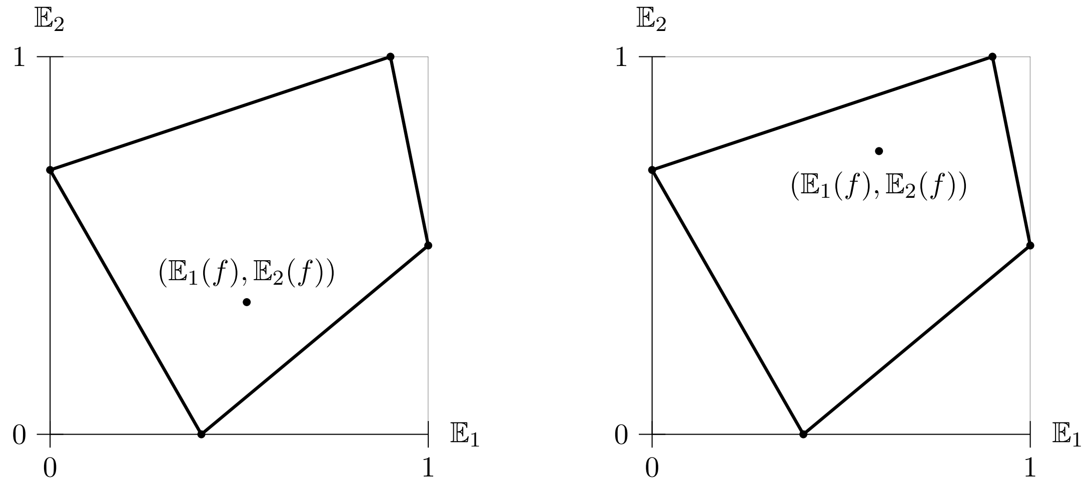

Figure 1: Illustration of independence of redundant acts. Given the preferences of two agents, each act corresponds to a point in $\mathbb{R}^2$ with the horizontal (vertical) coordinate equal to the expected utility of the first (second) agent. Assume the image of all acts under this (in general non-injective) map is the indicated quadrangle on the left when normalizing the utilities to the unit interval. A set of acts $\mathcal{A}'$ is co-redundant if its image is the entire quadrangle since then each act is unanimously indifferent to an act in $\mathcal{A}'$. If the agents change their preferences over acts that are not in $\mathcal{A}'$ so that still every act gets mapped to a point in the quadrangle, then $\mathcal{A}'$ is also co-redundant for the new preference profile. Independence of redundant acts implies that the preferences of society over $\mathcal{A}'$ must not change. Such a preference change is shown on the right for an act f that is not in $\mathcal{A}'$.

right. Without anonymity, the axioms characterize the class of social welfare functions that are weighted belief-averaging and weighted utilitarian: assign two positive (and possibly different) weights to every agent—one for the belief and one for the utility function—and derive the preferences of any society from the weighted mean of its members' beliefs and the weighted sum of their utility functions. Second, additionally dropping independence of redundant acts allows the weights of an agent to depend on the agent's own preferences. More precisely, the weights of every agent can now be arbitrary continuous and positive functions of its preferences. However, the weights of an agent cannot depend on the other agents' preferences. Section 4 discusses the necessity of the axioms for all three results. Section 5 gives an overview of related work. The proofs are in the appendix.

1. Preferences and Social Welfare Functions

Let Ω be a set of **states** of the world and $\mathcal{E}$ be a sigma-algebra over Ω. We refer to elements of $\mathcal{E}$ as events. A probability measure π on $(\Omega, \mathcal{E})$ is **non-atomic** if for every $E \in \mathcal{E}$ with $\pi(E) > 0$, there is $F \subset E$, $F \in \mathcal{E}$, with $0 < \pi(F) < \pi(E)$. A **belief** is a non-atomic and countably additive probability measure; Π denotes the set of all beliefs. Let O be a set of **outcomes** endowed with a sigma-algebra. We assume throughout that $|O| \geq 4$. An **act** is a measurable function $f\colon \Omega \to O$ that maps states to outcomes; $\mathcal{A}$ is the set of all acts. A subset of acts $\mathcal{A}'$ is called **regular** if $\mathcal{A}' = \mathcal{A}(\mathcal{E}', O') = \{f \in \mathcal{A}\colon f \text{ is } \mathcal{E}'\text{-measurable and } f(\Omega) \subset O'\}$ for some sub-sigma-algebra $\mathcal{E}' \subset \mathcal{E}$ and subset of outcomes $O' \subset O$. A **utility function** is a Borel-measurable and bounded function $u\colon O \to \mathbb{R}$. We denote by $\mathcal{U}$ the set of all utility functions that are normalized to the unit interval, that is, $\inf\{u(x)\colon x \in O\} = 0$ and $\sup\{u(x)\colon x \in O\} = 1$; $\bar{\mathcal{U}}$ consists of $\mathcal{U}$ plus the utility function that is constant equal to 0.

A **preference relation** $\succcurlyeq \subset \mathcal{A} \times \mathcal{A}$ is a binary relation over acts. The strict and symmetric part of $\succcurlyeq$ are $\succ$ and $\sim$, respectively. We say that $\succcurlyeq$ maximizes **subjective expected utility** if there is a belief π and a utility function u so that

$$f \succcurlyeq g \text{ if and only if } \int_\Omega (u \circ f) d\pi \geq \int_\Omega (u \circ g) d\pi$$

for all acts f, g. In that case, π and u represent $\succcurlyeq$. We denote by $\bar{\mathcal{R}}$ the set of all preference relations that maximize subjective expected utility; $\mathcal{R}$ consists of $\bar{\mathcal{R}}$ minus complete indifference. For every preference relation in $\mathcal{R}$, the belief in a representation is unique. The utility function is unique up to positive affine transformations. Hence, for every $\succcurlyeq \in \mathcal{R}$, there is a unique belief $\pi \in \Pi$ and a unique utility function $u \in \mathcal{U}$ that represent $\succcurlyeq$. Complete indifference $\succcurlyeq \in \bar{\mathcal{R}} - \mathcal{R}$ can be represented by an arbitrary belief π and the utility function $u \in \bar{\mathcal{U}} - \mathcal{U}$ that is constant equal to 0. In either case, we write $\mathbb{E}_{\succcurlyeq}(f) = \int_\Omega (u \circ f) d\pi$ for the expected utility of an act f under π and u (which does not depend on the choice of π if $\succcurlyeq$ is complete indifference).

We postulate an infinite set of potential agents N. A **society** I consists of a non-empty and finite subset of agents; the collection of all societies is $\mathcal{I}$. Symbols in bold face refer to tuples indexed by a set of agents. Every agent has a preference relation $\succcurlyeq_i \in \mathcal{R}$. (Notice that no agent may be completely indifferent.) Denote by $\pi_i \in \Pi$ and $u_i \in U$ the unique belief and utility function representing $\succcurlyeq_i$. A **preference profile** $\boldsymbol{\succcurlyeq} \in \mathcal{R}^I$ for agents in I specifies the preferences of each agent in I. For $i \in N - I$ and $\succcurlyeq_i \in \mathcal{R}$, we obtain a preference profile $\boldsymbol{\succcurlyeq}_{+i}$ for the agents in $I \cup \{i\}$ by adding $\succcurlyeq_i$ to $\boldsymbol{\succcurlyeq}$. Similarly, when $|I| > 1$ and $i \in I$, $\boldsymbol{\succcurlyeq}_{-i}$ is the profile where $\succcurlyeq_i$ is deleted.

A **social welfare function** (SWF) $\Phi\colon \bigcup_{I \in \mathcal{I}} \mathcal{R}^I \to \bar{\mathcal{R}}$ maps every preference profile for every society to an element of $\bar{\mathcal{R}}$. Φ is belief-averaging and relative utilitarianism (BARU) if for all $I \in \mathcal{I}$ and $\boldsymbol{\succcurlyeq} \in \mathcal{R}^I$, $\Phi(\boldsymbol{\succcurlyeq})$ is represented by

$$\frac{1}{|I|} \sum_{i \in I} \pi_i \qquad \text{and} \qquad \sum_{i \in I} u_i \qquad \text{(belief-averaging and relative utilitarianism)}$$

The table in Appendix D provides an overview of the notation, which is intended as a reference for the reader.

2. Axioms for Social Welfare Functions

We introduce six axioms for SWFs. The first two, restricted monotonicity and independence of redundant acts, carry the most power in the sense that they rule out other SWFs that have been proposed.

The restricted monotonicity axiom applies if a society I is indifferent between two acts f and g and is joined by an agent i for whose belief either of f and g induces the same probability distribution over outcomes as for the belief of the society. In that case, the augmented society $I \cup \{i\}$ should prefer f to g if and only if i does. Formally, for all $I \in \mathcal{I}$, $i \in N - I$, $\succcurlyeq \in \mathcal{R}^I$, and $\succcurlyeq_i \in \mathcal{R}$ with $\succcurlyeq = \Phi(\succcurlyeq)$ and $\succcurlyeq_{+i} = \Phi(\succcurlyeq_{+i})$,

$$f \sim g,\ f \succcurlyeq_i g,\ \pi \circ f^{-1} = \pi_i \circ f^{-1}, \text{ and } \pi \circ g^{-1} = \pi_i \circ g^{-1} \text{ implies } f \succcurlyeq_{+i} g \qquad \text{(restricted monotonicity)}$$

Moreover, a strict preference between f and g for agent i implies a strict preference for the society $I \cup \{i\}$.

Independence of redundant acts prescribes that if two profiles agree on a set of acts that makes every other act redundant, then the corresponding preferences of society over that set also agree. We say that a regular set of acts $\mathcal{A}' = \mathcal{A}(\mathcal{E}', O')$ is **co-redundant** for a profile $\succcurlyeq \in \mathcal{R}^I$ if

(i) every act is unanimously indifferent to some act in $\mathcal{A}'$, that is, for every $f \in \mathcal{A}$, there is $g \in \mathcal{A}'$ such that $f \sim_i g$ for all $i \in I$ and

(ii) for all $i \in I$, $\succcurlyeq_i|_{\mathcal{A}'}$ maximizes expected utility for a belief that is non-atomic on $\mathcal{E}'$[5]

Then, Φ satisfies independence of redundant acts if for all $I \in \mathcal{I}$, $\succcurlyeq, \succcurlyeq' \in \mathcal{R}^I$, and all $\mathcal{A}' \subset \mathcal{A}$ that are co-redundant for $\succcurlyeq$ and $\succcurlyeq'$,

$$\succcurlyeq|_{\mathcal{A}'} = \succcurlyeq'|_{\mathcal{A}'} \text{ implies } \Phi(\succcurlyeq)|_{\mathcal{A}'} = \Phi(\succcurlyeq')|_{\mathcal{A}'} \qquad \text{(independence of redundant acts)}$$

By part (i) of co-redundancy, the image of $\mathcal{A}'$ in utility space is the same as that of $\mathcal{A}$ for both $\succcurlyeq$ and $\succcurlyeq'$. That is, $\{(\mathbb{E}_{\succcurlyeq_i}(g))_{i\in I} \colon g \in \mathcal{A}'\} = \{(\mathbb{E}_{\succcurlyeq_i}(f))_{i\in I} \colon f \in \mathcal{A}\} \subset [0,1]^I$ and similarly for $\succcurlyeq'$.

The remaining four axioms are mostly standard. Faithfulness requires that the preferences of a single-agent society are those of its sole member. For all $i \in N$ and $\succcurlyeq_i \in \mathcal{R}$,

$$\Phi(\succcurlyeq_i) = \succcurlyeq_i \qquad \text{(faithfulness)}$$

The axioms above allow for SWFs that ignore the beliefs of some of the agents. To rule this out, it suffices to assume that no agent can impose its belief on a society. That is, the belief of a society is not identical to that of one of its members unless the rest of the society would arrive at that belief anyway. Formally, for all $I \in \mathcal{I}$ with $|I| > 1$, $i \in I$, and $\succcurlyeq \in \mathcal{R}^I$ where $\Phi(\succcurlyeq)$ and $\Phi(\succcurlyeq_{-i})$ are not complete indifference and π and π_{-i} are the beliefs associated with $\Phi(\succcurlyeq)$ and $\Phi(\succcurlyeq_{-i})$,

$$\pi_{-i} \neq \pi_i \text{ implies } \pi \neq \pi_i \qquad \text{(no belief imposition)}$$

[5]Part (ii) of co-redundancy requires that there exists a non-atomic probability measure π_i' on $(\Omega, \mathcal{E}')$ and a utility function $u_i' \colon O' \to \mathbb{R}$ such that π_i' and u_i' represent $\succcurlyeq_i|_{\mathcal{A}'}$. This ensures that the sigma-algebra $\mathcal{E}'$ is not too small.

Continuity requires that small changes in the agents' preferences can only lead to small changes in the preferences of society. To make this precise, we equip $\mathcal{R}$ and $\bar{\mathcal{R}}$ with topologies. The uniform metric $\sup\{|\,\mathbb{E}_{\succcurlyeq}(f) - \mathbb{E}_{\succcurlyeq'}(f)|\colon f \in \mathcal{A}\}$ induces a topology on $\mathcal{R}$. The set of profiles $\mathcal{R}^I$ has the product topology of $\mathcal{R}$. The topology on $\bar{\mathcal{R}}$ is that of $\mathcal{R}$ plus the entire set $\bar{\mathcal{R}}$ (which is thus the only neighborhood of the relation expressing complete indifference). Thus, the closure of $\mathcal{R}$ is $\bar{\mathcal{R}}$.

$$\Phi \text{ is continuous} \qquad \text{(continuity)}$$

Lastly, anonymity prescribes that relabeling the agents within a society does not change the society's preferences. For all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$,

$$\Phi(\succcurlyeq) = \Phi(\succcurlyeq \circ\, \eta) \text{ for all permutations } \eta \text{ on } I \qquad \text{(anonymity)}$$

Notice that anonymity as defined here is in general weaker than allowing η to be a bijection between two societies I and J of the same size.

The results in Section 3 remain valid if we require all axioms except restricted monotonicity (and, of course, faithfulness) to hold only for societies of size 2. Independence of redundant acts, continuity, and anonymity are used only for profiles with two agents. The assumption that Φ does not allow belief imposition is used for larger profiles, but it is not hard to check that this could be avoided.

3. Characterization of Belief-Averaging and Relative Utilitarianism

Conceptually, our main result is a characterization of BARU. We will obtain it as a corollary of Theorem 1, which uses the first five axioms (thus excluding anonymity) to characterize *weighted belief-averaging and weighted utilitarian* SWFs. These functions assign two positive weights to every agent, one for their belief and one for their utility function, and determine the preferences of every society from the weighted average of the agents' beliefs and the weighted sum of their utility functions. Crucially, the weights are constant across all profiles, that is, they cannot depend on an agent's own preferences, the preferences of any other agent, or the identities of the agents in the society. This is a *much* stronger conclusion than the linear aggregation of Gilboa et al. (2004), which allows the weights to vary arbitrarily across profiles. For example, if the agents' utility functions span the space of all utility functions, any utility function can be written as a linear combination of the agents' utility functions and linear aggregation of utilities holds trivially. This situation is generic for profiles with at least as many agents as outcomes. The difference between Gilboa et al.'s result and Theorem 1 is analogous to the difference between Harsanyi's social aggregation theorem and Dhillon and Mertens's characterization of relative utilitarianism. The cost one incurs is of course a stronger set of axioms.

Theorem 1. *For a social welfare function Φ, the following are equivalent.*

(i) Φ satisfies restricted monotonicity, independence of redundant acts, faithfulness, no belief imposition, and continuity

(ii) *There are* $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}^N_{++}$ *such that for all* $I \in \mathcal{I}$ *and* $\succcurlyeq \in \mathcal{R}^I$, $\Phi(\succcurlyeq)$ *is represented by* $\frac{1}{\sum_{i\in I} v_i} \sum_{i\in I} v_i\pi_i$ *and* $\sum_{i\in I} w_i u_i$

The weights $\boldsymbol{v}$ and $\boldsymbol{w}$ are unique up to multiplication by a positive constant. If Φ is also anonymous, it follows at once that the weights of all agents have to be equal. To see this, consider the two-agent society $I = \{i, j\}$ and any profile $\succcurlyeq \in \mathcal{R}^I$ so that $\pi_i \neq \pi_j$ and $u_i \neq u_j$. Anonymity requires that if the agents swap their preferences, society's preferences do not change. Equivalently, if they swap their beliefs and utility functions, the belief and utility function of society must not change. Hence, $v_i\pi_i + v_j\pi_j = v_i\pi_j + v_j\pi_i$ and $w_iu_i + w_ju_j = w_iu_j + w_ju_i$. These equalities can hold only if $v_i = v_j$ and $w_i = w_j$. Since multiplication of all weights by the same positive constant does not change the preferences of society, we may assume that all weights are equal to 1. This gives a characterization of BARU as the only SWF that satisfies all our axioms. It derives the preferences of society by averaging the beliefs and adding up the normalized utility functions of all agents.

