Bayesian social aggregation with almost-objective uncertainty*

Marcus Pivato and Élise Flore Tchouante†

November 15, 2022

Abstract

We consider collective decisions under uncertainty, when agents have *generalized Hurwicz* preferences, a broad class allowing many different ambiguity attitudes, including subjective expected utility preferences. We consider sequences of acts that are "almost-objectively uncertain" in the sense that asymptotically, all agents almost-agree about the probabilities of the underlying events. We introduce a Pareto axiom which applies *only* to asymptotic preferences along such almost-objective sequences. This axiom implies that the social welfare function is utilitarian, but it does not impose any constraint on collective beliefs. On the other hand, a Pareto axiom for "dichotomous" acts implies that collective beliefs are contained in the closed convex hull of individual beliefs, but imposes no constraints on the social welfare function. Neither axiom entails any link between individual and collective ambiguity attitudes.

Keywords. Bayesian social aggregation; almost-objective uncertainty; generalized Hurwicz; Bewley preferences; utilitarian.

JEL class: D70; D81.

Anyone who considers arithmetical methods of producing random digits is, of course, in a state of sin. —John von Neumann

1 Introduction

From a democratic point of view, collective decisions should be made by aggregating the preferences or opinions of the affected individuals. But almost all nontrivial decisions involve uncertainty. Normative decision theory considers the question of how rational agents

*We are grateful to Eric Danan, Nicolas Gravel and Bill Zwicker for their helpful comments. We would also like to thank the Associate Editor and three reviewers for exceptionally thorough, detailed, and helpful reports, which greatly improved the paper. This research was supported by Labex MME-DII (ANR11-LBX-0023-01) and CHOp (ANR-17-CE26-0003).

†CY Cergy Paris Université, CNRS, THEMA, F-95000 Cergy, France, email: marcus-pivato@cyu.fr and elise-flore.tchouante-ngamo@cyu.fr

should cope with such uncertainty. Bayesian social aggregation combines these two ingredients: it aims for collective decisions that are both rational and democratic. The foundational result is Harsanyi's (1955) Social Aggregation Theorem. Harsanyi considered a society in which all agents are von Neumann-Morgenstern (vNM) expected utility maximizers. He showed that if the vNM preferences of the social planner satisfy an ex ante Pareto axiom relative to the vNM preferences of the individuals, then the social welfare function —that is, the vNM utility function of the social planner —must be a weighted average of the individual vNM utility functions. Harsanyi interpreted this as a strong argument for utilitarianism.

Harsanyi's result is highly influential in social choice theory, but its dependence on the vNM framework curtails its applicability. The vNM framework assumes that all risks can be quantified with known, objective probabilities. But in many complex decision problems (e.g. macroeconomics, climate change, pandemics), it is not clear how to assign precise probabilities to the relevant contingencies. Indeed, when considering *sui generis* events in the future (e.g. hypothetical wars or financial crises in 2060), it is not clear that "objective" probabilities even exist. This led Savage (1954) to propose an approach to decision-making based on the maximization of *subjective* expected utility (SEU) —that is, expected utility computed using the agent's own "subjective" probabilistic beliefs.

A central tenet of the Savagean framework is that different rational agents may reasonably hold *different* subjective beliefs. But Mongin (1995) showed that Harsanyi's theorem breaks down in settings with heterogeneous beliefs. Mongin (1997) diagnosed the root of the problem as *spurious unanimity*: different individuals might have different utility functions *and* different beliefs, but these differences might "cancel out" to yield a unanimous ex ante preference amongst them for one act over another, thereby entailing (via the ex ante Pareto axiom) a corresponding ex ante social preference.

This suggests that to avoid Mongin's impossibility theorem, one should weaken the ex ante Pareto axiom to avoid cases of spurious unanimity. This strategy was realized in a landmark paper by Gilboa et al. (2004), who proposed a "restricted" ex ante Pareto axiom that only applied to acts for which all agents have the *same* probabilistic beliefs about the underlying events. Gilboa et al. showed that this restricted Pareto axiom has two consequences: (1) the social welfare function (SWF) must be a weighted sum of individual utility functions, and (2) the social beliefs must be a weighted average of individual beliefs.[1]

One objection to Gilboa et al.'s result is that it is not always appropriate to construct social beliefs as an arithmetic average of individual beliefs. For example, this way of aggregating beliefs does not interact well with Bayesian updating. In response, Dietrich (2021) has recently obtained a result similar to that of Gilboa et al. (2004), in which social beliefs are a weighted *geometric* average of individual beliefs. This ensures compatibility with Bayesian updating. But it does not address a broader issue. Different belief-aggregation rules are suitable in different contexts, and the criteria that determine the appropriate belief-aggregation rule are not necessarily the criteria that determine the correct social

[1]See §6 for a more detailed discussion of Gilboa et al. (2004). Recently, Brandl (2021) has obtained a similar result, but in his case, the SWF is *relative* utilitarian: it is a sum of the utility functions of individuals rescaled to range from 0 to 1. See also Billot and Qu (2021).

welfare function. The specification of collective beliefs is an *epistemic* problem, whereas the specification of the SWF is an *ethical* problem; there is no reason that these two problems should be solved by the same theorem.[2] For this reason, Mongin and Pivato (2020) and Pivato (2022) have recently introduced weak Pareto axioms which entail a utilitarian SWF, but which do not impose *any* constraints on collective beliefs. They thus concentrate on the ethical problem, leaving the epistemic problem to be solved later by other methods.

The present paper takes up this challenge: it addresses both problems, but deals with them independently of one another. We exploit the phenomenon of *almost-objective uncertainty* (due to Poincaré 1912 and Machina 2004, 2005), which involves a sequence of partitions $\mathfrak{G}^1, \mathfrak{G}^2, \mathfrak{G}^3, \ldots$ such that even agents with very different beliefs will assign increasingly similar probabilities to the cells of $\mathfrak{G}^n$ as $n \to \infty$. We propose a weak Pareto axiom, which only applies to asymptotic preferences for sequences of acts measurable with respect to these partitions. Our first main result says that this axiom is both necessary and sufficient for the SWF to be a weighted sum of individual utility functions (Theorem 1). But unlike results in the aforementioned literature, it does not impose any relationship between individual and collective beliefs.

We then turn to belief aggregation. We consider a second weak Pareto axiom, which only applies to preferences between two-valued acts for which all agents have the same preferences over the outcomes. Our second main result (Theorem 2) connects this axiom to the social aggregation of individual beliefs. But it does not impose any constraint on the SWF. Thus, the two theorems decouple the ethical problem from the epistemic problem, and deal with them separately.

The previous paragraph was vague about the belief aggregation in Theorem 2. If all agents have SEU preferences, then Theorem 2 says the social beliefs are a weighted average of individual beliefs, as in Gilboa et al. (2004). But to fully explain the result, we must broaden our perspective. All of the aforementioned literature assumes that all agents have SEU preferences. But in ambiguous decision environments, this might be inappropriate; it might be difficult to specify *any* single probability measure over contingencies as an adequate description of the uncertainty faced by an agent. This objection is both normative and descriptive. At a descriptive level, many agents might simply be *unable* to condense their uncertainty into a single probability measure. At a normative level, it is perhaps not even *rational* for an agent to resort to such a probabilistic description. These concerns have inspired a variety of *non-SEU* models of decision making. Typically such models represent an agent's beliefs not with a single probability measure but with an *ensemble* of probability measures, and in addition to her utility function, they often involve other parameters. For succinctness, we shall describe this entire package (i.e. a non-SEU decision model and its associated parameters) as the agent's *ambiguity attitude*.

This raises the question of whether non-SEU ambiguity attitudes can be incorporated into collective decisions. But just as different agents can reasonably hold different probabilistic beliefs, different agents can reasonably adopt different ambiguity attitudes. Such heterogeneity leads once again to impossibility theorems (Chambers and Hayashi, 2006; Gajdos et al., 2008; Mongin and Pivato, 2015; Zuber, 2016). In general, to satisfy the

[2]See §4.7 of Pivato (2022) for further elaboration of these points.

ex ante Pareto axiom, all agents must not only have the same beliefs, but the same ambiguity attitudes —indeed, they must be SEU maximizers.[3] Once again, to escape this undesirable conclusion, one must weaken the ex ante Pareto axiom; this strategy has been explored in a series of elegant papers by Alon and Gayer (2016), Danan et al. (2016), Qu (2017) and Hayashi and Lombardi (2019).[4] Like the foundational result of Gilboa et al. (2004), these more recent papers axiomatically characterize not only a SWF, but a procedure for aggregating individual beliefs into a collective belief. As already noted, non-SEU models generally represent agents' beliefs by ensembles of probability measures, so these procedures aggregate these ensembles. Thus, they are vulnerable to the same objections earlier raised against Gilboa et al. (2004) and Dietrich (2021): different belief-aggregation rules are appropriate in different environments, and in any case, collective beliefs should not necessarily be determined at the same time as the social welfare function. Furthermore, these theorems generally impose a particular ambiguity attitude on society (either in their hypotheses or in their conclusions).

Aside from heterogeneity of beliefs, another problem confronts the SEU framework adopted by Mongin (1995) and Gilboa et al. (2004): that of *state-dependent utility*. In certain situations, it may be perfectly reasonable for an agent's utility function to depend upon what state of nature is realized.[5] This creates two problems for Bayesian social aggregation. First, it makes it unclear how to impute probabilistic beliefs to the individual based on her ex ante preferences, as noted by Schervish et al. (1990) and Karni (1996), among others (see Baccelli (2017) for an excellent recent discussion of this problem). Second, in the specification of the SWF, it raises the question of *which utility function* we should impute to each individual. For these reasons (among others) Duffie (2014) and Sprumont (2018, 2019) have rejected the approach pioneered by Gilboa et al. (2004) of weakening ex ante Pareto so as to separately aggregate beliefs and utilities. Instead Sprumont (2018, 2019) uses the full-strength ex ante Pareto axiom to characterize two approaches to Bayesian social aggregation based entirely the aggregation of individuals' ex ante preferences. The cost of these purely ex ante approaches is a loss of collective rationality: social decisions are no longer consistent with SEU maximization.[6]