Corollary 1. *For a social welfare function* Φ*, the following are equivalent.*

(i) Φ *satisfies restricted monotonicity, independence of redundant acts, faithfulness, no belief imposition, continuity, and anonymity*

(ii) Φ *is* BARU

The proof of Theorem 1 proceeds as follows. First, we only consider the implications of restricted monotonicity, faithfulness, no belief imposition, and continuity. The former two imply the restricted Pareto condition of Gilboa et al. (2004) and thus that beliefs and utility functions are aggregated linearly (with positive weights for utility functions by the strict part of restricted monotonicity). The additional strength of restricted monotonicity lies in the fact that if an agent joins a society, the belief of the augmented society is an affine combination of the belief of the original society and that of the agent. Assuming all beliefs are affinely independent, it follows that the relative weights of the agents in the original society cannot change. Thus, their relative weights are independent of the preferences of the agent who joins the society. The analogous statement holds for utility functions.

Now, given any profile, we can apply this conclusion to every agent and the subprofile excluding this agent. Some algebra shows that the magnitude of the weight for the belief and the utility function of an agent (relative to that of other agents) can only depend on the agent's own preferences. The signs may however depend on the preferences of the other agents. Since the weights for utility functions have to be positive, any dependence on other agents' preferences vanishes for the weights of utility functions. Continuity allows us to get the same conclusion for beliefs. If the weight for an agent's belief ever were to change sign, by continuity, it would have to be zero in some (two-agent) profile. But then the second agent would get to impose its belief, which is ruled out. We conclude that the weight for an agent's belief cannot change sign. If it were to be negative regardless of the agent's preferences, we could find a profile where the belief of society assigns a negative probability to some event, which gives a contradiction.

In summary, we derive the following intermediate result, which is interesting in its own right.

Proposition 1. *For a social welfare function Φ, the following are equivalent.*

(i) Φ satisfies restricted monotonicity, faithfulness, no belief imposition, and continuity

(ii) There are continuous functions $\boldsymbol{\nu}, \boldsymbol{\omega} \colon \mathcal{R} \to \mathbb{R}^N_{++}$ such that for all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$, $\Phi(\succcurlyeq)$ is represented by $\frac{1}{\sum_{i\in I} \nu_i(\succcurlyeq_i)} \sum_{i\in I} \nu_i(\succcurlyeq_i)\pi_i$ and $\sum_{i\in I} \omega_i(\succcurlyeq_i)u_i$

The last step is to show that for the SWFs characterized in Proposition 1, independence of redundant acts implies that the weights of an agent have to be constant, that is, that the functions ν_i and ω_i are constant. By applying independence of redundant acts to a suitable two-agent profile, one can show that the weights of an agent cannot depend on its belief. To conclude that they are independent of the utility function as well, we consider a two-agent profile in which every act is unanimously indifferent to an act with a range of only three outcomes $\{x_0, x_1, x^*\}$. Hence, the set of acts with range $\{x_0, x_1, x^*\}$ is co-redundant. If the focal agent changes its utility for other outcomes in a way that does not change the image of the profile in utility space, this set remains co-redundant and we can apply independence of redundant acts to conclude that the preferences of society over acts in the co-redundant set do not change. By construction of the profile, this can hold only if the weights of the focal agent remain the same. Lastly, we construct a path between any pair of utility functions so that neighboring utility functions result in the same weights by the preceding conclusion.

4. Necessity of the Axioms

We discuss the necessity of the axioms for Corollary 1 first since the same examples will also work for Theorem 1 and Proposition 1.

To see that restricted monotonicity cannot even be weakened to the restricted Pareto condition, consider the following example adapted from Dhillon and Mertens (1999). For any profile, the belief of society is the average of the agents' beliefs. Consider the closure of the image of the profile in utility space, that is, the closure of $\{(\mathbb{E}_{\succcurlyeq_i}(f))_{i\in I} \colon f \in \mathcal{A}\} \subset [0,1]^I$, and let $(u^i)_{i\in I} \in \mathbb{R}^I$ be the unique point that maximizes the product of utilities. Let the utility function of society be the linear combination of the agents' utility functions where the weight of agent i is $\Pi_{j\in I\setminus\{i\}} u^j$. It is not hard to see that this weight can depend on the utility functions of all agents. Hence, Proposition 1 does not hold if restricted monotonicity is weakened to the restricted Pareto condition. We get the same conclusion for Corollary 1 by observing that the function above satisfies independence of redundant acts since the weights for an agent are the same in any two profiles with the same image in utility space.

The class of functions that satisfy all axioms but independence of redundant acts and anonymity is characterized by Proposition 1. Anonymity holds if and only if the weight functions ν_i and ω_i are the same for all agents.

Without faithfulness, we could have a "phantom agent" whose belief and utility function are always added on top of the beliefs and utility functions of actual agents with a constant weight.

If Φ is not continuous, agents with the same preferences could be handled specially. For example, let (α_n) be a strictly increasing sequence of positive numbers. Then if $\Phi(\boldsymbol{\succcurlyeq})$ is represented by $\frac{1}{\sum_{\succcurlyeq \in \mathcal{R}} \alpha_{n(\succcurlyeq)}} \sum_{\succcurlyeq \in \mathcal{R}} \alpha_{n(\succcurlyeq)} \pi_i$ and $\sum_{\succcurlyeq \in \mathcal{R}} \alpha_{n(\succcurlyeq)} u_i$, where $n(\succcurlyeq)$ is the number of agents in the profile $\boldsymbol{\succcurlyeq}$ with preferences $\succcurlyeq$, then Φ satisfies all axioms but continuity.

It is open whether no belief imposition is necessary for the conclusion of Corollary 1. In the absence of anonymity, that is, for Theorem 1, it is necessary, however. Here, we lose the decomposable form for the belief of society if Φ allows belief imposition. For example, we could have that the belief of society I is π_1 whenever $1 \in I$ and $\frac{1}{|I|} \sum_{i \in I} \pi_i$ otherwise. Thus, whether the belief of an agent gets non-zero weight could depend on whether some particular other agent is present.

For Proposition 1, continuity is even more vital than the above example suggests. Without it, the weight of an agent's belief can be negative. Consider $\Omega = [0, 1]$ equipped with the Borel sigma-algebra $\mathcal{B}([0, 1])$. Let $\tilde{\pi}$ be the uniform distribution on Ω and, for a non-atomic measure π on Ω, let $\rho(\pi) = \sup\{\frac{\pi(E)}{\tilde{\pi}(E)} \colon E \in \mathcal{B}([0, 1])\}$. (Since non-atomic measures have continuous density functions, $\rho(\pi)$ is finite.) For $i \geq 2$, let $\nu_i(\succcurlyeq_i) = \frac{1}{3^i \rho(\pi_i)}$, and ν_1 as well as all ω_i be constant at 1. If $\Phi(\boldsymbol{\succcurlyeq})$ is represented as in Proposition 1 except that the belief of society is $\pi_1 - \sum_{i \in I - \{1\}} \nu_i(\succcurlyeq_i) \pi_i$ (suitably scaled) whenever $1 \in I$ and $\pi_1 = \tilde{\pi}$, then Φ satisfies all axioms but continuity.

Our proof requires $|O| \geq 4$ since it relies on profiles with three linearly independent utility functions. It is open if the results hold when $|O| = 3$.

5. Related Literature on Preference Aggregation Under Uncertainty

Most closely related to the present paper are three multi-profile results for decision-making under uncertainty. Sprumont (2019) studies preference aggregation when agents are subjective expected utility maximizers but societies need not be. He characterizes *ex ante* relative utilitarianism, which ranks acts according to the sum of the agents' expected utilities (when normalizing their utility functions to the unit interval). That is, according to $\sum_{i \in I} \mathbb{E}_{\succcurlyeq_i}(f) = \sum_{i \in I} \int_\Omega (u_i \circ f) d\pi_i$. By contrast, BARU derives a belief of society and ranks acts by the sum of the agents' expected utilities under the belief of society instead of their own belief. Hence, it yields the ranking induced by $\int_\Omega \left(\sum_{i \in I} u_i \circ f\right) d\pi$, where $\pi = \frac{1}{|I|} \sum_{i \in I} \pi_i$. The difference is whether the expectation is taken with respect to an agent's own belief or the aggregate belief. Sprumont assumes the full Pareto condition, independence of inessential expansions (a strengthening of independence of redundant acts),[6] belief irrelevance (the ranking

[6] Independence of inessential expansions requires that if two profiles agree on a set of acts $\mathcal{A}'$ so that $\mathcal{A}'$ contains a most-preferred and a least-preferred act for every agent, then the preferences of society over acts in $\mathcal{A}'$ are the same for both profiles. It is stronger than independence of redundant acts since it also applies sets of acts that are not co-redundant.

of constant acts is independent of the beliefs), and that the preferences of society satisfy Savage's sure-thing principle and depend continuously on the agents' preferences. As a high-level summary, one could say that the assumptions about the SWF are stronger whereas those about the preferences of society are weaker.

Dietrich (2019) assumes that agents as well as societies are subjective expected utility maximizers as does the present paper (albeit that he works in the framework of Anscombe and Aumann (1963) instead of that of Savage). He requires that preference aggregation is consistent with Bayesian updating. That is, preference aggregation and Bayesian updating commute. This asserts that the preferences of society do not depend on whether new information arrives before or after the aggregation. Dietrich shows that consistency with Bayesian updating, continuity, and a weakening of the restricted Pareto condition imply geometric-linear aggregation: society's belief is a *geometric* mean of the agents' beliefs and society's utility function is a linear combination of the agents' utility functions. The weight of an agent in either of these combinations can depend on the profile of utility functions, but not on the beliefs.

Early work on collective decision-making under uncertainty by Hylland and Zeckhauser (1979) studies functions that output a collective *choice* instead of a preference relation of society. Again, all preferences are subjective expected utility-maximizing. They show that no social choice function satisfies the following properties: acts in the choice set are expected utility-maximizing for a belief and utility function of society, this belief (utility function) of society depends only on the agents' beliefs (utility functions), Pareto-dominated acts are never chosen, and no agent can dictate the belief.

Much of the literature on single-profile results is inspired by Mongin (1995) and Gilboa et al. (2004) and focuses on the Pareto condition. Recall that Mongin proved that the Pareto condition can hold only if either the preferences of society coincide with those of one of the agents or all agents have the same belief. In subsequent work, Mongin (1998) showed that this result persists in Anscombe and Aumann's model of subjective expected utility as well as if utilities may be state-dependent as long as the preferences identify a unique belief. Gilboa et al. (2014) consider no-betting-Pareto dominance, which is in between the full and the restricted Pareto condition. It states that a society should prefer one act to another if Pareto dominance in the usual sense holds and there is a belief such that if all agents held that belief, they would also unanimously prefer the first act to the second. They argue that no-betting-Pareto dominance characterizes situations in which agents can benefit from trade, but do not seek to determine how it restricts preference aggregation. Pivato (2020) considers preference aggregation when beliefs are stochastic processes and the state of the world is revealed over time. He shows that linear aggregation of the agents' utility functions is equivalent to eventual Pareto optimality: if all agents eventually prefer one act over another, then society should not eventually have the opposing preference. Alon and Gayer (2016) assume that societies have max-min expected utility preferences (cf. Gilboa and Schmeidler, 1989), where acts are compared based on their minimal expected utility within a set of beliefs. In addition to the restricted Pareto condition, they require a Pareto condition for beliefs: if all agents believe that one event is more likely than another, then so does the society. These two axioms imply that the utility functions are aggregated linearly and the set of beliefs of society are convex combinations of the agents' beliefs. We recommend the discussion of Mongin and Pivato (2016) of different Pareto conditions for more details.

6. Discussion

We conclude the paper with three remarks.

Interpretation of beliefs Two of the axioms are not formulated purely on the level of preferences: restricted monotonicity and no belief imposition reference the belief in the expected utility representation of a preference relation. While it is possible to state these axioms in terms of preferences only, such a formulation would be clumsy and obscure the intuitions. So for these axioms to be meaningful, it is not enough that the preferences can be *represented* by expected utility maximization (that is, satisfy Savage's postulates 1 through 7). The representing beliefs and utility functions also need to be the decision maker's *rationale* for the preference relation.

One prerequisite for the latter is that the utility functions are state-independent as assumed in Savage's framework. Whether state-independent utilities are a tenable assumption depends on the context. State-independence requires that the decision-maker is only affected by the outcomes and not by the state. This is true in good approximation for, say, investment decisions where the decision-maker is only interested in the realized asset value and not the causes for the outcome. By contrast, if the decision-maker cares about the causes for an outcome, as, for example, when the decision-maker's health depends on the state, state-independence may fail. We refer to Drèze and Rustichini (1999) for a critical discussion of state-(in)dependent utilities. If utilities were allowed to be state-dependent, the belief in the expected utility representation would not be unique, thus making it ambiguous how to interpret an axiom that references a single belief. To regain the uniqueness of beliefs, one needs additional structural assumptions on the preferences (see, e.g., Karni et al., 1983).

A more general critique is that beliefs and utility functions are nothing but mathematical objects used to concisely represent preferences. This is a general critique of subjective expected utility maximization and not specific to this paper. Gilboa et al. (2014, Section 1.3) and Gayer et al. (2014, Section 5.1) address the subject is some detail and we will not engage in it further.

Properties of belief-averaging and relative utilitarianism One can ask which axiomatic properties BARU satisfies beyond those assumed in Corollary 1 and which ones it violates. As mentioned, restricted monotonicity and faithfulness together imply the restricted Pareto condition of Gilboa et al. (2004). It is also clear from the impossibility result of Mongin (1997) that BARU does not satisfy the full Pareto condition (since for almost all preference profiles, society's preferences are not identical to the preferences of one of the agents).

It is less obvious that the axioms imply that the beliefs and the utility functions are aggregated separately. That is, society's belief does not depend on the agents' utility functions and likewise for society's utility function. The social welfare functions characterized in Proposition 1 do not, in general, have this property. So the separability relies on the independence of redundant acts axiom. Similarly, BARU is neutral in the sense that it does not distinguish between the outcomes. A priori, it could not be ruled out that, say, a larger weight is given to agents who rank acts that yield certain outcomes at the top of their preferences. Independence of redundant acts again plays a key role in establishing neutrality.

Alternative interpretations of the model We have interpreted the formal framework as modeling collective decisions under uncertainty. Here are two alternative interpretations.

The first is the aggregation of expert judgments. Consider a decision-maker who has to decide between various acts and seeks the advice of multiple experts. Each expert provides a belief and a utility function, which specify the expert's judgment about the probabilities of the states and the desirability of the outcomes (for the decision-maker). The task of the decision-maker is to form preferences by aggregating the experts' judgments. This interpretation has been suggested by Raiffa (1968), who seems to be in favor of belief-averaging and (relative) utilitarianism. The translation between collective decision-making and aggregation of expert judgements is as follows.

$$\begin{aligned} \text{agents} &\longleftrightarrow \text{experts} \\ \text{society} &\longleftrightarrow \text{decision-maker} \end{aligned}$$

The second interpretation is the aggregation of inter-temporal preferences. Suppose there is a continuous time axis and at each point in time, some outcome is realized. Moreover, there are multiple agents, each of which has preferences over the outcome streams ranking them by their discounted utility for some discounting function and some utility function on the outcomes. The agents' preferences over the outcome streams are then to be aggregated into a collective preference. For example, the members of a household may try to decide on a household consumption plan over some (possibly indefinite) time horizon. Jackson and Yariv (2014) discuss further examples and employ a similar model.

This situation fits our framework as follows. Each point in time corresponds to a state. Hence, outcome streams are identified with acts and discounting functions with beliefs.

$$\begin{aligned} \text{states} &\longleftrightarrow \text{points in time} \\ \text{acts} &\longleftrightarrow \text{outcome streams} \\ \text{beliefs} &\longleftrightarrow \text{discounting functions} \end{aligned}$$

We point out that the analogue of averaging beliefs is averaging *normalized* discounting functions. Since beliefs are probability measures, the corresponding normalization is to scale discounting functions so that their integral over time is 1 (so discounting functions have to be in

L^1).[7]

There is a vast literature on either of the two interpretations and no attempt is made to cover those here. Whether our axioms should hold and BARU is suited for the aggregation task is, in either case, a separate issue. Generally speaking, aggregating expert judgments seems closer to our default interpretation than aggregating inter-temporal preferences. For example, there is an inherent ordering on points in time but none on an abstract state space. The purpose of this discussion to merely to highlight that the presented framework is general enough to model a range of applications.