The results of the present paper are compatible with both heterogeneity of beliefs *and* heterogeneity of ambiguity attitudes, and even compatible with certain forms of state-dependent utility. Theorem 1 is formulated for *generalized Hurwicz* preferences, a broad class that includes SEU preferences, maximin SEU preferences, Hurwicz preferences, and second-order SEU preferences, among others. Theorem 2 is formulated for *Bewley preferences*. In both preference classes, each agent's beliefs are described by a set of probability measures. The precise statement of Theorem 2 is that the belief set underlying collective preferences must be contained in the closed convex hull of the union of the belief sets underlying the individual preferences. Importantly, neither Theorem 1 nor Theorem 2

[3]In fact, when all agents have maximin SEU preferences, or all have Hurwicz preferences, Hayashi (2021) has shown that ex ante Pareto implies dictatorship, *even if* all agents have the same beliefs.

[4]See Mongin and Pivato (2016) or Fleurbaey (2018) for reviews of this literature.

[5]See e.g. Section 2.8 and Appendix 2A of Drèze (1987).

[6]See also Ceron and Vergopoulos (2019) for an interesting hybrid of ex ante and ex post.

imposes any relationship between individual ambiguity attitudes and collective ambiguity attitudes. We see this as an advantage. Just as the specification of the SWF is an ethical problem, and the specification of collective beliefs is an epistemic problem, the specification of collective ambiguity attitudes is a problem of *prudential rationality*. It is better to disentangle these three problems. This paper focuses on the first two problems, leaving the prudential problem for future work.

The rest of this paper is organized as follows. Section 2 introduces generalized Hurwicz representations. Section 3 introduces almost-objective uncertainty, and provides several sufficient conditions for the existence of almost-objective uncertainty. Section 4 turns to social welfare; it introduce a concept of "asymptotic preferences" based on almost-objective uncertainty and a corresponding Pareto axiom, along with the statement of Theorem 1 and several corollaries. Section 5 turns to belief aggregation, and contains our second Pareto axiom and Theorem 2. Section 6 discusses some possible variations and extensions of our results, and compares them to some prior literature. All proofs are in the Appendices.

2 Generalized Hurwicz representations

Let $\mathcal{S}$ and $\mathcal{X}$ be measurable spaces —i.e. sets equipped with sigma-algebras.[7] We shall refer to $\mathcal{S}$ as the ***state space*** and $\mathcal{X}$ as the ***outcome space***. Let $\Delta(\mathcal{S})$ be the set of all countably additive probability measures on $\mathcal{S}$. An ***act*** is a measurable function $\alpha : \mathcal{S} \longrightarrow \mathcal{X}$ that takes only finitely many values. Let $\mathcal{A}$ be the set of all acts. Let $\succcurlyeq$ be a preference order on $\mathcal{A}$. In the Savage model of uncertainty, $\mathcal{X}$ is a set of "outcomes", while $\mathcal{S}$ is a set of possible "states of nature"; the true state is unknown. The order $\succcurlyeq$ describes an agent's ex ante preferences. A ***representation*** of $\succcurlyeq$ is a function $V : \mathcal{A} \longrightarrow \mathbb{R}$ such that

$$\text{for all } \alpha, \beta \in \mathcal{A}, \quad \Big(\alpha \succcurlyeq \beta\Big) \quad \Longleftrightarrow \quad \Big(V(\alpha) \geqslant V(\beta)\Big). \tag{1}$$

In particular, V is a ***subjective expected utility*** (SEU) ***representation*** if there is some $\rho \in \Delta(\mathcal{S})$ and a bounded measurable function $u : \mathcal{X} \longrightarrow \mathbb{R}$ such that

$$V(\alpha) \quad = \quad \int_{\mathcal{S}} u \circ \alpha \, \mathrm{d}\rho, \qquad \text{for all } \alpha \in \mathcal{A}. \tag{2}$$

Here, ρ is interpreted as the agent's *subjective beliefs* about the unknown state of nature, while u describes the utility she would obtain from each outcome. But as noted in Section 1, in situations of ambiguity, it might be inappropriate to represent an agent's beliefs as a single probability measure over $\mathcal{S}$. This has led to classes of preferences that use an *ensemble* of probability measures. This paper will focus on a broad class of such preferences: those admitting a *generalized Hurwicz* representation.

[7] For simplicity, we shall not make these sigma-algebras explicit in our notation. A set will never be equipped with more than one sigma-algebra in this paper.

A representation V is ***generalized Hurwicz*** (GH) if there is a closed convex subset $\mathcal{P} \subseteq \Delta(\mathcal{S})$ and a bounded measurable function $u : \mathcal{X} \longrightarrow \mathbb{R}$, such that

$$\text{for all } \alpha \in \mathcal{A}, \quad \underline{V}(\alpha) \leqslant V(\alpha) \leqslant \overline{V}(\alpha), \tag{3}$$
$$\text{where} \quad \underline{V}(\alpha) := \inf_{\rho \in \mathcal{P}} \int_{\mathcal{S}} u \circ \alpha \, \mathrm{d}\rho \quad \text{and} \quad \overline{V}(\alpha) := \sup_{\rho \in \mathcal{P}} \int_{\mathcal{S}} u \circ \alpha \, \mathrm{d}\rho.$$

The idea here is that the agent is not only unsure of the true state of nature, but also unsure about the correct probability distribution to put on $\mathcal{S}$; the set $\mathcal{P}$ contains all probabilities that she considers *possible*. The GH representation (3) encompasses a wide gamut of preferences. It reduces to the SEU representation (2) if $\mathcal{P}$ is a singleton. It obviously includes the class of *maximin SEU* (or *multiple priors*) preferences characterized by Gilboa and Schmeidler (1989) (for which $V(\alpha) = \underline{V}(\alpha)$, for all $\alpha \in \mathcal{A}$), and also the classical *Hurwicz* (or *α-maximin*) preferences introduced by Hurwicz (1951) and recently characterized by Chateauneuf et al. (2020) and Hartmann (2022) (for which $V(\alpha) = q\,\underline{V}(\alpha) + (1-q)\,\overline{V}(\alpha)$, for all $\alpha \in \mathcal{A}$, for some constant $q \in [0,1]$). It also includes the class of *second order SEU* (or *smooth ambiguity*) preferences characterized by Klibanoff et al. (2005) and the *Choquet expected utility* preferences of Schmeidler (1989). More generally, Cerreia-Vioglio et al. (2011, Prop.5) show that any *monotone, Bernoullian, Archimedean* (MBA) preference admits a GH representation like (3), generalizing an earlier result of Ghirardato et al. (2004, Prop.7) for *invariant biseparable* preferences.[8]

Let $\mathcal{M}(\mathcal{S})$ be the vector space of all signed measures on $\mathcal{S}$. This becomes a Banach space when equipped with the total variation norm

$$\|\mu\|_{\mathrm{vr}} := \sup_{\substack{\mathcal{H}_1,\ldots,\mathcal{H}_N \subseteq \mathcal{S} \\ \text{disjoint measurable}}} \sum_{n=1}^{N} |\mu[\mathcal{H}_n]| \,. \tag{4}$$

We will say that a GH representation (3) is ***compact*** if $\mathcal{P}$ is compact in this norm. We shall say it is ***nonatomic*** if all elements of $\mathcal{P}$ are nonatomic measures. We shall say that a representation V is ***contiguous*** if its image $V(\mathcal{A})$ is a dense subset of an interval in $\mathbb{R}$. For example, if $\mathcal{X}$ is a connected topological space and $u : \mathcal{X} \longrightarrow \mathbb{R}$ is continuous, then any GH representation (3) with u as its utility function is contiguous.[9]

The goal of this paper is not to axiomatically characterize GH representations. We shall simply *assume* that the agents' preference have such representations; in light of the generality of this class, this is a reasonable assumption. But different agents might have *different* representations, with different u and $\mathcal{P}$. Thus, our framework allows great diversity in the beliefs and ambiguity attitudes of the agents.

Uniqueness. Cerreia-Vioglio et al. (2011, Proposition 5(iii)) have shown that any MBA preference has *unique* GH representation, up to positive affine transformation of the utility

[8] These results allow $\mathcal{P}$ to include *finitely* additive measures. But one can easily constrain $\mathcal{P}$ to countably additive measures by imposing a suitable continuity axiom; see e.g. Proposition B.1 of Ghirardato et al. (2004).

[9] To see this, let α range over all constant-valued acts, to deduce that $V(\mathcal{A}) = u(\mathcal{X})$.

function. But Cerreia-Vioglio et al. work in the Anscombe-Aumann framework, whereas the present paper assumes the Savage framework, and is not restricted to MBA preferences. Thus, a preference order $\succcurlyeq$ might have many GH representations. How much do they have in common?

First, note that if $V : \mathcal{A} \longrightarrow \mathbb{R}$ is a GH representation for $\succcurlyeq$, then the utility function u in expression (3) is entirely determined by V: for any $x \in \mathcal{X}$, we have $u(x) = V(\kappa_x)$, where $\kappa_x \in \mathcal{A}$ is just the constant act with value x. Conversely, suppose that $\succcurlyeq$ satisfies the following mild condition:

Certainty equivalents. For any $\alpha \in \mathcal{A}$, there is some $x \in \mathcal{X}$ such that $\alpha \approx \kappa_x$.

(For example, if $\mathcal{X}$ is connected and $u : \mathcal{X} \longrightarrow \mathbb{R}$ is continuous, then $\succcurlyeq$ satisfies Certainty equivalents.) In this case, V is also entirely determined by u, because for any $\alpha \in \mathcal{A}$ we have $V(\alpha) = u(x)$, where $x \in \mathcal{X}$ is any outcome such that $\alpha \approx \kappa_x$. Thus, for preferences satisfying Certainty equivalents, V and u codetermine each other.