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 5969/1-1. The author thanks Jean Baccelli, Francesc Dilmé, David Easley, Loren Fryxell, Johannes Hörner, Ian Jewitt, Michel Le Breton, Wolfgang Pesendorfer, Dominik Peters, Evgenii Safonov, and Omer Tamuz for helpful comments. A previous version of this paper has been presented in the Public Economics and Microeconomic Theory Seminar at Cornell University and the Microeconomic Theory Seminar at Princeton University.

References

S. Alon and G. Gayer. Utilitarian preferences with multiple priors. *Econometrica*, 84(3):1181–1201, 2016.

F. J. Anscombe and R. J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205, 1963.

K. J. Arrow. *Social Choice and Individual Values.* New Haven: Cowles Foundation, 1st edition, 1951. 2nd edition 1963.

T. Börgers and Y.-M. Choo. Revealed relative utilitarianism. CESifo Working Paper No. 6613, 2017.

A. Dhillon. Extended Pareto rules and relative utilitarianism. *Social Choice and Welfare*, 15 (4):521–542, 1998.

A. Dhillon and J.-F. Mertens. Relative utilitarianism. *Econometrica*, 67(3):471–498, 1999.

F. Dietrich. Fully Bayesian aggregation. 2019. Working paper.

J. H. Drèze and A. Rustichini. State-dependent utility and decision theory. In S. Barberà, P. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume 2, pages 839–892. Springer, Boston, MA, 1999.

[7]This is also the normalization we would have derived for beliefs had we allowed them to be finite positive measures instead of probability measures.

G. Gayer, I. Gilboa, L. Samuelson, and D. Schmeidler. Pareto efficiency with different beliefs. *The Journal of Legal Studies*, 43(2):151–171, 2014.

I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

I. Gilboa, D. Samet, and D. Schmeidler. Utilitarian aggregation of beliefs and tastes. *Journal of Political Economy*, 112(4):932–938, 2004.

I. Gilboa, L. Samuelson, and D. Schmeidler. No-betting-Pareto dominance. *Econometrica*, 82 (4):1405–1442, 2014.

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

A. Hylland and R. Zeckhauser. The impossibility of Bayesian group decision making with separate aggregation of beliefs and values. *Econometrica*, 47(6):1321–1336, 1979.

M. O. Jackson and L. Yariv. Present bias and collective dynamic choice in the lab. *American Economic Review*, 104(12):4184–4204, 2014.

E. Karni. Impartiality: Definition and representation. *Econometrica*, 66(6):1405–1415, 1998.

E. Karni, D. Schmeidler, and K. Vind. On state dependent preferences and subjective probabilities. *Econometrica*, 51(4):1021–1031, 1983.

A. Liapounoff. Sur les fonctions-vecteurs complètement additives. *Bulletin of the American Mathematical Solciety URSS. Sér. Math. [Izv. Akad. Nauk SSSR]*, 10(3):465–478, 1940.

P. Mongin. Consistent Bayesian aggregation. *Journal of Economic Theory*, 66(2):313–351, 1995.

P. Mongin. Spurious unanimity and the Pareto principle. THEMA, Université de Cergy-Pontoise, 1997.

P. Mongin. The paradox of the Bayesian experts and state-dependent utility theory. *Journal of Mathematical Economics*, 29(3):331–361, 1998.

P. Mongin and M. Pivato. Social evaluation under risk and uncertainty. In M. D. Adler and M. Fleurbaey, editors, *The Oxford Handbook of Well-Being and Public Policy*. Oxford University Press, 2016.

M. Pivato. Bayesian social aggregation with accumulating evidence. 2020.

H. Raiffa. *Decision analysis; introductory lectures on choices under uncertainty*. Addison-Wesley, 1968.

L. J. Savage. *The Foundations of Statistics*. Wiley Publications in Statistics. Wiley and Sons, 1954.

U. Segal. Let's agree that all dictatorships are equally bad. *Journal of Political Economy*, 108 (3):569–589, 2000.

Y. Sprumont. Relative utilitarianism under uncertainty. *Social Choice and Welfare*, 53:621–639, 2019.

APPENDIX: Proofs

An SWF Φ gives rise to two functions ϕ and ψ, which take a profile $\succcurlyeq \in \mathcal{R}^I$ for an arbitrary society I to a belief $\phi(\succcurlyeq) \in \Pi$ and a utility function $\psi(\succcurlyeq) \in \bar{\mathcal{U}}$ of society. Whenever the utility function of society is trivial, that is, $\psi(\succcurlyeq) = 0$, $\phi(\succcurlyeq)$ is not uniquely determined. In fact, it can be arbitrary.

Preferences are invariant under multiplication of the belief by a positive constant and positive affine transformations of utility functions. For measures π, π' on Ω and utility functions u, u' on O, we write $\pi \equiv \pi'$ if $\pi = \alpha\pi'$ for some $\alpha > 0$ and $u \equiv u'$ if $u = \alpha u' + \beta$ for $\alpha > 0$ and $\beta \in \mathbb{R}$.

Theorem 1 requires that we find $\boldsymbol{v}, \boldsymbol{w} \in R^N_{++}$ such that $\phi(\succcurlyeq) \equiv \sum_{i\in I} v_i\pi_i$ and $\psi(\succcurlyeq) \equiv \sum_{i\in I} w_i u_i$ for all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$. The proof proceeds in three main steps. First, we examine the implications of restricted monotonicity in conjunction with faithfulness and no belief imposition. These axioms imply that in almost all profiles, the weights assigned to an agent's belief and utility function can only depend on the agent's own preferences. The weights for beliefs may be negative however. Second, we add continuity, which allows us to rule out negative weights and to extend the obtained representation to all profiles (Proposition 1). Lastly, independence of redundant acts implies that the weights of an agent cannot depend on the agent's preferences either and thus have to be constant across all profiles.

A. Implications of Restricted Monotonicity

The proofs that the weight of the belief and the weight of the utility function of an agent can only depend on the agent's preferences in Sections A.1 and A.2 proceed along the same lines. For the most part, the proof for beliefs requires more work, since we cannot rule out negative weights. Thus, we advise readers interested in the proofs to take a look at Appendix A.2 first.

A.1. Aggregation of Beliefs

The first lemma is the basis from which we will derive all further conclusions about belief aggregation. It states that if an agent joins a society, the belief of the augmented society is an affine combination of the belief of original society and the belief of the agent. The restricted Pareto condition of Gilboa et al. (2004) already implies that the belief of society is an affine combination of its members' beliefs. The additional strength of this conclusion lies in the fact that no matter which belief the new agent holds, it is always combined with the same belief of the original society. If we assume that an agent cannot impose its belief on the society, the agent's weight in the affine combination cannot be 1.

Lemma 1. *Assume that* Φ *satisfies restricted monotonicity and rules out belief imposition. Let* $I \in \mathcal{I}$, $i \in I$, *and* $\succcurlyeq \in \mathcal{R}^I$ *with* $\psi(\succcurlyeq) \neq 0$. *Then,* $\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-i}) + \alpha\pi_i$ *for some* $\alpha \in \mathbb{R} - \{1\}$.

Proof. Let $\succsim = \Phi(\succcurlyeq)$ and $\succsim_{-i} = \Phi(\succcurlyeq_{-i})$. Restricted monotonicity implies that $f \sim g$ whenever $f \sim_{-i} g$, $f \sim_i g$ and $\phi(\succcurlyeq_{-i}) \circ f^{-1} = \pi_i \circ f^{-1}$ and $\phi(\succcurlyeq_{-i}) \circ g^{-1} = \pi_i \circ g^{-1}$. Thus, (the two-agent case of) Theorem 1 of Gilboa et al. (2004) implies that $\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-i}) + \alpha\pi_i$ for some $\alpha \in \mathbb{R}$. If $\pi_i = \phi(\succcurlyeq_{-i})$, we can choose α arbitrarily. Otherwise, $\pi_i \neq \phi(\succcurlyeq)$, since Φ rules out belief imposition, and so $\alpha \neq 1$. □

Lemma 2. *Assume that* Φ *satisfies restricted monotonicity and faithfulness and rules out belief imposition. Let* $I \in \mathcal{I}$ *and* $\succcurlyeq \in \mathcal{R}^I$ *with* $\psi(\succcurlyeq) \neq 0$. *Then* $\phi(\succcurlyeq) = \sum_{i \in I} v_i \pi_i$ *for some* $\boldsymbol{v} \in \mathbb{R}^I$ *with* $\sum_{i \in I} v_i = 1$. *Moreover, if* $(\pi_i)_{i \in I}$ *are affinely independent, then* $\boldsymbol{v} \in (\mathbb{R} - \{0\})^I$ *and* $\boldsymbol{v}$ *is unique.*

Proof. Since Φ is faithful, we have that $\phi(\succsim_1) = \pi_1$. Now let one agent after another join. We apply Lemma 1 at each step and get $\phi(\succcurlyeq) = \sum_{i \in I} v_i \pi_i$ for some $\boldsymbol{v} \in \mathbb{R}^I$ with $\sum_{i \in I} v_i = 1$.

If $(\pi_i)_{i \in I}$ are affinely independent, $\boldsymbol{v}$ is unique. We prove by induction over $|I|$ that $v_i \neq 0$ for all i. If $|I| = 1$, then $\boldsymbol{v} = 1$ is forced. Now suppose that $|I| > 1$ and let $i, j \in I$. By the induction hypothesis, we have $\phi(\succcurlyeq_{-i}) = \sum_{k \in I - \{i\}} v'_k \pi_k$ and $\phi(\succcurlyeq_{-j}) = \sum_{k \in I - \{j\}} v''_k \pi_k$ for some $\boldsymbol{v}' \in (\mathbb{R} - \{0\})^{I - \{i\}}$ and $\boldsymbol{v}'' \in (\mathbb{R} - \{0\})^{I - \{j\}}$. Lemma 1 implies that

$$\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-i}) + \alpha\pi_i = (1-\beta)\phi(\succcurlyeq_{-j}) + \beta\pi_j$$

for some $\alpha, \beta \in \mathbb{R} - \{1\}$. We set $\boldsymbol{v} = ((1-\alpha)\boldsymbol{v}', \alpha) = ((1-\beta)\boldsymbol{v}'', \beta)$, where α and β appear in position i and j respectively. Since $\alpha \neq 1$ and $v'_k \neq 0$, it follows that $v_k \neq 0$ for all $k \in I - \{i\}$. Similarly, $\beta \neq 1$ implies that $v_k \neq 0$ for all $k \in I - \{j\}$. □

We define the dimension of a vector of beliefs $\boldsymbol{\pi} \in \Pi^I$ as the maximal number of affinely independent probability distributions in $\{\pi_i \colon i \in I\}$. Equivalently, the dimension of $\boldsymbol{\pi}$ is the dimension of the subset $\{(\pi_i(E))_{i \in I} \colon E \in \mathcal{E}\}$ of $\mathbb{R}^I$. For later use, we prove a fact for $\boldsymbol{\pi}$ with dimension at least 3.

Lemma 3. *Assume that* Φ *satisfies restricted monotonicity and rules out belief imposition. Let* $I \in \mathcal{I}$ *and* $\succcurlyeq \in \mathcal{R}^I$ *with* $\psi(\succcurlyeq) \neq 0$. *If* $\boldsymbol{\pi}$ *has dimension at least 3, there are distinct* $i, j \in I$ *such that* $\phi(\succcurlyeq_{-i,j}), \pi_i$, *and* π_j *are affinely independent.*

Proof. Suppose $\{1, 2, 3\} \subset I$. Since $\boldsymbol{\pi}$ has dimension at least 3, we may assume that π_1, π_2, and π_3 are affinely independent. So $\phi(\succcurlyeq)$ cannot be in the affine hull of all three pairs from $\{\pi_1, \pi_2, \pi_3\}$, for if say $\phi(\succcurlyeq)$ is in the affine hull of $\{\pi_1, \pi_2\}$ and $\{\pi_1, \pi_3\}$, then $\phi(\succcurlyeq) = \pi_1$ and so is not in the affine hull of $\{\pi_2, \pi_3\}$. Assume that $\phi(\succcurlyeq)$ is not in the affine hull of $\{\pi_1, \pi_2\}$. Then

Lemma 1 implies that $\phi(\succcurlyeq_{-1,2})$ is not in the affine hull of $\{\pi_1, \pi_2\}$. Since $\pi_1 \neq \pi_2$, $\phi(\succcurlyeq_{-1,2}), \pi_1$, and π_2 are affinely independent. □

Lemma 2 ensures that the belief of a society is always an affine combination of the agents' beliefs. To show that ϕ has the form claimed in Theorem 1, we have to prove that the relative weight of an agent in this affine combination depends only on the agent's own belief and utility function. For now, we have to be contempt with a weaker conclusion, which allows negative weights. For the rest of this section, we will assume that *beliefs and utility functions are pairwise distinct in all profiles.*

Lemma 4. *Assume that Φ satisfies restricted monotonicity and faithfulness and rules out belief imposition. Then there are $\boldsymbol{\nu}\colon \mathcal{R} \to (\mathbb{R} - \{\mathbf{0}\})^N$ and for all $I \in \mathcal{I}$, $\boldsymbol{\sigma}^I\colon \mathcal{R}^I \to \{-1, 1\}^I$ such that for all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$, $\phi(\succcurlyeq) \equiv \sum_{i\in I} \sigma_i^I(\succcurlyeq)\nu_i(\succcurlyeq_i)\pi_i$. Moreover, $\frac{\sigma_i^I(\succcurlyeq)}{\sigma_j^I(\succcurlyeq)} = \frac{\sigma_i^{\{i,j\}}(\succcurlyeq_i,\succcurlyeq_j)}{\sigma_j^{\{i,j\}}(\succcurlyeq_i,\succcurlyeq_j)}$ for all I, $i, j \in I$, and $\succcurlyeq \in \mathcal{R}^I$.*

Proof. For $l \in \mathbb{R} - \{1, 2\}$, let $I_l = \{1, 2, l\}$ and $\mathcal{R}_l \subset \mathcal{R}^{I_l}$ be the set of all profiles for agents I_l such that π_1, π_2, and π_l are affinely independent. Let $l \in N - \{1, 2\}$ be arbitrary and fix some $\tilde{\succcurlyeq} \in \mathcal{R}_l$. By Lemma 2, there is a unique function $\boldsymbol{\kappa}\colon \mathcal{R}_l \to (\mathbb{R} - \{\mathbf{0}\})^{I_l}$ such that $\phi(\succcurlyeq) = \sum_{i\in I_l} \kappa_i(\succcurlyeq)\pi_i$ for all $\succcurlyeq \in \mathcal{R}_l$. For $i, j \in I_l$ and $\succcurlyeq \in \mathcal{R}_l$, let $\lambda^{i,j}(\succcurlyeq) = \frac{\kappa_j(\succcurlyeq_{-j},\tilde{\succcurlyeq}_j)\kappa_i(\succcurlyeq)}{\kappa_j(\tilde{\succcurlyeq})\kappa_i(\succcurlyeq_{-j},\tilde{\succcurlyeq}_j)}$. The fact that $\boldsymbol{\kappa}$ maps to $(\mathbb{R} - \{\mathbf{0}\})^{I_l}$ ensures that $\lambda^{i,j}$ is well-defined. Then, let $\nu_i(\succcurlyeq_i) = \frac{|\kappa_i(\succcurlyeq)|}{|\lambda^{i,j}(\succcurlyeq)|}$ and $\sigma_i^{I_l}(\succcurlyeq) = \text{sign}(\kappa_i(\succcurlyeq))$.

We show that ν_i is independent of j and $\succcurlyeq_{-i}$ and thus well-defined. We proceed in three steps. Before, note that the projection of $\mathcal{R}_l$ to $\mathcal{R}$ that returns the preferences of i is onto, and so ν_i is a function on all of $\mathcal{R}$.

Step 1. Let $k \in I_l - \{i, j\}$. We show that $\frac{\kappa_i(\succcurlyeq)}{\kappa_j(\succcurlyeq)}$ is independent of $\succcurlyeq_k$. To this end, let $\succcurlyeq' \in \mathcal{R}_l$ such that $\succcurlyeq_{-k}' = \succcurlyeq_{-k}$. By Lemma 1, we have that

$$\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-k}) + \alpha\pi_k = \kappa_i(\succcurlyeq)\pi_i + \kappa_j(\succcurlyeq)\pi_j + \kappa_k(\succcurlyeq)\pi_k, \text{ and}$$
$$\phi(\succcurlyeq') = (1-\beta)\phi(\succcurlyeq_{-k}') + \beta\pi_k' = \kappa_i(\succcurlyeq')\pi_i + \kappa_j(\succcurlyeq')\pi_j + \kappa_k(\succcurlyeq')\pi_k',$$

for some $\alpha, \beta \in \mathbb{R} - \{1\}$. Affine independence of π_1, π_2, π_l and π_1, π_2, π_l' implies that $\kappa_i(\succcurlyeq)\pi_i + \kappa_j(\succcurlyeq)\pi_j \equiv \phi(\succcurlyeq_{-k}) = \phi(\succcurlyeq_{-k}') \equiv \kappa_i(\succcurlyeq')\pi_i + \kappa_j(\succcurlyeq')\pi_j$. In particular, $\frac{\kappa_i(\succcurlyeq)}{\kappa_j(\succcurlyeq)} = \frac{\kappa_i(\succcurlyeq')}{\kappa_j(\succcurlyeq')}$. Repeated application yields the desired independence.