Also, for a given GH representation V, there is a unique *minimal* subset $\mathcal{P}_* \subseteq \Delta(\mathcal{S})$ satisfying the inequalities (3). To see this, let $\mathfrak{P}$ be the set of all subsets of $\Delta(\mathcal{S})$ that satisfy (3), and let $\mathcal{P}_* := \bigcap_{\mathcal{P} \in \mathfrak{P}} \mathcal{P}$. It is clear that $\mathcal{P}_*$ also satisfies (3). In the rest of the paper, we always assume that we are working with this minimal $\mathcal{P}_*$.

These remarks show that u and $\mathcal{P}_*$ are unique for a given representation V. But couldn't $\succcurlyeq$ have two *different* representations V and V', described by two different utility functions u and u' and two different minimal belief sets $\mathcal{P}_*$ and $\mathcal{P}'_*$? The next result addresses this question.

Proposition 1 *Suppose $\succcurlyeq$ satisfies* Certainty equivalents*. If V and V' are compact, nonatomic GH representations for $\succcurlyeq$, then they have the same minimal belief set $\mathcal{P}_*$, and there are constants $a > 0$ and $b \in \mathbb{R}$ such that $V' = a\,V + b$.*

3 Almost-objective uncertainty

A *measurable partition* of $\mathcal{S}$ is a finite collection $\mathfrak{G} = \{\mathcal{G}_n\}_{n=1}^N$ of disjoint measurable subsets such that $\mathcal{S} = \bigsqcup_{n=1}^N \mathcal{G}_n$. For any $K \in \mathbb{N}$, let $\Delta^K := \{\mathbf{q} = (q_1, \ldots, q_K) \in \mathbb{R}_+^K;\ \sum_{k=1}^K q_k = 1\}$, the set of K-dimensional probability vectors.

Let $\mathcal{R}$ be a collection of probability measures on $\mathcal{S}$. Let $K \in \mathbb{N}$ and let $\mathbf{q} \in \Delta^K$. For all $n \in \mathbb{N}$, let $\mathfrak{G}^n := \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$ be a K-element measurable partition of $\mathcal{S}$. We shall say that the sequence of partitions $(\mathfrak{G}^n)_{n=1}^\infty$ is $\mathcal{R}$-*almost-objectively uncertain* and *subordinate to* $\mathbf{q}$ if, for all $\rho \in \mathcal{R}$, we have

$$\lim_{n\to\infty} \rho(\mathcal{G}_k^n) \quad = \quad q_k, \qquad \text{for all } k \in [1 \ldots K]. \tag{5}$$

For example, let $\mathcal{S} = [0, 1]$, and let $\mathcal{R}$ be the set of all probability measures that are absolutely continuous with respect to the Lebesgue measure, with continuous density functions. Suppose $\mathbf{q} = (0.1, 0.2, 0.3, 0.4)$. For any number $s \in [0, 1]$ and $n \in \mathbb{N}$, let $s_{(n)}$ be the nth digit in the decimal expansion of s.[10] For all $n \in \mathbb{N}$, let $\mathfrak{G}^n := \{\mathcal{G}_1^n, \mathcal{G}_2^n, \mathcal{G}_3^n, \mathcal{G}_4^n\}$, where $\mathcal{G}_1^n := \{s \in [0,1]; \ s_{(n)} = 0\}$, $\mathcal{G}_2^n := \{s \in [0,1]; \ s_{(n)} \in \{1,2\}\}$ $\mathcal{G}_3^n := \{s \in [0,1]; \ s_{(n)} \in \{3,4,5\}\}$, and $\mathcal{G}_4^n := \{s \in [0,1]; \ s_{(n)} \in \{6,7,8,9\}\}$. It is easily seen that $(\mathfrak{G}^n)_{n=1}^{\infty}$ is $\mathcal{R}$-almost-objectively uncertain and subordinate to $\mathbf{q}$.

Almost-objective uncertainty was first introduced by Poincaré (1912) to explain why it is reasonable to hold particular epistemic probabilities regarding a physical randomization device such as a roulette wheel, even if we do not have an exact understanding of how this apparent randomness is generated. Its first application to decision-making under ambiguity was due to Machina (2004, 2005), who also coined the term "almost-objective uncertainty". Poincaré and Machina considered almost-objective uncertainty on the unit interval $[0,1]$, as in the above example. We will now generalize this concept to a much broader collection of state spaces and probability measures. Let $\mathcal{S}$ be a measurable space, and let $\mathcal{R} \subseteq \Delta(\mathcal{S})$. We shall say that $\mathcal{R}$ is *consilient* if, for any $K \in \mathbb{N}$ and $\mathbf{q} \in \Delta^K$, there is an $\mathcal{R}$-almost-objectively uncertain sequence of partitions $(\mathfrak{G}^n)_{n=1}^{\infty}$ subordinate to $\mathbf{q}$. The results in this section give sufficient conditions for consilience. We need some terminology. A subset $\mathcal{R} \subseteq \Delta(\mathcal{S})$ is *nonatomic* if all elements of $\mathcal{R}$ are nonatomic. It is *separable* if it has a countable dense subset in the topology of the total variation norm (4).