Step 2. We show that $\lambda^{i,j}(\succcurlyeq)$ is independent of i and j. This is tedious, but only uses Step 1.

Let $k \in I_l - \{i, j\}$. First we show independence of j.

$$\begin{aligned}
\lambda^{i,j}(\boldsymbol{\succcurlyeq}) &= \frac{\kappa_j(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)\kappa_i(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)} \\
&= \frac{\kappa_j(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)\kappa_i(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\pi}})\kappa_i(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)} \\
&= \frac{\kappa_j(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)\kappa_i(\boldsymbol{\succcurlyeq})\kappa_k(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)\kappa_k(\boldsymbol{\succcurlyeq}_{-j,k}, \tilde{\succcurlyeq}_k, \tilde{\succcurlyeq}_j)} \\
&= \frac{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_k(\boldsymbol{\succcurlyeq}_{-k}, \tilde{\succcurlyeq}_k)\kappa_i(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq}_{-k}, \tilde{\succcurlyeq}_k)\kappa_k(\tilde{\boldsymbol{\succcurlyeq}})} = \nu^{i,k}(\boldsymbol{\succcurlyeq})
\end{aligned}$$

Verifying independence of i is very similar.

$$\begin{aligned}
\lambda^{i,j}(\boldsymbol{\succcurlyeq}) &= \frac{\kappa_j(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)\kappa_i(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)} \\
&= \frac{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)\kappa_i(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)} \\
&= \frac{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)\kappa_i(\boldsymbol{\succcurlyeq})\kappa_k(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)\kappa_k(\tilde{\boldsymbol{\succcurlyeq}}_{-k,i}, \succcurlyeq_k, \succcurlyeq_i)} \\
&= \frac{\kappa_j(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)\kappa_i(\boldsymbol{\succcurlyeq})\kappa_k(\boldsymbol{\succcurlyeq})}{\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq})\kappa_k(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)} = \nu^{k,j}(\boldsymbol{\succcurlyeq})
\end{aligned}$$

Step 3. We show that $\nu_i(\boldsymbol{\succcurlyeq})$ is independent of $\boldsymbol{\succcurlyeq}_{-i}$ and j.

$$\nu_i(\succcurlyeq_i) = \frac{|\kappa_i(\boldsymbol{\succcurlyeq})|}{|\lambda^{i,j}(\boldsymbol{\succcurlyeq})|} = \frac{|\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)|}{|\kappa_j(\boldsymbol{\succcurlyeq}_{-j}, \tilde{\succcurlyeq}_j)|} = \frac{|\kappa_j(\tilde{\boldsymbol{\succcurlyeq}})\kappa_i(\tilde{\boldsymbol{\succcurlyeq}}_{-i}, \succcurlyeq_i)|}{|\kappa_j(\tilde{\boldsymbol{\succcurlyeq}}_{-i}, \succcurlyeq_i)|},$$

where we use Step 1 for the last equality. The resulting term is independent of $\boldsymbol{\succcurlyeq}_{-i}$ and, by Step 2, of j.

Now it is easy to see that

$$\phi(\boldsymbol{\succcurlyeq}) = \sum_{i \in I_l} \kappa_i(\boldsymbol{\succcurlyeq})\pi_i \equiv \sum_{i \in I_l} \frac{\kappa_i(\boldsymbol{\succcurlyeq})}{|\nu^{i,j_i}(\boldsymbol{\succcurlyeq})|}\pi_i = \sum_{i \in I_l} \sigma_i^{I_l}(\boldsymbol{\succcurlyeq})\nu_i(\succcurlyeq_i)\pi_i,$$

where $j_i \in I_l - \{i\}$ for all i. For the second equality, we used the fact that $\lambda^{i,j}$ is independent of i and j.

Since l was arbitrary, we have now defined ν_i for each $i \in N$. However, we have defined ν_1 and ν_2 multiple times, once for each $l \in N - \{1, 2\}$. So we have to check that these definitions are not conflicting. It follows from Lemma 1 that the ratio between ν_1 and ν_2 is the same for each triple $\{1, 2, l\}$. Thus, we can define ν_1 and ν_2 as obtained for, say, $l = 3$ and scale the triples (ν_1, ν_2, ν_l) obtained for the remaining l appropriately.

Step 4. Now we define $\boldsymbol{\sigma}$ for the remaining profiles. Our strategy will be to first define it for two-agent profiles, then inductively for all profiles such that $\boldsymbol{\pi}$ has dimension at least 3, and then for the remaining profiles. At each point, we maintain that $\frac{\sigma_i^I(\boldsymbol{\succcurlyeq})}{\sigma_j^I(\boldsymbol{\succcurlyeq})} = \frac{\sigma_i^{\{i,j\}}(\succsim_i, \succsim_j)}{\sigma_j^{\{i,j\}}(\succsim_i, \succsim_j)}$, which we will refer to as the ratio condition on $\boldsymbol{\sigma}$. We omit the superscript in expressions like $\sigma_i^I(\boldsymbol{\succcurlyeq})$ from now on, since it is clear from the profile.

Let $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$. If $|I| = 2$, say $I = \{1,2\}$, then $\phi(\succcurlyeq) = \alpha_1\pi_1 + \alpha_2\pi_2$ for a unique $\boldsymbol{\alpha} \in (\mathbb{R} - \{0\})^I$. We define $\sigma_i(\succcurlyeq) = \operatorname{sign}(\alpha_i)$ for $i \in I$.

Now assume that $|I| \geq 3$ and $\boldsymbol{\pi}$ has dimension at least 3; assume further that we have defined $\boldsymbol{\sigma}$ for all profiles of dimension three on fewer than $|I|$ agents such that the ratio condition holds. Let $i \in I$ such that $\pi_i \neq \phi(\succcurlyeq_{-i})$, which exists by Lemma 3. We show that there is $s \in \{-1,1\}$ such that $\frac{s}{\sigma_j(\succcurlyeq_{-i})} = \frac{\sigma_i(\succsim_i,\succsim_j)}{\sigma_j(\succsim_i,\succsim_j)}$ for all $j \in I - \{i\}$. If not, there are j,k which require $s = 1$ and $s = -1$ respectively. It is not hard to see that then there must be j,k with this property such that (π_i,π_j,π_k) has dimension 3. Then we have

$$\frac{\sigma_j(\succcurlyeq_{-i})}{\sigma_k(\succcurlyeq_{-i})} = \frac{\sigma_j(\succsim_j,\succsim_k)}{\sigma_k(\succsim_j,\succsim_k)} \neq \frac{\sigma_j(\succsim_i,\succsim_j)}{\sigma_i(\succsim_i,\succsim_j)}\frac{\sigma_i(\succsim_i,\succsim_k)}{\sigma_k(\succsim_i,\succsim_k)} = \frac{\sigma_j(\succcurlyeq_{i,j,k})}{\sigma_i(\succcurlyeq_{i,j,k})}\frac{\sigma_i(\succcurlyeq_{i,j,k})}{\sigma_k(\succcurlyeq_{i,j,k})} = \frac{\sigma_j(\succcurlyeq_{i,j,k})}{\sigma_k(\succcurlyeq_{i,j,k})}$$

where we use the fact that $\boldsymbol{\sigma}^{I-\{i\}}$ satisfies the ratio condition for the first equality. This is a contradiction, since $\boldsymbol{\sigma}^{\{i,j,k\}}$ also satisfies the ratio condition. Thus we can find s as required.

Lemma 1 implies that $\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-i}) + \alpha\pi_i$ for some unique $\alpha \in \mathbb{R} - \{1\}$. If $\alpha < 1$, we set $\boldsymbol{\sigma}^I = (\boldsymbol{\sigma}^{I-\{i\}}, s)$; if $\alpha > 1$, set $\boldsymbol{\sigma}^I = -(\boldsymbol{\sigma}^{I-\{i\}}, s)$.

Lastly, if $\boldsymbol{\pi}$ has dimension 2, let $i \in N - I$ and consider a profile $\succcurlyeq'$ for agents in $I \cup \{i\}$ such that ${\succcurlyeq_{-i}}' = \succcurlyeq$ and $\boldsymbol{\pi}'$ has dimension 3. By Lemma 2, $\pi_i' \neq \phi(\succcurlyeq)$ and so $\phi(\succcurlyeq') = (1-\alpha)\phi(\succcurlyeq) + \alpha\pi_i'$ for some $\alpha \neq 1$. If $\alpha < 1$, we set $\sigma_j(\succcurlyeq) = \sigma_j(\succcurlyeq')$ for all $j \in I$; if $\alpha > 1$, set $\sigma_j(\succcurlyeq) = -\sigma_j(\succcurlyeq')$.

We still have to make sure that these definitions of $\boldsymbol{\nu}$ and $\boldsymbol{\sigma}$ are consistent with ϕ. For $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$, let $\bar{\phi}(\succcurlyeq) \equiv \sum_{i\in I} \sigma_i^I(\succcurlyeq)\nu_i(\succcurlyeq_i)\pi_i$. We show by induction over $|I|$ that ϕ and $\bar{\phi}$ agree on all profiles $\succcurlyeq$. First we assume that $\boldsymbol{\pi}$ has dimension at least 3. Later we will take care of the remaining profiles.

We start with two observations.

Step 5. Let $\succcurlyeq$ be a profile for agents in I such that $\boldsymbol{\pi}$ has dimension 3 and $\boldsymbol{\pi}_{-i}$ has dimension 2. If ϕ and $\bar{\phi}$ agree on $\succcurlyeq$, then they also agree on $\succcurlyeq_{-i}$. By Lemma 1 and the assumption, we have

$$\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-i}) + \alpha\pi_i \equiv \sum_{j\in I-\{i\}} \sigma_j(\succcurlyeq)\nu_j(\succcurlyeq_j)\pi_j + \sigma_i(\succcurlyeq)\nu_i(\succcurlyeq_i)\pi_i$$

for some $\alpha \in \mathbb{R} - \{1\}$. Since π_i is not in the affine hull of $(\pi_j)_{j\in I-\{i\}}$, we have to have $(1-\alpha)\phi(\succcurlyeq_{-i}) \equiv \sum_{j\in I-\{i\}} \sigma_j(\succcurlyeq)\nu_j(\succcurlyeq_j)\pi_j$. If $\alpha < 1$, then by definition, $\sigma_j(\succcurlyeq_{-i}) = \sigma_j(\succcurlyeq)$ for all $j \in I - \{i\}$, and so $\phi(\succcurlyeq_{-i}) \equiv \sum_{j\in I-\{i\}} \sigma_j(\succcurlyeq_{-i})\nu_j(\succcurlyeq_j)\pi_j \equiv \bar{\phi}(\succcurlyeq_{-i})$. If $\alpha > 1$, then $\sigma_j(\succcurlyeq_{-i}) = -\sigma_j(\succcurlyeq)$ for all j, and again $\phi(\succcurlyeq_{-i}) = \bar{\phi}(\succcurlyeq_{-i})$ follows.

Step 6. Let $\succcurlyeq$ be a profile for agents in I and $i,j \in I$. If ϕ and $\bar{\phi}$ agree on $\succcurlyeq_{-i}$ and $\succcurlyeq_{-j}$ and $\bar{\phi}(\succcurlyeq_{-i,j}), \pi_i$, and π_j are affinely independent, then they also agree on $\succcurlyeq$. Lemma 1 implies that

$$\phi(\succcurlyeq) = (1-\alpha)\phi(\succcurlyeq_{-j}) + \alpha\pi_j = (1-\beta)\phi(\succcurlyeq_{-i}) + \beta\pi_i$$

for some $\alpha, \beta \in \mathbb{R} - \{1\}$. To make notation less cumbersome, we write $\sigma_k(\succcurlyeq_{-J}) = \sigma_k^J$ and $\nu_k(\succcurlyeq_k) = \nu_k$ for $k \in I$ and $J \subset I - \{i\}$ for the rest of this step. Four cases arise, depending on whether α and β are greater of smaller than 1.

Case 1. Assume $\alpha, \beta < 1$. By definition of $\boldsymbol{\sigma}$, we have that $\sigma_k = \sigma_k^j$ for all $k \in I - \{j\}$ and $\sigma_k = \sigma_k^i$ for all $k \in I - \{i\}$. In particular, $\sigma_k^i = \sigma_k^j$ for $k \in I - \{i, j\}$. Moreover, either $\sigma_k^{ij} = \sigma_k^i$ for all $k \in I - \{i, j\}$ or $\sigma_k^{ij} = -\sigma_k^i$ for all k. Let $s = 1$ in the former case and $s = -1$ otherwise. Then,

$$\phi(\succcurlyeq) \equiv \underbrace{s \left(\sum_{k \in I - \{i,j\}} \sigma_k^{ij} \nu_k \pi_k \right)}_{\phi(\succcurlyeq_{-i,j})} + \sigma_i^j \nu_i \pi_i + \alpha' \pi_j = s \left(\sum_{k \in I - \{i,j\}} \sigma_k^{ij} \nu_k \pi_k \right) + \beta' \pi_i + \sigma_j^i \nu_j \pi_j$$

for some $\alpha', \beta' \in \mathbb{R}$. Affine independence implies that $\alpha' = \sigma_j^i \nu_j = \sigma_j \nu_j$. Moreover, $\sigma_i^j = \sigma_i$ and $s\sigma_k^{ij} = \sigma_k^i = \sigma_k$. So $\phi(\succcurlyeq) = \sum_{k \in I} \sigma_k \nu_k \pi_k$, which concludes this case.

Case 2. Assume $\alpha > 1$ and $\beta < 1$. By definition of $\boldsymbol{\sigma}$, we have that $\sigma_k = -\sigma_k^j$ for all $k \in I - \{j\}$ and $\sigma_k = \sigma_k^i$ for all $k \in I - \{i\}$. In particular, $\sigma_k^i = -\sigma_k^j$ for $k \in I - \{i, j\}$. Moreover, either $\sigma_k^{ij} = \sigma_k^j$ for all $k \in I - \{i, j\}$ or $\sigma_k^{ij} = -\sigma_k^j$ for all k. Let $s = 1$ in the former case and $s = -1$ otherwise. Then,

$$\begin{aligned} \phi(\succcurlyeq) &\equiv -\phi(\succcurlyeq_{-j}) + \alpha' \pi_j \equiv -s \sum_{k \in I - \{i,j\}} \sigma_k^{ij} \nu_k \pi_k - \sigma_i^j \nu_i \pi_i + \alpha' \pi_j, \text{ and} \\ &\equiv \phi(\succcurlyeq_{-i}) + \beta' \pi_i \equiv -s \sum_{k \in I - \{i,j\}} \sigma_k^{ij} \nu_k \pi_k + \beta' \pi_i + \sigma_j^i \nu_j \pi_j \end{aligned}$$

for some $\alpha', \beta' \in \mathbb{R}$. The second equality in the second line follows from $-s\sigma_k^{ij} = -\sigma_k^j = \sigma_k^i$ for $k \in I - \{i, j\}$. Affine independence implies that $\alpha' = \sigma_j^i \nu_j = \sigma_j \nu_j$. Moreover, $-\sigma_i^j = \sigma_i$ and $-s\sigma_k^{ij} = -\sigma_k^j = \sigma_k$. So $\phi(\succcurlyeq) = \sum_{k \in I} \sigma_k \nu_k \pi_k$.

The remaining two cases are analogous to the two we have examined and therefore omitted.

Step 7. We have shown that ϕ and $\bar{\phi}$ agree for societies $I = \{1, 2, l\}$ for all l, which we use for the base case $|I| = 3$. Let $I = \{1, i, j\}$ for distinct $i, j \in N - \{1\}$ and $\succcurlyeq \in \mathcal{R}^I$ such that π_1, π_i, and π_j are affinely independent. Observe that ϕ and $\bar{\phi}$ agree on the subprofiles $\succcurlyeq_{-i}$ and $\succcurlyeq_{-j}$ of $\succcurlyeq$ by Step 5. Since $\bar{\phi}(\succcurlyeq_{-i,j}) = \pi_1, \pi_i$, and π_j are affinely independent, Step 6 implies that ϕ and $\bar{\phi}$ agree on $\succcurlyeq$. With a second application of the same argument, we get that ϕ and $\bar{\phi}$ agree on all profiles of three agents with affinely independent beliefs.

In the rest of the proof, we deal with the case $|I| \geq 4$. Moreover, we assume that $\boldsymbol{\pi}$ has dimension at least 3 for now.