Proposition 2 *If $\mathcal{R}$ is nonatomic and separable, then $\mathcal{R}$ is consilient.*

It is sometimes convenient to have a consilient set that is closed under Bayesian updating. For any subset $\mathcal{R} \subseteq \Delta(\mathcal{S})$, let $\langle\mathcal{R}\rangle := \{\mu \in \Delta(\mathcal{S}); \ \mu$ is absolutely continuous with respect to some $\rho \in \mathcal{R}$, and the Radon-Nikodym derivative $\frac{d\mu}{d\rho}$ is bounded$\}$. In particular, $\langle\mathcal{R}\rangle$ includes all measures that arise from a Bayesian update of some element of $\mathcal{R}$. Let us say that $\mathcal{R}$ is *strongly consilient* if $\langle\mathcal{R}\rangle$ is consilient.

The next result gives two sufficient conditions for strong consilience. First we need some terminology. A probability measure μ on $\mathcal{S}$ is *separable* if there is a countable set of events $\{\mathcal{E}_n\}_{n=1}^{\infty}$ that is *dense*: for any measurable $\mathcal{B} \subseteq \mathcal{S}$, and any $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that $\mathcal{B}$ is "ϵ-approximated" by $\mathcal{E}_n$ in the sense that $\mu[\mathcal{B}\backslash\mathcal{E}_n] < \epsilon$ and $\mu[\mathcal{E}_n\backslash\mathcal{B}] < \epsilon$. (Equivalently, μ is separable if the Banach space $\mathcal{L}^1(\mathcal{S}, \mu)$ is separable.) For example, the Lebesgue probability measure on $[0,1]$ is separable. Most probability spaces that arise in practical applications are separable.

A *standard Borel space* is a measurable space $\mathcal{S}$ that is measurably isomorphic to a complete, separable metric space $\mathcal{S}'$ (e.g. a closed subset of $\mathbb{R}^N$), endowed with its Borel sigma algebra. (That is: there is a measurable bijection from $\mathcal{S}$ to $\mathcal{S}'$ whose inverse is also measurable.) Every Polish space is a standard Borel space. But a standard Borel space need not have a Polish topology (or indeed, any topology at all). Almost every measurable space encountered in applications is standard Borel.[11]

Proposition 3 *Suppose that $\mathcal{R} \subseteq \Delta(\mathcal{S})$ is nonatomic and separable, and suppose that either* (a) *Every element of $\mathcal{R}$ is separable; or* (b) *$\mathcal{S}$ is a standard Borel space. Then $\mathcal{R}$ is strongly consilient.*

[10]There is a countable subset of $[0,1]$ of numbers with non-unique decimal expansions, for whom $s_{(n)}$ is not well-defined. But it has Lebesgue measure zero, so it is irrelevant to this construction.

[11]For a good introduction to standard Borel spaces, see §424, p.158 of Fremlin (2006a).

Although the scopes of Propositions 2 and 3 are already very broad, there are many other examples of consilient collections of measures. To illustrate this, let $\widehat{\mathcal{S}}$ and $\mathcal{S}$ be two measurable spaces, and let $\phi : \widehat{\mathcal{S}} \longrightarrow \mathcal{S}$ be any measurable function. This induces a function $\phi_* : \mathcal{M}(\widehat{\mathcal{S}}) \longrightarrow \mathcal{M}(\mathcal{S})$ where, for any $\widehat{\mu} \in \mathcal{M}(\widehat{\mathcal{S}})$ and any measurable $\mathcal{B} \subseteq \mathcal{S}$, we define $\phi_*(\widehat{\mu})[\mathcal{B}] := \widehat{\mu}[\phi^{-1}(\mathcal{B})]$.

Proposition 4 *Let $\widehat{\mathcal{S}}$ and $\mathcal{S}$ be measurable spaces, and let $\phi : \widehat{\mathcal{S}} \longrightarrow \mathcal{S}$ be measurable. Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$, and let $\widehat{\mathcal{R}} := (\phi_*)^{-1}(\mathcal{R}) \subseteq \Delta(\widehat{\mathcal{S}})$. If $\mathcal{R}$ is (strongly) consilient, then $\widehat{\mathcal{R}}$ is (strongly) consilient.*

Consilience in dynamical systems. Dynamical systems are mathematical models of systems evolving deterministically in time. They arise frequently in the study of ordinary differential equations, difference equations, and all parts of applied mathematics. Formally, a ***(measurable) dynamical system*** is a pair $(\mathcal{S}, \phi)$, where $\mathcal{S}$ is a measurable space and $\phi : \mathcal{S} \longrightarrow \mathcal{S}$ is a measurable function. A probability measure μ on $\mathcal{S}$ is ***ϕ-invariant*** if $\phi_*(\mu) = \mu$. The triple $(\mathcal{S}, \mu, \phi)$ is then called a ***measure-preserving dynamical system*** (MPDS). A wide variety of dynamical systems admit invariant measures, and hence can be treated as MPDS. For example, if $\mathcal{S}$ is any compact metric space and $\phi : \mathcal{S} \longrightarrow \mathcal{S}$ is continuous, then the Krylov-Bogolyubov theorem yields an invariant measure for ϕ (Walters, 1982, §6.2, p.152).