Case 1. Suppose $\boldsymbol{\pi}_{-k}$ has dimension 2 for some $k \in I$. Thus, all beliefs in $\boldsymbol{\pi}_{-k}$ are linear combinations of π_i and π_j for distinct but otherwise arbitrary $i, j \in I - \{k\}$. Since $\boldsymbol{\pi}$ has dimension 3, π_k is not in the affine hull of the beliefs in $\boldsymbol{\pi}_{-k}$. So any subprofile of $\boldsymbol{\pi}$ with at least three agents one of which is k has dimension 3. By the induction hypothesis, ϕ and $\bar{\phi}$ agree on such profiles except for possibly $\boldsymbol{\pi}$ itself. In particular, they agree on $\succcurlyeq_{-i}$ and $\succcurlyeq_{-j}$. Moreover, $\bar{\phi}(\succcurlyeq_{-i,j}) = \alpha\pi_i + \beta\pi_j + \gamma\pi_k$ for $\alpha, \beta, \gamma \in \mathbb{R}$ with $\gamma \neq 0$ by Lemma 2. So $\bar{\phi}(\succcurlyeq_{-i,j}), \pi_i$, and π_j are affinely independent. By Step 6, we get that ϕ and $\bar{\phi}$ agree on $\succcurlyeq$.

Case 2. The remaining case is that $\boldsymbol{\pi}_{-k}$ has dimension 3 for all $k \in I$. The induction hypothesis implies that ϕ and $\bar{\phi}$ agree on $\succcurlyeq_{-j}$ and $\succcurlyeq_{-i}$. The argument in the proof of Lemma 3 also applies to $\bar{\phi}$, and so we can choose distinct $i, j \in I$ such that $\bar{\phi}(\succcurlyeq_{-i,j}), \pi_j$, and π_i are affinely independent. We again conclude from Step 6 that ϕ and $\bar{\phi}$ agree on $\succcurlyeq$.

Now let $\succcurlyeq$ be an arbitrary profile for agents in I. We have covered the case when $\boldsymbol{\pi}$ has dimension at least 3. If $\boldsymbol{\pi}$ has dimension 2, let $i \in N - I$ and $(\tilde{\succcurlyeq})$ be a profile for agents in $I \cup \{i\}$ such that $\tilde{\succcurlyeq}_{-i} = \succcurlyeq$ and $\tilde{\boldsymbol{\pi}}$ has dimension 3. Then $\phi(\tilde{\succcurlyeq}) = \bar{\phi}(\tilde{\succcurlyeq})$ and so Step 5 applies, which gives $\phi(\succcurlyeq) = \bar{\phi}(\succcurlyeq)$. Single-agent profiles are covered by Lemma 2. □

A.2. Aggregation of Utility Functions

To begin with, it is useful to clarify the linear algebra on $\bar{\mathcal{U}}$. Elements of $\bar{\mathcal{U}}$ are normalized representatives of a class of utility functions, consisting of all its positive affine transformations. Thus, we say that $(u_i)_{i \in I}$ are linearly independent if their span does not include any utility function that is equivalent to the 0 element of $\bar{\mathcal{U}}$, that is, any constant utility function.

We show that the utility function of a society containing agent i is a linear combination of the utility function of i and that of the society without i. If the latter two utility functions are not equal to completely opposed, then i has positive weight in this linear combination. In Lemma 6, we leverage this fact to prove that the utility function of any society is a *positive* linear combination of the utility functions of its members.

Lemma 5. *Assume that Φ satisfies restricted monotonicity. Let $I \in \mathcal{I}$, $i \in I$, and $\succcurlyeq \in \mathcal{R}^I$. Then $\psi(\succcurlyeq) = \alpha\psi(\succcurlyeq_{-i}) + \beta u_i$ for some $\alpha, \beta \in \mathbb{R}$. Moreover, if $u_i \neq \pm\psi(\succcurlyeq_{-i})$, then $\beta > 0$ and β is unique.*

Proof. Let $\succsim = \Phi(\succcurlyeq)$, $\succsim_{-i} = \Phi(\succcurlyeq_{-i})$, and $\pi_{-i} = \phi(\succcurlyeq_{-i})$. We start in the same way as for Lemma 1. Restricted monotonicity implies that $f \sim g$ whenever $f \sim_{-i} g$, $f \sim_i g$, $\pi_{-i} \circ f^{-1} = \pi_i \circ f^{-1}$, and $\pi_{-i} \circ g^{-1} = \pi_i \circ g^{-1}$. Thus, it follows from Theorem 1 of Gilboa et al. (2004) that $\psi(\succcurlyeq) = \alpha\psi(\succcurlyeq_{-i}) + \beta u_i$ for some $\alpha, \beta \in \mathbb{R}$.

If $u_i \neq \pm\psi(\succcurlyeq_{-i})$, then β is unique. Moreover, we can find probability distributions p and q on O with finite support such that $\psi(\succcurlyeq_{-i})(p) = \psi(\succcurlyeq_{-i})(q)$ and $u_i(p) > u_i(q)$. A theorem of Liapounoff (1940) allows us to construct acts f and g with the following properties: they induce the distributions p and q under π_{-i} and π_i, that is, $p = \pi_{-i} \circ f^{-1} = \pi_i \circ f^{-1}$ and $q = \pi_{-i} \circ g^{-1} = \pi_i \circ g^{-1}$; Thus, $f \sim_{-i} g$ and $f \succ_i g$. Since Φ is monotonic, we get that $f \succ g$. From Lemma 1, we know that $\phi(\succcurlyeq)$ is an affine combination of π_{-i} and π_i, and so $\phi(\succcurlyeq) \circ f^{-1} = p$ and $\phi(\succcurlyeq) \circ g^{-1} = q$. It follows that $f \succ g$ if and only if $\psi(\succcurlyeq)(p) > \psi(\succcurlyeq)(q)$. Thus, $\beta > 0$. □

Lemma 6. *Assume that Φ satisfies restricted monotonicity and faithfulness. Let $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$. Then $\psi(\succcurlyeq) = \sum_{i \in I} w_i u_i$ for some $\boldsymbol{w} \in \mathbb{R}^I$. If $(u_i)_{i \in I}$ are linearly independent, $\boldsymbol{w} \in \mathbb{R}^I_{++}$ and $\boldsymbol{w}$ is unique.*

Proof. The first part is a straightforward corollary of Lemma 5. For the second part, assume that $(u_i)_{i \in I}$ are linearly independent. Let $\boldsymbol{w} \in \mathbb{R}^I$ such that $\psi(\succcurlyeq) = \sum_{i \in I} w_i u_i$. Linear independence implies that $\boldsymbol{w}$ is unique and $u_i \neq \pm\psi(\succcurlyeq_{-i})$ for all $i \in I$. Thus, Lemma 5 implies that $\psi(\succcurlyeq) = \alpha\psi(\succcurlyeq_{-i}) + \beta u_i$ for $\alpha \in \mathbb{R}$ and $\beta > 0$. Since $\psi(\succcurlyeq_{-i})$ is a linear combination of $(u_j)_{j\in I-\{i\}}$ and $\boldsymbol{w}$ is unique, it follows that $w_i = \beta > 0$. □

The dimension of an vector of utility functions $\boldsymbol{u} \in \mathcal{U}^I$ is the maximal number of linearly independent utility functions in $\{u_i : i \in I\}$. The next lemma is the analogue of Lemma 3. Its proof is similar and therefore omitted.

Lemma 7. *Assume that Φ satisfies restricted monotonicity. Let $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$ with $\psi(\succcurlyeq) \neq 0$. If $\boldsymbol{u}$ has dimension at least 3, there are distinct $i, j \in I$ such that $\psi(\succcurlyeq_{-i,j}), u_i$, and u_j are linearly independent.*

In general, $\boldsymbol{w}$ may depend on $\succcurlyeq$. The content of the next lemma is that w_i must not depend on $\succcurlyeq_{-i}$. For the rest of this section, we assume that *beliefs and utility functions are pairwise distinct in any profile.*

Lemma 8. *Assume that $|O| \geq 4$ and Φ satisfies restricted monotonicity and faithfulness. Then there is $\boldsymbol{\omega} : \mathcal{R} \to \mathbb{R}^N_{++}$ such that $\psi(\succcurlyeq) \equiv \sum_{i\in I} \omega_i(\succcurlyeq_i)u_i$ for all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$.*

Proof. The first part is very similar to the construction of $\boldsymbol{\nu}$ in the proof of Lemma 4. For $l \in N - \{1, 2\}$, let $I_l = \{1, 2, l\}$ and $\mathcal{R}_l$ be the set of all $\succcurlyeq \in \mathcal{R}^{I_l}$ such that u_1, u_2, and u_l are linearly independent. Let $l \in N - \{1, 2\}$ be arbitrary and fix some $\tilde{\succcurlyeq} \in \mathcal{R}_l$. By Lemma 6, there is a unique function $\boldsymbol{\nu} : \mathcal{R}_l \to \mathbb{R}^{I_l}_{++}$ such that $\psi(\succcurlyeq) = \sum_{i\in I_l} \nu_i(\succcurlyeq)u_i$ for all $\succcurlyeq \in \mathcal{R}_l$. For $i, j \in I_l$ and $\succcurlyeq \in \mathcal{R}_l$, let $\lambda^{i,j}(\succcurlyeq) = \frac{\nu_j(\succcurlyeq_{-j}, \tilde{\succcurlyeq}_j)\nu_i(\succcurlyeq)}{\nu_j(\tilde{\succcurlyeq})\nu_i(\succcurlyeq_{-j}, \tilde{\succcurlyeq}_j)}$. The fact that $\boldsymbol{\nu}$ maps to $\mathbb{R}^{I_l}_{++}$ ensures that $\lambda^{i,j}$ is well-defined and positive. Then, let $\omega_i(\succcurlyeq_i) = \frac{\nu_i(\succcurlyeq)}{\lambda^{i,j}(\succcurlyeq)}$. Note that the projection of $\mathcal{R}_l$ to $\mathcal{R}$ that returns the preferences of i is onto, and so ω_i is a function on all of $\mathcal{R}$. Here we use the assumption that $|O| \geq 4$, as otherwise $\mathcal{R}_l$ is empty.

With the same arguments as in the proof of Lemma 4, we can show that $\frac{\nu_i(\succcurlyeq)}{\nu_j(\succcurlyeq)}$ is independent of $\succcurlyeq_k$ for $k \in I_l - \{i, j\}$, that $\lambda^{i,j}$ is independent of i and j, and that ω_i is well-defined. Then we have

$$\psi(\succcurlyeq) = \sum_{i\in I_l} \nu_i(\succcurlyeq)u_i \equiv \sum_{i\in I} \frac{\nu_i(\succcurlyeq)}{\lambda^{i,j_i}(\succcurlyeq)} u_i = \sum_{i\in I} \omega_i(\succcurlyeq_i)u_i,$$

where $j_i \in I - \{i\}$ for all $i \in I$.

Since l was arbitrary, we have now defined ω_i for each $i \in N$. However, we have defined ω_1 and ω_2 multiple times, once for each $l \in N - \{1, 2\}$. So we have to check that these definitions are not conflicting. It follows from Lemma 1 that the ratio between ω_1 and ω_2 is the same for each triple $\{1, 2, l\}$. Thus, we can define ω_1 and ω_2 as obtained for, say, $l = 3$ and scale the triples $(\omega_1, \omega_2, \omega_l)$ obtained for the remaining l appropriately.

$\boldsymbol{\omega}$ defines a function that returns a utility function of society for every profile. For $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$, let $\bar{\psi}(\succcurlyeq) \equiv \sum_{i\in I} \omega_i(\succcurlyeq_i)u_i$. The following two observations will carry us a long way in the rest of the proof.

Step 1. Let $\succcurlyeq$ be a profile for agents in I such that $\boldsymbol{u}$ has dimension 3 and $\boldsymbol{u}_{-i}$ has dimension 2. If ψ and $\bar{\psi}$ agree on $\succcurlyeq$, then they also agree on $\succcurlyeq_{-i}$. By Lemma 5 and the assumption, we have

$$\psi(\succcurlyeq) = \alpha\psi(\succcurlyeq_{-i}) + \beta u_i \equiv \sum_{j\in I-\{i\}} \omega_j(\succcurlyeq_j)u_j + \omega_i(\succcurlyeq_i)u_i$$

for some $\alpha, \beta \in \mathbb{R}$. Since u_i is not in the span of $(u_j)_{j\in I-\{i\}}$, we get that $\psi(\succcurlyeq_{-i}) \equiv \sum_{j\in I-\{i\}} \omega_j(\succcurlyeq_j)u_j \equiv \bar{\psi}(\succcurlyeq_{-i})$.

Step 2. Let $\succcurlyeq$ be a profile for agents in I and $i, j \in I$. If $\bar{\psi}(\succcurlyeq_{-i,j}), u_i$, and u_j are linearly independent and ψ and $\bar{\psi}$ agree on $\succcurlyeq_{-i}$ and $\succcurlyeq_{-j}$, then they also agree on $\succcurlyeq$.

Lemma 5 and Lemma 6 imply that

$$\psi(\succcurlyeq) \equiv \psi(\succcurlyeq_{-j}) + \alpha u_j \equiv \overbrace{\sum_{k\in I-\{i,j\}} \omega_k(\succcurlyeq_k)u_k + \omega_i(\succcurlyeq_i)u_i}^{\equiv\bar{\psi}(\succcurlyeq_{-j})} + \alpha' u_j, \text{ and}$$

$$\psi(\succcurlyeq) \equiv \psi(\succcurlyeq_{-i}) + \beta u_i \equiv \underbrace{\sum_{k\in I-\{i,j\}} \omega_k(\succcurlyeq_k)u_k}_{\equiv\bar{\psi}(\succcurlyeq_{-i,j})} + \beta' u_i + \omega_j(\succcurlyeq_j)u_j$$

for some $\alpha, \alpha'\beta, \beta' \in \mathbb{R}_{++}$. Linear independence of $\bar{\psi}(\succcurlyeq_{-i,j}), u_i$, and u_j implies that $\alpha' = \omega_j(\succcurlyeq_j)$, and so ψ and $\bar{\psi}$ agree on $\succcurlyeq$.

With all this in place, we can finish the proof. First we show by induction over $|I|$ that ψ and $\bar{\psi}$ agree on all profiles where $\boldsymbol{u}$ has dimension at least 3. Later we take care of the remaining profiles later.

The base case is $|I| = 3$. Let $\succcurlyeq \in \mathcal{R}^I$. First assume that $I = \{1, i, j\}$ for distinct $i, j \in N - \{1\}$. We have shown that ψ and $\bar{\psi}$ agree for societies of the form $\{1, 2, l\}$ for any l. Thus Step 1 implies that they agree on $\succcurlyeq_{-i}$ and $\succcurlyeq_{-j}$. Moreover, $\bar{\psi}(\succcurlyeq_{-i,j}) = u_1, u_i$, and u_j are linearly independent. So Step 2 implies that ψ and $\bar{\psi}$ agree on $\succcurlyeq$. A second iteration of the same argument implies that they agree on profiles for three arbitrary agents.

Now we deal with the case $|I| \geq 4$ and again assume that $\boldsymbol{u}$ has dimension at least 3.

Case 1. Suppose $\psi(\succcurlyeq) = 0$ or $\bar{\psi}(\succcurlyeq) = 0$. We show that $\psi(\succcurlyeq) = 0$ if and only if $\bar{\psi}(\succcurlyeq) = 0$. Assume for contradiction that $\psi(\succcurlyeq) \neq 0$ and $\bar{\psi}(\succcurlyeq) = 0$. By Lemma 7, we can find distinct i, j such that $\psi(\succcurlyeq_{-i,j}), u_i$, and u_j are linearly independent. Since $|I| \geq 4$, $\boldsymbol{u}_{-i,j}$ has dimension at least 2. If it has dimension exactly 2, then either $\boldsymbol{u}_{-i}$ or $\boldsymbol{u}_{-j}$ has dimension 3, as otherwise $\boldsymbol{u}$ would have dimension 2. Suppose $\boldsymbol{u}_{-j}$ has dimension at least 3. By the induction hypothesis, $\psi(\succcurlyeq_{-j}) = \bar{\psi}(\succcurlyeq_{-j})$. Since $\psi(\succcurlyeq_{-i,j}), u_i$, and u_j are linearly independent, we can conclude that $u_j \neq \pm\psi(\succcurlyeq_{-j})$. But $\bar{\psi}(\succcurlyeq) \equiv \bar{\psi}(\succcurlyeq_{-j}) + \alpha u_j$ for some $\alpha > 0$, and so since $\pm u_j \neq \psi(\succcurlyeq_{-j}) = \bar{\psi}(\succcurlyeq_{-j})$, we get $\bar{\psi}(\succcurlyeq) \neq 0$.