An MPDS $(\mathcal{S}, \mu, \phi)$ is ***mixing*** if, for all measurable subsets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{S}$, we have $\lim_{t\to\infty} \mu[\mathcal{A} \cap \phi^{-t}(\mathcal{B})] = \mu(\mathcal{A}) \cdot \mu(\mathcal{B})$. Many MPDS are mixing —in particular, ones which exhibit so-called "chaotic" behaviour. For example, let $\mathcal{S} = [0, 1]$. The ***tent map*** $\phi : [0, 1] \longrightarrow [0, 1]$ is defined:

$$\phi(s) \quad = \quad \begin{cases} 2\,s & \text{if } s \leqslant \frac{1}{2}; \\ 1 - 2\,s & \text{if } s > \frac{1}{2}. \end{cases}$$

The Lebesgue measure on $[0, 1]$ is ϕ-invariant, and the resulting MPDS is mixing.

Proposition 5 *Let $(\mathcal{S}, \mu, \phi)$ be any mixing MPDS. Let $\mathcal{R} := \{\rho \in \Delta(\mathcal{S});\ \ \rho \ll \mu$ and $\frac{\mathrm{d}\rho}{\mathrm{d}\mu} \in \mathcal{L}^2(\mathcal{S}, \mu)\}$. Then $\mathcal{R}$ is strongly consilient.*

This result addresses a possible concern about Propositions 2 and 3. Whereas the almost-objectively uncertain partition sequences constructed in Propositions 2 and 3 might seem somewhat exotic, the sequences constructed in Proposition 5 are extremely natural: they take a single partition of $\mathcal{S}$ and shift it into the far future via ϕ. Many standard examples of "effectively random" questions have this form, such as, "What will the temperature in Times Square be at 12:00 PM on April 1, 2062?"[12] It is not implausible that such questions could arise in collective decisions. This provides additional motivation for the Almost-objective Pareto axiom that we will introduce in Section 4.

[12]Here we assume that global weather patterns can be described as a chaotic dynamical system.

4 Social aggregation of utility

As noted in Section 1, a central problem in Bayesian social aggregation is that different agents might have different probabilistic beliefs and different attitudes towards ambiguity. We shall now use almost-objective uncertainty to obviate these problems.

Almost-objective acts. Let $\mathcal{R}$ be a consilient collection of probability measures on a measurable space $\mathcal{S}$. Let $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^{\infty}$ be a sequence of acts. We shall say that $\boldsymbol{\alpha}$ is an *$\mathcal{R}$-almost-objective act* if there is a K-tuple of outcomes $\mathbf{x} \in \mathcal{X}^K$ (for some $K \in \mathbb{N}$), and an $\mathcal{R}$-almost-objectively uncertain sequence of K-cell partitions $\boldsymbol{\mathcal{G}} = (\mathfrak{G}^n)_{n=1}^{\infty}$, with $\mathfrak{G}^n := \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$ for all $n \in \mathbb{N}$, such that for all $n \in \mathbb{N}$ and $k \in [1 \ldots K]$ we have $\alpha^n(s) = x_k$ for all $s \in \mathcal{G}_k^n$. If $\boldsymbol{\mathcal{G}}$ is subordinate to the probability vector $\mathbf{q} \in \Delta^K$, then we shall say that $\boldsymbol{\alpha}$ is *subordinate to* $(\mathbf{q}, \mathbf{x})$.

Let $\boldsymbol{\beta} = (\beta^n)_{n=1}^{\infty}$ be another almost-objective act. We shall say that $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are *compatible* if β^n is also measurable with respect to $\mathfrak{G}^n$ for all $n \in \mathbb{N}$.

Asymptotic preferences. Let $\succcurlyeq$ be a preference order on $\mathcal{A}$. Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be almost-objective acts. We shall say $\succcurlyeq$ *asymptotically prefers* $\boldsymbol{\alpha}$ to $\boldsymbol{\beta}$, and write $\boldsymbol{\alpha} \succ^{\infty} \boldsymbol{\beta}$ if there exist $\alpha', \beta' \in \mathcal{A}$ and $N \in \mathbb{N}$ such that $\alpha^n \succ \alpha' \succ \beta' \succ \beta^n$ for all $n \geqslant N$.

Almost-objective Pareto. Let $\mathcal{I}$ be a set of individuals. Let o be another agent, representing a social planner or social observer. Let $\mathcal{J} = \mathcal{I} \sqcup \{o\}$. For all $j \in \mathcal{J}$, let $\succcurlyeq_j$ be a preference order on $\mathcal{A}$. We shall require $\succcurlyeq_o$ to satisfy the following axiom, relative to $\{\succcurlyeq_i\}_{i \in \mathcal{I}}$ and $\mathcal{R}$:

$\mathcal{R}$-Almost-objective Pareto. If $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are compatible $\mathcal{R}$-almost-objective acts, and $\boldsymbol{\alpha} \succ_i^{\infty} \boldsymbol{\beta}$ for all $i \in \mathcal{I}$, then $\boldsymbol{\alpha} \not\prec_o^{\infty} \boldsymbol{\beta}$.