The proof is similar if $\psi(\succcurlyeq) = 0$ and $\bar{\psi}(\succcurlyeq) \neq 0$. Note however, that we find i, j such that $\bar{\psi}(\succcurlyeq_{-i,j}), u_i$, and u_j are linearly independent not directly by Lemma 7, but by the same argument as in its proof.

Case 2. Suppose $\boldsymbol{u}_{-k}$ has dimension 2 for some $k \in I$. Thus, all beliefs in $\boldsymbol{u}_{-k}$ are linear combinations of u_i and u_j for distinct but otherwise arbitrary $i, j \in I - \{k\}$. Since $\boldsymbol{u}$ has dimension 3, u_k is not in the span of the utility functions in $\boldsymbol{u}_{-k}$. So any subprofile of $\boldsymbol{u}$ with at least 3 agents one of which is k has dimension 3. By the induction hypothesis, ψ and $\bar{\psi}$ agree on such profiles except for possibly $\boldsymbol{\succcurlyeq}$ itself. In particular, they agree on $\boldsymbol{\succcurlyeq}_{-i}$ and $\boldsymbol{\succcurlyeq}_{-j}$. Moreover, $\bar{\psi}(\boldsymbol{\succcurlyeq}_{-i,j}) = \alpha u_i + \beta u_j + \gamma u_k$ for $\alpha, \beta, \gamma \in \mathbb{R}$ with $\gamma > 0$ by definition of $\bar{\psi}$. So $\bar{\psi}(\boldsymbol{\succcurlyeq}_{-i,j}), u_i$, and u_j are linearly independent. Step 2 implies that $\psi(\boldsymbol{\succcurlyeq}) = \bar{\psi}(\boldsymbol{\succcurlyeq})$.

Case 3. The remaining case is that $\boldsymbol{u}_{-k}$ has dimension at least 3 for all $k \in I$. The induction hypothesis implies that ψ and $\bar{\psi}$ agree on $\boldsymbol{\succcurlyeq}_{-i}$ and $\boldsymbol{\succcurlyeq}_{-j}$. By Lemma 7 we can choose distinct $i, j \in I$ such that $\psi(\boldsymbol{\succcurlyeq}_{-i,j}), u_i$, and u_j are linearly independent. If $\boldsymbol{u}_{-i,j}$ has dimension 3, the induction hypothesis implies that ψ and $\bar{\psi}$ agree on $\boldsymbol{\succcurlyeq}_{-i,j}$. If $\boldsymbol{u}_{-i,j}$ has dimension 2, then we use the fact that $\boldsymbol{u}_{-i}$ has dimension 3 to apply Step 1 and conclude that ψ and $\bar{\psi}$ agree on $\boldsymbol{\succcurlyeq}_{-i,j}$. In either case, $\bar{\psi}(\boldsymbol{\succcurlyeq}_{-i,j}), u_i$, and u_j are linearly independent. So Step 2 implies that ψ and $\bar{\psi}$ agree on $\boldsymbol{\succcurlyeq}$.

Now let $\boldsymbol{\succcurlyeq}$ be an arbitrary profile for agents in I. We have covered the case when $\boldsymbol{u}$ has dimension at least 3. If $\boldsymbol{u}$ has dimension 2, let $i \in N - I$ and $\tilde{\boldsymbol{\succcurlyeq}}$ be a profile for agents in $I \cup \{i\}$ such that $\tilde{\boldsymbol{\succcurlyeq}}_{-i} = \boldsymbol{\succcurlyeq}$ and $\tilde{\boldsymbol{u}}$ has dimension 3. Then $\psi(\tilde{\boldsymbol{\succcurlyeq}}) = \bar{\psi}(\tilde{\boldsymbol{\succcurlyeq}})$, and so Step 1 implies $\psi(\boldsymbol{\succcurlyeq}) = \bar{\psi}(\boldsymbol{\succcurlyeq})$. Single-agent profiles are covered by Lemma 6. □

B. Implications of Continuity

Recall that the topology on $\mathcal{R}$ is induced by the uniform metric $\sup\{|\mathbb{E}_{\succcurlyeq}(f) - \mathbb{E}_{\succcurlyeq'}(f)| \colon f \in \mathcal{A}\}$. The topology on $\bar{\mathcal{R}}$ consists of the open sets in the topology on $\mathcal{R}$ plus the set $\bar{\mathcal{R}}$. We define topologies on Π and $\bar{\mathcal{U}}$ analogously. For $\pi, \pi' \in \Pi$, the uniform metric $\sup\{|\pi(E) - \pi'(E)| \colon E \in \mathcal{E}\}$ gives a topology on π. For $u, u' \in \mathcal{U}$ (note the absence of constant utility function 0), we also use the uniform metric $\sup\{|u(x) - u'(x)| \colon x \in O\}$. The topology on $\bar{\mathcal{U}}$ is that of $\mathcal{U}$ plus the entire set $\bar{\mathcal{U}}$. So the only neighborhood of the utility function that is constantly 0 is the set $\bar{\mathcal{U}}$ itself. This is the topology $\bar{\mathcal{U}}$ inherits from the space of all utility functions equipped with the uniform metric when taking the quotient by positive affine transformations.

These topologies make the mappings from preference relations to beliefs and utility functions continuous. Likewise, the inverse operation mapping a pair of belief and utility function to a preference relation is continuous. Lemma 5 of Dietrich (2019) is the equivalent of this statement in the framework of Anscombe and Aumann (1963). To ease the notation in the proof of the next lemma, when E is an event and x, y are outcomes, we write xEy for the act which yields x for states in E and y for states in $\Omega - E$.

Lemma 9. *The correspondence $\pi(\succcurlyeq)$ and the function $u(\succcurlyeq)$ mapping $\succcurlyeq \in \bar{\mathcal{R}}$ to the beliefs and the utility function representing $\succcurlyeq$ are (upper-hemi) continuous. Moreover, the function $\succcurlyeq(\pi, u)$ mapping each pair of belief and utility function to the preference relation it induces is continuous.*

Proof. Let $(\succcurlyeq^n)$ be a sequence that converges to $\succcurlyeq$ in $\bar{\mathcal{R}}$. For each n, let $\pi^n \in \pi(\succcurlyeq^n)$ and $u^n = u(\succcurlyeq^n)$.

First we show that (u^n) converges to $u = u(\succcurlyeq)$. Let $x \in O$ and f_x be the act that returns x in all states. We have $\sup\{|u^n(x) - u(x)| \colon x \in O\} = \sup\{|\mathbb{E}_{\succcurlyeq^n}(f_x) - \mathbb{E}_{\succcurlyeq}(f_x)| \colon x \in O\}$, and so (u^n) converges uniformly to u.

Second, we need to show that if (π^n) converges to $\pi' \in \Pi$, then $\pi' \in \pi(\succcurlyeq)$. If $u = 0$, then $\pi(\succcurlyeq) = \Pi$ and there is nothing to show. Otherwise, $\pi(\succcurlyeq) = \{\pi\}$ for some $\pi \in \Pi$. We show that (π^n) converges to π, which implies $\pi = \pi'$. Since $u \neq 0$, we can choose $x, y \in O$ such that $u(x) > u(y)$. Then, for large enough n,

$$\begin{aligned}\sup\{|\pi^n(E) - \pi(E)| : E \in \mathcal{E}\} &= \sup\left\{\left|\frac{\mathbb{E}_{\succcurlyeq^n}(xEy) - u^n(y)}{u^n(x) - u^n(y)} - \frac{\mathbb{E}_{\succcurlyeq}(xEy) - u(y)}{u(x) - u(y)}\right| : E \in \mathcal{E}\right\} \\ &\leq \frac{2}{u(x) - u(y)} \sup\{|\mathbb{E}_{\succcurlyeq^n}(xEy) - \mathbb{E}_{\succcurlyeq}(xEy)| : E \in \mathcal{E}\}\end{aligned}$$

and so (π^n) converges uniformly to π.

Conversely, assume that (π^n) and (u^n) converge to π and u, respectively. Let $\succcurlyeq^n = \succcurlyeq(\pi^n, u^n)$ and $\succcurlyeq = \succcurlyeq(\pi, u)$ be the induced preference relations. If $u = 0$, then $\succcurlyeq$ is complete indifference. Since $\bar{\mathcal{R}}$ is the only neighborhood of complete indifference in $\bar{\mathcal{R}}$, $(\succcurlyeq^n)$ trivially converges to $\succcurlyeq$. Suppose $u \neq 0$. For $\epsilon > 0$, let $n_0 \in \mathbb{N}$ such that for $n \geq n_0$, $\sup\{|\pi^n(E) - \pi(E)| \colon E \in \mathcal{E}\} < \frac{\epsilon}{2}$ and $\sup\{|u^n(x) - u(x)| \colon x \in O\} < \frac{\epsilon}{2}$. ($n_0$ exists since $u \neq 0$.) Then, for all $n \geq n_0$ and $f \in \mathcal{A}$,

$$\begin{aligned}|\mathbb{E}_{\succcurlyeq^n}(f) - \mathbb{E}_{\succcurlyeq}(f)| &= \left|\int_\Omega (u^n \circ f) d\pi^n - \int_\Omega (u \circ f) d\pi\right| \\ &\leq \left|\int_\Omega (u^n \circ f - u \circ f) d\pi^n\right| + \left|\int_\Omega (u \circ f) d\pi^n - \int_\Omega (u \circ f) d\pi\right| \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon\end{aligned}$$

Thus $\mathbb{E}_{\succcurlyeq^n}$ converges uniformly to $\mathbb{E}_{\succcurlyeq}$. □

The basis for Proposition 1 are Lemma 4 and Lemma 8. Lemma 4 allows negative weights for the beliefs and that both lemmas only apply if the beliefs and the utility functions are pairwise distinct. Assuming that the SWF is continuous eliminates both of these issues.

Proposition 1. *For a social welfare function Φ, the following are equivalent.*

(i) Φ satisfies restricted monotonicity, faithfulness, no belief imposition, and continuity

(ii) There are continuous functions $\boldsymbol{\nu}, \boldsymbol{\omega} \colon \mathcal{R} \to \mathbb{R}^N_{++}$ such that for all $I \in \mathcal{I}$ and $\boldsymbol{\succcurlyeq} \in \mathcal{R}^I$, $\Phi(\boldsymbol{\succcurlyeq})$ is represented by $\frac{1}{\sum_{i\in I} \nu_i(\succcurlyeq_i)} \sum_{i\in I} \nu_i(\succcurlyeq_i)\pi_i$ and $\sum_{i\in I} \omega_i(\succcurlyeq_i)u_i$

Proof. One can check easily that (ii) implies (i). The rest of the proof will establish that (i) implies (ii).

Let $\boldsymbol{\nu}$, $\boldsymbol{\omega}$, and for all $I \in \mathcal{I}$, $\boldsymbol{\sigma}^I$ be the functions obtained from Lemma 4 and Lemma 8.

Step 1. We show that $\boldsymbol{\omega}$ is continuous. Let $i, j \in N$ and $\succcurlyeq \in \mathcal{R}^{\{i,j\}}$ such that $\pi_i \neq \pi_j$ and $u_i \neq \pm u_j$; let $(\succcurlyeq_i{}^n)$ be a sequence in $\mathcal{R}$ converging to $\succcurlyeq_i$ and $\succcurlyeq^n = (\succcurlyeq_i{}^n, \succcurlyeq_j)$. By assumption, Φ is continuous, and by Lemma 9, the mapping from preference relations to the corresponding utility functions is continuous. Thus, $u^n = \psi(\succcurlyeq^n) \equiv \omega_i(u_i^n)u_i^n + \omega_j(u_j)u_j$ converges to $u = \psi(\succcurlyeq) \equiv \omega_i(u_i)u_i + \omega_j(u_j)u_j$. First, $\omega_i(\succcurlyeq_i{}^n)$ is bounded, as otherwise, a subsequence of (u^n) would converge to u_i. But this is impossible, since $\omega_j(u_j) \neq 0$ and $u_i \neq \pm u_j$. Now if α is an accumulation point of $(\omega_i(u_i^n))$, then $\alpha u_i + \omega_j(u_j)u_j \equiv u$, since (u^n) converges to u. But $\alpha u_i + \omega_j(u_j)u_j \equiv \beta u_i + \omega_j(u_j)u_j$ if and only if $\alpha = \beta$. So $(\omega_i(u_i^n))$ is bounded and has a unique accumulation point. Thus, it converges to $\omega_i(u_i)$.

Step 2. Let $i, j \in N$ and $\succcurlyeq \in \mathcal{R}^{\{i,j\}}$ such that $\pi_i \neq \pi_j$, $u_i \neq \pm u_j$, and $\Phi(\succcurlyeq)$ is not complete indifference. We show that $\sigma_i^{\{i,j\}}\nu_i$ is continuous at $\succcurlyeq$. (For convenience, we will omit the superscript $\{i,j\}$ from now on.) Let $(\succcurlyeq^n)$ be a sequence of profiles converging to $\succcurlyeq$. Let $\alpha^n = \sigma_i(\succcurlyeq^n)\nu_i(\succcurlyeq_i{}^n)$ and $\beta^n = \sigma_j(\succcurlyeq^n)\nu_j(\succcurlyeq_j{}^n)$, and $\alpha = \sigma_i(\succcurlyeq)\nu_i(\succcurlyeq_i)$ and $\beta = \sigma_j(\succcurlyeq)\nu_j(\succcurlyeq_j)$.

First we prove convergence when j's preferences remain constant at $\succcurlyeq_j$. Let $\bar{\succcurlyeq}^n = (\succcurlyeq_i{}^n, \succcurlyeq_j)$, $\bar{\alpha}^n = \sigma_i(\bar{\succcurlyeq}^n)\nu_i(\succcurlyeq_i{}^n)$, and $\bar{\beta}^n = \sigma_j(\bar{\succcurlyeq}^n)\nu_j(\succcurlyeq_j)$. We need to show that $(\bar{\alpha}^n)$ converges to α. Note that $\bar{\beta}^n$ can only vary in sign but not in absolute value. Since Φ and the correspondence mapping preference relations to the corresponding beliefs are continuous, we have that $\bar{\alpha}^n\pi_i^n + \bar{\beta}^n\pi_j \equiv \phi(\bar{\succcurlyeq}^n) \to \phi(\succcurlyeq) \equiv \alpha\pi_i + \beta\pi_j$. With the same reasoning as in Step 1, we get that $(\bar{\alpha}^n)$ is bounded and has a unique accumulation point. Thus, it converges to α. Similarly, $\sigma_j(\succcurlyeq_i, \succcurlyeq_j{}^n)\nu_j(\succcurlyeq_j{}^n)$ converges to β.

Now we show that (α^n) converges α. We already know that the sequences of absolute values of (α^n) and (β^n) converge to α and β, respectively. So any subsequence $(\alpha^{n_k}, \beta^{n_k})$ such that all α^{n_k} and all β^{n_k} have the same sign converges. By the same reasoning as in the previous paragraph, we conclude that $\alpha^{n_k}\pi_i^{n_k} + \beta^{n_k}\pi_j^{n_k} \equiv \phi(\succcurlyeq^{n_k}) \to \phi(\succcurlyeq) \equiv \alpha\pi_i + \beta\pi_j$. Since $\pi_i \neq \pi_j$, this implies that $(\alpha^{n_k}, \beta^{n_k})$ converges to (α, β). Thus (α^n, β^n) converges to (α, β).

Step 3. Now we deduce that σ_i is always equal to 1. Assume for contradiction that there is a profile $\tilde{\succcurlyeq} \in \mathcal{R}^{\{i,j\}}$ such that $\tilde{\pi}_i \neq \tilde{\pi}_j$, $\tilde{u}_i \neq \pm\tilde{u}_j$, $\Phi(\tilde{\succcurlyeq})$ is not complete indifference, and $\sigma_i(\tilde{\succcurlyeq}) = -1$. Since $\sigma_i\nu_i$ and $\sigma_j\nu_j$ are continuous at $\tilde{\succcurlyeq}$ by Step 2, we can find a neighborhood of $\tilde{\succcurlyeq}$ such that $\sigma_i(\succcurlyeq) = -1$ for all profiles $\succcurlyeq$ contained in it. In particular, we can find $\epsilon > 0$ such that $\sigma_i(\succcurlyeq) = -1$ whenever $\succcurlyeq_i = \tilde{\succcurlyeq}_i$, $u_j = \tilde{u}_j$, and $\sup\{|\pi_j(E) - \tilde{\pi}_j(E)| \colon E \in \mathcal{E}\} < \epsilon$. Let $\tilde{\mathcal{P}}$ be this set of profiles. By Liapounoff's theorem, we can find an event E such that $\tilde{\pi}_i(E) = \tilde{\pi}_j(E) = \frac{\epsilon}{2}$. But then $\tilde{\mathcal{P}}$ contains a profile $\succcurlyeq$ such that $\sigma_i(\succcurlyeq) = -1$, $\pi_i(E) = \tilde{\pi}_i(E) = \frac{\epsilon}{2}$ and $\pi_j(E) = 0$. This is not possible, since $\sigma_i(\succcurlyeq)\nu_i(\succcurlyeq_i)\pi_i(E) + \sigma_j(\succcurlyeq)\nu_j(\succcurlyeq_j)\pi_j(E)$ would be negative.

Since j was arbitrary and $\boldsymbol{\sigma}^I$ satisfies the restriction on the ratio of σ_i^I and σ_j^I stated in Lemma 4, it follows that σ_i^I is constant at 1 for all I. Since we have shown that $\sigma_i\nu_i$ is continuous, so is ν_i.

(Alternatively, one could show that the set of profiles where beliefs and utility functions are pairwise distinct and Φ is not complete indifference is connected. Then the fact that $\sigma_i\nu_i$ is continuous and never 0 implies that it cannot change sign.)

Step 4. Let $\bar{\Phi}$ be the SWF where $\bar{\Phi}(\succcurlyeq)$ is represented by $\bar{\phi}(\succcurlyeq) \equiv \sum_{i\in I} \nu_i(\succcurlyeq_i)\pi_i$ and $\bar{\psi}(\succcurlyeq) \equiv \sum_{i\in I} \omega_i(\succcurlyeq_i)u_i$ for every profile $\succcurlyeq$ with agents in $I \in \mathcal{I}$. We know that Φ and $\bar{\Phi}$ agree on all profiles for which beliefs and utility functions are pairwise distinct. These profiles are dense in $\mathcal{R}^I$ for all I. Our task is to show that they agree on an arbitrary profile $\succcurlyeq \in \mathcal{R}^I$.

Case 1. Suppose neither $\Phi(\succcurlyeq)$ nor $\bar{\Phi}(\succcurlyeq)$ is complete indifference. By Lemma 9, Step 1, and Step 2, $\phi, \psi, \bar{\phi}$, and $\bar{\psi}$ are continuous at $\succcurlyeq$. Moreover, the pairs ϕ and $\bar{\phi}$ and ψ and $\bar{\psi}$ agree on a set of profiles with $\succcurlyeq$ in its closure. Thus, $\phi(\succcurlyeq) = \bar{\phi}(\succcurlyeq)$ and $\psi(\succcurlyeq) = \bar{\psi}(\succcurlyeq)$. It follows that $\Phi(\succcurlyeq) = \bar{\Phi}(\succcurlyeq)$.

Case 2. Suppose that $\Phi(\succcurlyeq)$ is complete indifference. (The proof is analogous if $\bar{\Phi}(\succcurlyeq)$ is complete indifference.) Let $i \in N - I$ with preferences $\succcurlyeq_i$ such that $u_i \neq \pm\bar{\psi}(\succcurlyeq)$. Then for $\succcurlyeq_{+i} = (\succcurlyeq, \succcurlyeq_i)$, by Lemma 5, $\psi(\succcurlyeq_{+i}) = u_i$ and $\bar{\psi}(\succcurlyeq_{+i}) \equiv \bar{\psi}(\succcurlyeq) + \alpha u_i$ for some $\alpha > 0$. In particular, $\psi(\succcurlyeq_{+i}), \bar{\psi}(\succcurlyeq_{+i}) \neq 0$. Thus Case 1 implies that $\psi(\succcurlyeq_{+i}) = \bar{\psi}(\succcurlyeq_{+i})$. Since $u_i \neq \pm\bar{\psi}(\succcurlyeq)$, this can only be if $\bar{\psi}(\succcurlyeq) = 0$, and hence $\bar{\Phi}(\succcurlyeq)$ is complete indifference.

□

C. Implications of Independence of Redundant Actions

Using independence of redundant acts, we derive a lemma which, together with Proposition 1, concludes the proof of Theorem 1. But first we need two auxiliary statements. Recall that a function is simple if it has finite range.

Lemma 10. *Let $I \in \mathcal{I}$ and $i \in I$; let $\succcurlyeq \in \mathcal{R}^I$ such that u_j is simple for all $j \in I - \{i\}$. Then for every act f, there is a simple act g such that $f \sim_j g$ for all $j \in I$.*

Proof. Put differently, we want to show that for every act f, there is a simple act g such that $(\mathbb{E}_{\succcurlyeq_j}(f))_{j\in I} = (\mathbb{E}_{\succcurlyeq_j}(g))_{j\in I}$.

We first show that the sets $O^+ = \{x \in O\colon u_i(x) \geq \mathbb{E}_{\succcurlyeq_i}(f)\}$ and $O^- = \{x \in O\colon u_i(x) \leq \mathbb{E}_{\succcurlyeq_i}(f)\}$ are non-empty. If O^+ is empty, then $\Omega = \bigcup_{k\in N}\{s \in \Omega\colon u_i(f(s)) \leq \mathbb{E}_{\succcurlyeq_i}(f) - \frac{1}{k}\}$. Note that all sets in this union are measurable. Since π_i is countably additive, there is k_0 such that $\pi_i(\{s \in \Omega\colon u_i(f(s)) \leq \mathbb{E}_{\succcurlyeq_i}(f) - \frac{1}{k_0}\}) = \epsilon > 0$. The fact that O^+ is empty then gives $\mathbb{E}_{\succcurlyeq_i}(f) \leq \mathbb{E}_{\succcurlyeq_i}(f) - \frac{\epsilon}{k_0}$, which is a contradiction. Similarly, one shows that O^- is non-empty.

Now let $V = \{\boldsymbol{u}_{-i}(x)\colon x \in O\} \subset \mathbb{R}^{I-\{i\}}$ be the range of $\boldsymbol{u}_{-i}$. Since all u_j are simple, V is finite. We partition the set of outcomes O into the measurable sets $O_v = \boldsymbol{u}_{-i}^{-1}(v)$ and the set of states Ω into the measurable sets $E_v = f^{-1}(O_v)$ with v ranging over V. To define g, consider two cases. If $\pi_i(E_v) = 0$, choose $x_v \in O_v$ arbitrarily and let $g(s) = x_v$ for all $s \in E_v$. If $\pi_i(E_v) > 0$, then $\frac{1}{\pi_i(E_v)}\pi_i|_{E_v}$ is a probability measure on E_v, where $\pi_i|_{E_v}$ is π_i restricted to events contained in E_v. By the previous paragraph, the sets $O_v^+ = \{x \in O_v\colon u_i(x) \geq \frac{1}{\pi_i(E_v)}\int_{E_v}(u_i \circ f)d\pi_i|_{E_v}\}$ and $O_v^- = \{x \in O_v\colon u_i(x) \leq \frac{1}{\pi_i(E_v)}\int_{E_v}(u_i \circ f)d\pi_i|_{E_v}\}$ are non-empty. Thus, there are $x^+ \in O_v^+$ and $x^- \in O_v^-$ and $\alpha \in [0,1]$ such that $\alpha u_i(x^+) + (1-\alpha)u_i(x^-) = \frac{1}{\pi_i(E_v)}\int_{E_v}(u_i \circ f)d\pi_i|_{E_v}$. Since π_i is non-atomic, there is $E_v^+ \subset E_v$ such that $\pi_i(E_v^+) = \alpha\pi_i(E_v)$. We define $g(s) = x^+$ for

$s \in E_v^+$ and $g(s) = x^-$ for $s \in E_v - E_v^+$. This gives

$$\int_{E_v} (u_i \circ f) d\pi_i|_{E_v} = \pi_i(E_v) \left(\alpha u_i(x^+) + (1-\alpha) u_i(x^-)\right)$$
$$= \pi_i(E_v^+) u_i(x^+) + \pi_i(E_v - E_v^+) u_i(x^-) = \int_{E_v} (u_i \circ g) d\pi_i|_{E_v}$$

Also, since $\boldsymbol{u}_{-i}$ is constant on O_v, $\int_{E_v} (u_j \circ f) d\pi_j|_{E_v} = \int_{E_v} (u_j \circ g) d\pi_j|_{E_v}$ for all $j \in I - \{i\}$. In summary, we have $(\mathbb{E}_{\succcurlyeq_j}(f))_{j \in I} = (\mathbb{E}_{\succcurlyeq_j}(g))_{j \in I}$. ☐

Lemma 11. *Let π and π' be non-atomic probability measures on $(\Omega, \mathcal{E})$. Then there is a sub-sigma-algebra $\mathcal{E}'$ of $\mathcal{E}$ such that $\pi|_{\mathcal{E}'} = \pi'|_{\mathcal{E}'}$ ($\pi(E) = \pi'(E)$ for all $E \in \mathcal{E}'$) and $\pi|_{\mathcal{E}'}, \pi'|_{\mathcal{E}'}$ are non-atomic on $(\Omega, \mathcal{E}')$.*

Proof. We first construct an increasing sequence of sub-sigma-algebras $\mathcal{E}_k \subset \mathcal{E}$ so that π and π' agree on each $\mathcal{E}_k$. In the second step, we show that the sigma-algebra generated by $\bigcup_{k \geq 0} \mathcal{E}_k$ has the desired properties.

We define $\mathcal{E}_k$ inductively so that $\mathcal{E}_k \subset \mathcal{E}_{k+1}$ and $\mathcal{E}_k$ is the sigma-algebra generated by a partition $\{E_k^1, \ldots, E_k^{2^k}\}$ of Ω with $\pi(E_k^m) = \pi'(E_k^m) = \frac{1}{2^k}$ for all m. Let $\mathcal{E}_0 = \{\emptyset, \Omega\}$ (that is, the sigma-algebra generated by $\{\Omega\}$). Clearly, $\mathcal{E}_0$ has the required properties. Now let $k \geq 1$ and assume we have constructed $\mathcal{E}_l$ with the required properties for $l < k$. Let $\{E_{k-1}^1, \ldots, E_{k-1}^{2^{k-1}}\}$ be the partition generating $\mathcal{E}_{k-1}$. For each E_{k-1}^m, we can use Liapounoff's theorem to divide it into two equal halves when measured by π and π'. That is, we can find $E_k^{2m-1}, E_k^{2m} \in \mathcal{E}$ such that $\{E_k^{2m-1}, E_k^{2m}\}$ partitions E_{k-1}^m and $\pi(E_k^{2m-1}) = \pi'(E_k^{2m-1}) = \frac{1}{2^k}$. Then if we let $\mathcal{E}_k$ be the sigma-algebra generated by the partition $\{E_k^1, \ldots, E_k^{2^k}\}$, it has the required properties.

Now let $\mathcal{F} = \bigcup_{k \geq 0} \mathcal{E}_k$ and $\mathcal{E}'$ be the sigma-algebra generated by $\mathcal{F}$. We note two facts. First, $\mathcal{F}$ is an algebra.[8] Second, $\mathcal{M} = \{E \in \mathcal{E} \colon \pi(E) = \pi'(E)\}$ is a monotone class containing $\mathcal{F}$.[9] Thus, the monotone class theorem implies that $\mathcal{E}' \subset \mathcal{M}$.

It remains to show that $\pi|_{\mathcal{E}'}$ and $\pi'|_{\mathcal{E}'}$ are non-atomic on $(\Omega, \mathcal{E}')$. Since $\pi|_{\mathcal{E}'} = \pi'|_{\mathcal{E}'}$, it suffices to prove the statement for π. Let $E \in \mathcal{E}'$ with $\pi(E) > 0$. Choose k such that $\frac{1}{2^k} < \pi(E)$. By construction of $\mathcal{E}_k$, we can partition Ω into sets $E^1, \ldots, E^{2^k} \in \mathcal{E}_k \subset \mathcal{E}'$ such that $\pi(E^m) = \frac{1}{2^k}$ for all m. Note that $E \cap E^m \in \mathcal{E}'$ for all m since $\mathcal{E}'$ is a sigma-algebra. Then the sets $E \cap E^1, \ldots, E \cap E^{2^k}$ partition E and so $\pi(E \cap E^m) > 0$ for some m. It follows that $0 < \pi(E \cap E^m) \leq \pi(E^m) = \frac{1}{2^k} < \pi(E)$, which proves non-atomicity. ☐

Proposition 1 characterizes SWFs that aggregate beliefs and utilities linearly so that the weight of each agent in either linear combination is a function of the agent's own preferences only. The final lemma shows that for any such SWF Φ that additionally satisfies independence of redundant acts, the weights have to be constant.

[8] A collection of subsets $\mathcal{F}$ of Ω is an algebra if (i) $\emptyset \in \mathcal{F}$, (ii) $\mathcal{F}$ is closed under taking complements, and (iii) $\mathcal{F}$ is closed under *finite* unions.

[9] A collection of subsets $\mathcal{M}$ of Ω is a monotone class if it is closed under countable monotone unions and countable monotone intersections. That is, $(E_k)_{k \geq 0} \subset \mathcal{M}$ with $E_0 \subset E_1 \subset \ldots$ implies $\bigcup_{k \geq 0} E_k \in \mathcal{M}$ and $(E_k)_{k \geq 0} \subset \mathcal{M}$ with $E_0 \supset E_1 \supset \ldots$ implies $\bigcap_{k \geq 0} E_k \in \mathcal{M}$.

Lemma 12. *Let $\boldsymbol{\nu}, \boldsymbol{\omega} \colon \mathcal{R} \to \mathcal{R}^N_{++}$ be continuous functions; let Φ be an SWF such that for every $I \in \mathcal{I}$ and $\boldsymbol{\succcurlyeq} \in \mathcal{R}^I$, $\Phi(\boldsymbol{\succcurlyeq})$ is represented by $\frac{1}{\sum_{i\in I}\nu_i(\succcurlyeq_i)}\sum_{i\in I}\nu_i(\succcurlyeq_i)\pi_i$ and $\sum_{i\in I}\omega_i(\succcurlyeq_i)u_i$. Then if Φ satisfies independence of redundant acts, $\boldsymbol{\nu}$ and $\boldsymbol{\omega}$ are constant across all preferences.*

Proof. Let $i, j \in N$ and $\succcurlyeq_i, \succcurlyeq_i{}' \in \mathcal{R}$. We want to show that $\nu_i(\succcurlyeq_i) = \nu_i(\succcurlyeq_i{}')$ and $\omega_i(\succcurlyeq_i) = \omega_i(\succcurlyeq_i{}')$. In the first step, we show that ν_i and ω_i are independent of π_i. In the rest of the proof, we show that they are independent of u_i, too.

Step 1. Assume that $u_i = u_i'$. First we show $\nu_i(\succcurlyeq_i) = \nu_i(\succcurlyeq_i{}')$. By Lemma 11, we can find a sub-sigma-algebra $\mathcal{E}'$ of $\mathcal{E}$ such that $\pi_i|_{\mathcal{E}'} = \pi_i'|_{\mathcal{E}'}$ and $\pi_i|_{\mathcal{E}'}, \pi_i'|_{\mathcal{E}'}$ are non-atomic on $(\Omega, \mathcal{E}')$. We construct a belief for agent j that allows us to leverage independence of redundant acts. Let $E \in \mathcal{E}'$ such that $\pi_i(E) = \frac{1}{2}$. If $\pi_i = \pi_i'$, there is nothing to show. Otherwise, either $\pi_i(F) \neq \pi_i'(F)$ for some $F \subset E$ or $\pi_i(F) \neq \pi_i'(F)$ for some $F \subset \Omega - E$. Assume the former is true. Then define π_j so that $\pi_j(F) = 2\pi_i(F)$ for every $F \subset E$ and $\pi_j(\Omega - E) = 0$. Moreover, choose $u_j \in \mathcal{U}$ so that u_j is simple and $u_j \neq \pm u_i$ and let $\succcurlyeq_j$ be represented by π_j and u_j.

The set of acts to which we will apply independence of redundant acts is $\mathcal{A}' = \mathcal{A}(\mathcal{E}', O)$. To meet the antecedent of independence of redundant acts, we have to show that $\mathcal{A}'$ is co-redundant for the profiles $\boldsymbol{\succcurlyeq} = (\succcurlyeq_i, \succcurlyeq_j)$ and $\boldsymbol{\succcurlyeq}' = (\succcurlyeq_i{}', \succcurlyeq_j)$. Condition (ii) of co-redundancy is satisfied, since $\pi_i|_{\mathcal{E}'}$, $\pi_i'|_{\mathcal{E}'}$, and $\pi_j|_{\mathcal{E}'}$ are non-atomic by the choice of $\mathcal{E}'$. To verify condition (i), we show that for every $f \in \mathcal{A}$, there is $g \in \mathcal{A}'$ such that $f \sim_i g$ and $f \sim_j g$. (The choice of $\mathcal{A}'$ and $u_i = u_i'$ ensure that also $f \sim_i' g$.)

By Lemma 10, we may assume that f is simple. Define g as follows: let $f(\Omega) = \{x_1, \ldots, x_k\}$ be the range of f. For every x_l, let $\alpha_l = \pi_i(E \cap f^{-1}(x_l))$ and $\alpha_l^c = \pi_i((\Omega - E) \cap f^{-1}(x_l))$. (Note that $\pi_j(E \cap f^{-1}(x_l)) = 2\alpha_l$.) The non-atomicity of $\pi_i|_{\mathcal{E}'}$ allows us to find events $E_l \subset E$ and $E_l^c \subset \Omega - E$ in $\mathcal{E}'$ such that $\pi_i(E_l) = \alpha_l$ and $\pi_i(E_l^c) = \alpha_l^c$. In fact, we can partition E and $\Omega - E$ into $\{E_1, \ldots, E_k\}$ and $\{E_1^c, \ldots, E_k^c\}$, respectively. Then let $g(s) = x_l$ for $s \in E_l \cup E_l^c$. One can check that $\pi_j(E_l \cup E_l^c) = 2\pi_i(E_l) = 2\alpha_l$. Thus, $\mathbb{E}_{\succcurlyeq_i}(f) = \mathbb{E}_{\succcurlyeq_i}(g)$ and $\mathbb{E}_{\succcurlyeq_j}(f) = \mathbb{E}_{\succcurlyeq_j}(g)$ and so $f \sim_i g$ and $f \sim_j g$.

Let $\succcurlyeq = f(\boldsymbol{\succcurlyeq})$ and $\succcurlyeq' = f(\boldsymbol{\succcurlyeq}')$ and π, u, π', u' be the corresponding beliefs and utility functions. Independence of redundant acts applied to $\boldsymbol{\succcurlyeq}$ and $\boldsymbol{\succcurlyeq}'$ gives $g \succcurlyeq g'$ if and only if $g \succcurlyeq' g'$ for all $g, g' \in \mathcal{A}'$.

Assume that $\nu_i(\succcurlyeq_i) \neq \nu_i(\succcurlyeq_i{}')$. First, since $u_j \neq \pm u_i$ and $u \equiv \nu_i(u_i)u_i + \nu_j(u_j)u_j$, $\succcurlyeq$ cannot be complete indifference, and so we can find outcomes x and y such that $x \succ y$. Recall that $\pi_i(E) = \pi_i'(E) = \frac{1}{2}$ and $\pi_j(E) = 1$. It follows that $\pi(E) \neq \pi'(E)$ and $\pi(E), \pi'(E) > \frac{1}{2}$. So there is an event $E' \in \mathcal{E}'$ such that $E' \subset E$ and $\pi(E') \neq \pi'(E') = \frac{1}{2}$. Thus, $xE'y$ and $yE'x$ are acts in $\mathcal{A}'$ but $xE'y \not\sim yE'x$ and $xE'y \sim' yE'x$. This contradicts independence of redundant acts and so $\nu_i(\succcurlyeq_i) = \nu_i(\succcurlyeq_i{}')$.

Second, assume that $\omega_i(\succcurlyeq_i) \neq \omega_i(\succcurlyeq_i{}')$. Since $u_j \neq \pm u_i$, it follows that $u \neq u'$ and we can a find simple lotteries p and q on O such that $u(p) > u(q)$ but $u'(q) \geq u'(p)$. By the previous paragraph, $\nu_i(\succcurlyeq_i) = \nu_i(\succcurlyeq_i{}')$ and so $\pi|_{\mathcal{E}'} = \pi'|_{\mathcal{E}'}$. Moreover, $\pi|_{\mathcal{E}'}$ is non-atomic as a weighted mean of measures that are non-atomic on $\mathcal{E}'$. So we can find acts g and g' in $\mathcal{A}'$ with $g \circ \pi = p$

and $g' \circ \pi' = q$. This gives $g \succ g'$ but $g' \not\succcurlyeq' g$, which contradicts independence of redundant acts. We conclude that $\omega_i(\succcurlyeq_i) = \omega_i(\succcurlyeq_i')$.

Step 2. By Step 1, we can view ν_i and ω_i as functions $\nu_i(u_i)$ and $\omega_i(u_i)$ of u_i. We show that both these functions are constant.

Recall that $\mathcal{U}$ consists of utility functions which are normalized to the unit interval, that is, $\inf_x u(x) = 0$ and $\sup_x u(x) = 1$. Let $\mathcal{U}' = \{u \in \mathcal{U}\colon \text{there exist } x, y \in O \text{ with } u(x) = 0 \text{ and } u(y) = 1\}$ be those utility functions for which the infimum and the supremum are attained. Observe that the closure of $\mathcal{U}'$ is $\mathcal{U}$. Thus, since ν_i and ω_i are continuous, it suffices to show that they are constant on $\mathcal{U}'$. This we do now.

Let $u_i \in \mathcal{U}'$; let $x_0, x_1 \in O$ such that $u_i(x_0) = 0$ and $u_i(x_1) = 1$ and $x^* \in O - \{x_0, x_1\}$ be arbitrary; let u_i' be such that $u_i'(x) = u_i(x)$ for $x \in \{x_0, x_1, x^*\}$. We show that $\nu_i(u_i) = \nu_i(u_i')$ and $\omega_i(u_i) = \omega_i(u_i')$. Since $|O| \geq 4$, repeated application of this statement gives the same conclusion for all $u_i' \in \mathcal{U}'$.

Let u_i'' be such that

$$u_i''(x) = \begin{cases} u_i(x^*) & \text{if } u_i(x) < u_i(x^*) \text{ and } u_i'(x) > u_i(x^*) \\ u_i'(x) & \text{if } u_i(x) < u_i(x^*) \text{ and } u_i'(x) \leq u_i(x^*) \\ u_i(x) & \text{if } u_i(x) \geq u_i(x^*) \end{cases}$$

Note that $u_i''(x) = u_i'(x) = u_i(x)$ for $x \in \{x_0, x_1, x^*\}$. We want to apply independence of redundant acts to profiles with utility functions (u_i, u_j) and (u_i'', u_j) and the set of acts $\mathcal{A}' = \mathcal{A}(\mathcal{E}, \{x_0, x_1, x^*\})$. This requires choosing u_j appropriately. Let u_j be such that

$$u_j(x) = \begin{cases} 0 & \text{if } u_i(x) \leq u_i(x^*) \\ u_i(x) & \text{otherwise} \end{cases}$$

Figure 2 depicts the images of (u_i, u_j) and (u_i'', u_j) in utility space. From u_i to u_i'', we adjust the utility for outcomes with $u_i(x) \leq u_i(x^*)$ toward $u_i'(x)$ without raising it above $u_i(x^*)$. Setting u_j as we did, we can now apply independence of redundant acts to the corresponding profiles.

Let $\pi_i, \pi_j \in \Pi$ with $\pi_i \neq \pi_j$ and $\succcurlyeq_i, \succcurlyeq_i''$, and $\succcurlyeq_j$ be represented by the pairs $(\pi_i, u_i), (\pi_i, u_i'')$, and (π_j, u_j), respectively. First, since $u_i(x) = u_i''(x)$ for $x \in \{x_0, x_1, x^*\}$, it is clear that $\boldsymbol{\succcurlyeq} = (\succcurlyeq_i, \succcurlyeq_j)$ and $\boldsymbol{\succcurlyeq}'' = (\succcurlyeq_i'', \succcurlyeq_j)$ agree on the preferences over acts in $\mathcal{A}'$. Second, since $\boldsymbol{u}(x)$ is in the convex hull of $\{\boldsymbol{u}(x_0), \boldsymbol{u}(x_1), \boldsymbol{u}(x^*)\}$ for all $x \in O$, we have that for every act $f \in \mathcal{A}$, there is an act $g \in \mathcal{A}'$ such that $f \sim_i g$ and $f \sim_j g$. The analogous assertion holds for $\succcurlyeq_i''$ and $\succcurlyeq_j$. Thus, $\mathcal{A}'$ satisfies condition (i) of co-redundancy for $\boldsymbol{\succcurlyeq}$ and $\boldsymbol{\succcurlyeq}''$. Condition *(ii)* holds since π_i and π_j are (trivially) non-atomic on $\mathcal{E}$. It follows from independence of redundant acts that with $\succcurlyeq = \Phi(\boldsymbol{\succcurlyeq})$ and $\succcurlyeq'' = \Phi(\boldsymbol{\succcurlyeq}'')$, we have for all $g, g' \in \mathcal{A}'$, $g \succcurlyeq g'$ if and only if $g \succcurlyeq'' g'$. Let (π, u) and (π'', u'') be the beliefs and utility functions associated with $\succcurlyeq$ and $\succcurlyeq''$, respectively. Note that $u(x_0) = u''(x_0) = 0$ and $u(x_1) = u''(x_1) = 1$.

If $\nu_i(u_i) \neq \nu_i(u_i'')$, then $\pi \neq \pi''$ since $\pi_i \neq \pi_j$. So we can find an event E such that $\pi(E) = \frac{1}{2} \neq \pi''(E)$. It follows that $x_0 E x_1 \sim x_1 E x_0$ but $x_0 E x_1 \not\sim'' x_1 E x_0$, which is a contradiction since both acts are in $\mathcal{A}'$.

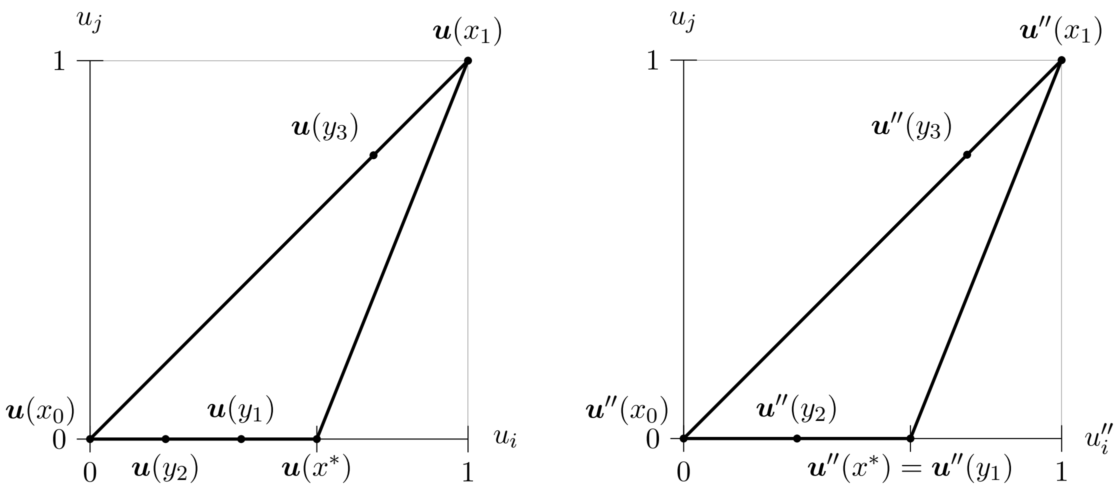

Figure 2: The images of $\boldsymbol{u} = (u_i, u_j)$ and $\boldsymbol{u}'' = (u_i'', u_j)$ in utility space. For example, $\boldsymbol{u}(x_0) = (u_i(x_0), u_j(x_0)) = (0, 0)$. The outcomes y_1, y_2, and y_3 are examples for the three cases in the definition of u_i''.

If $\omega_i(u_i) \neq \omega_i(u_i'')$, then $u(x^*) \neq u''(x^*)$, since $u_i(x^*) = u_i''(x^*) \neq u_j(x^*)$. Let E be an event such that $\pi(E) = \pi''(E) = u(x^*)$. Then $x^* \sim x_1 E x_0$ but $x^* \not\sim x_1 E x_0$, which is again a contradiction.

We conclude that $\nu_i(u_i) = \nu_i(u_i'')$ and $\omega_i(u_i) = \omega_i(u_i'')$. The function u_i'' is closer to u_i' than is u_i, since we have constructed it by moving utilities toward those in u_i'. Two more modifications of agent 1's utility function along the same lines will result in u_i'. To this end, we apply the same construction first to the profiles with utility functions (u_i'', u_j') and (u_i''', u_j') and then to profiles with utility functions (u_i''', u_j) and (u_i', u_j) (and the same beliefs π_i and π_j).

$$u_i'''(x) = \begin{cases} u_i''(x^*) & \text{if } u_i''(x) \geq u_i''(x^*) \text{ and } u_i'(x) < u_i''(x^*) \\ u_i'(x) & \text{if } u_i''(x) \geq u_i''(x^*) \text{ and } u_i'(x) \geq u_i''(x^*) \\ u_i''(x) & \text{if } u_i''(x) < u_i''(x^*) \end{cases} \quad u_j'(x) = \begin{cases} 1 & \text{if } u_i''(x) \geq u_i''(x^*) \\ u_i''(x) & \text{otherwise} \end{cases}$$

In summary, this gives $\nu_i(u_i) = \nu_i(u_i')$ and $\omega_i(u_i) = \omega_i(u_i')$ and proves the lemma.

□

We complete the proof of Theorem 1.

Theorem 1. *For a social welfare function Φ, the following are equivalent.*

(i) Φ satisfies restricted monotonicity, independence of redundant acts, faithfulness, no belief imposition, and continuity

(ii) There are $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}^N_{++}$ such that for all $I \in \mathcal{I}$ and $\succcurlyeq \in \mathcal{R}^I$, $\Phi(\succcurlyeq)$ is represented by $\frac{1}{\sum_{i \in I} v_i} \sum_{i \in I} v_i \pi_i$ and $\sum_{i \in I} w_i u_i$

Proof. It follows from Proposition 1 and Lemma 12 that (i) implies (ii).

To prove the converse, we need to show that Φ satisfies all axioms if it has the form stated in (ii). It is easy to see that restricted monotonicity, faithfulness, no belief imposition, and continuity hold. We verify that independence of redundant acts is also satisfied. Let $I \in \mathcal{I}$ and $\succcurlyeq, \succcurlyeq' \in \mathcal{R}^I$ be two profiles. Let $\mathcal{E}'$ be a sub-sigma-algebra of $\mathcal{E}$ and $O' \subset O$ be a set of outcomes so that $\mathcal{A}' = \mathcal{A}(\mathcal{E}', O')$ is co-redundant for $\succcurlyeq$ and $\succcurlyeq'$. For all $i \in I$, both $(\pi_i|_{\mathcal{E}'}, u_i|_{O'})$ and $(\pi_i'|_{\mathcal{E}'}, u_i'|_{O'})$ represent $\succcurlyeq_i|_{\mathcal{A}'} = \succcurlyeq_i'|_{\mathcal{A}'}$. By condition (ii) of co-redundancy, the representation of $\succcurlyeq_i|_{\mathcal{A}'}$ is unique up to positive affine transformations of the utility function. Hence, $\pi_i|_{\mathcal{E}'} = \pi_i'|_{\mathcal{E}'}$. Moreover, by condition (i) of co-redundancy, $\inf\{u_i(x) \colon x \in O'\} = 0$ and $\sup\{u_i(x) \colon x \in O'\} = 1$ and likewise for u_i'. Thus, $u_i|_{O'} = u_i'|_{O'}$. It follows that $\sum_{i \in I} v_i \pi_i|_{\mathcal{E}'} = \sum_{i \in I} v_i \pi_i'|_{\mathcal{E}'}$ and $\sum_{i \in I} w_i u_i|_{O'} = \sum_{i \in I} w_i u_i'|_{O'}$ and so $\Phi(\succcurlyeq)|_{\mathcal{A}'} = \Phi(\succcurlyeq')|_{\mathcal{A}'}$. □

D. Reference Table

Symbol	[elements]	Name	Mathematical object
Ω		states	set
$\mathcal{E}$	E	events	sigma-algebra on Ω
O	x, y	outcomes	set
$\mathcal{A}$	f, g	acts	measurable functions $\Omega \to O$
Π	π	beliefs	non-atomic probability measures on $(\Omega, \mathcal{E})$
$\mathcal{U}$ ($\bar{\mathcal{U}}$)	u	utility functions	measurable functions $O \to \mathbb{R}$ with inf 0 and sup 1 (plus the function constant at 0)
$\mathcal{R}$ ($\bar{\mathcal{R}}$)	$\succcurlyeq$	preference relations	SEU preferences on $\mathcal{A}$ (plus complete indifference)
N	i	potential agents	infinite set
$\mathcal{I}$	I	societies	non-empty finite subsets of N
$\mathcal{R}^I$	$\boldsymbol{\succcurlyeq}$	preference profiles	I-tuples with components in $\mathcal{R}$
$\succcurlyeq_i$		preferences of agent i	component of $\boldsymbol{\succcurlyeq}$ with index i
π_i		belief of agent i	unique element of Π representing $\succcurlyeq_i$
u_i		utility function of agent i	unique element of $\mathcal{U}$ representing $\succcurlyeq_i$
Φ		social welfare function	function from $\bigcup_{I \in \mathcal{I}} \mathcal{R}^I \to \bar{\mathcal{R}}$