This axiom does *not* require $\boldsymbol{\alpha} \succ_o^{\infty} \boldsymbol{\beta}$; it simply requires the social planner not to form the *opposite* asymptotic preference to that of the individuals.

Minimal agreement. Suppose that each of the preference orders $\{\succcurlyeq_j\}_{j \in \mathcal{J}}$ has a GH representation (3) with an associated utility function $u_j : \mathcal{X} \longrightarrow \mathbb{R}$. We shall say that the utility functions $\{u_i\}_{i \in \mathcal{I}}$ satisfy *Minimal Agreement* if there exist probability measures μ_1 and μ_2 on $\mathcal{X}$ such that $\int_{\mathcal{X}} u_i \, \mathrm{d}\mu_1 > \int_{\mathcal{X}} u_i \, \mathrm{d}\mu_2$ for all $i \in \mathcal{I}$. In other words, there exist two "objective lotteries" over outcomes, for which all individuals have the same strict preference. Versions of this condition are widespread in the literature on Bayesian social aggregation; see e.g. Mongin (1995, 1998), Alon and Gayer (2016), or Danan et al. (2016).

Utilitarianism and weak utilitarianism. Recall that u_o is the ex post utility function associated to the social preference order $\succcurlyeq_o$. We shall say that u_o is *weakly utilitarian* if there exist constants $c_i \geqslant 0$ for all $i \in \mathcal{I}$ and $b \in \mathbb{R}$ such that

$$u_o \quad = \quad b + \sum_{i \in \mathcal{I}} c_i \, u_i. \tag{6}$$

It is possible that $c_i = 0$ for some $i \in \mathcal{I}$; thus, the preferences of some individuals might be ignored. If $c_i > 0$ for *all* $i \in \mathcal{I}$, then u_o is ***utilitarian***. Under mild conditions, weak utilitarianism is equivalent to utilitarianism (see Proposition D.2 in Appendix D). So we focus on establishing *weak* utilitarianism. We now come to our main result.

Theorem 1 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be consilient. For all $j \in \mathcal{J}$, suppose $\succcurlyeq_j$ has a compact, contiguous GH representation (3) with $\mathcal{P}_j \subseteq \mathcal{R}$. Assume that $\{u_i\}_{i\in\mathcal{I}}$ satisfy* Minimal Agreement*. Then $\succcurlyeq_o$ satisfies $\mathcal{R}$-*Almost-objective Pareto *if and only if u_o is weakly utilitarian.*

The next result applies this to the original problem of Bayesian social aggregation.

Corollary 1 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be consilient. For all $j \in \mathcal{J}$, suppose $\succcurlyeq_j$ has a contiguous SEU representation (2) with $\rho_j \in \mathcal{R}$. Suppose $\{u_i\}_{i\in\mathcal{I}}$ satisfy* Minimal Agreement*. Then $\succcurlyeq_o$ satisfies $\mathcal{R}$-*Almost-objective Pareto *if and only if u_o is weakly utilitarian.*

State-dependent utility. As noted in Section 1, Bayesian social aggregation may encounter difficulties when individuals have state-dependent utilities.[13] The simplest version of state-dependent utility supposes that the agent has the same utility function in all states, up to some state-dependent scalar multiplier. In other words, the agent's state-dependent utility function $v : \mathcal{S} \times \mathcal{X} \longrightarrow \mathbb{R}$ has the form

$$v(s, x) \quad = \quad w(s)\, u(x), \quad \text{for all } s \in \mathcal{S} \text{ and } x \in \mathcal{X}, \tag{7}$$

where $u : \mathcal{X} \longrightarrow \mathbb{R}$ and $w : \mathcal{S} \longrightarrow \mathbb{R}_+$ are bounded measurable functions. Heuristically, u is an underlying state-*independent* utility function, while w assigns more "weight" to this utility in some states than in others. Let $\succcurlyeq$ be a preference on $\mathcal{A}$. Given a state-dependent utility function like (7), a *state-dependent SEU representation* is a representation (1) where

$$V(\alpha) \quad = \quad \int_{\mathcal{S}} v\left(s, \alpha(s)\right)\, \mathrm{d}\rho[s] \quad = \quad \int_{\mathcal{S}} w(s)\, u\left(\alpha(s)\right)\, \mathrm{d}\rho[s], \quad \text{for all } \alpha \in \mathcal{A}. \tag{8}$$

Here is a state-dependent generalization of Corollary 1.

Corollary 2 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be strongly consilient. For all $j \in \mathcal{J}$, suppose $\succcurlyeq_j$ has a contiguous state-dependent SEU representation (8) for some $u_j : \mathcal{X} \longrightarrow \mathbb{R}$, $w_j : \mathcal{S} \longrightarrow \mathbb{R}_+$ and $\rho_j \in \mathcal{R}$. Assume that $\{u_i\}_{i\in\mathcal{I}}$ satisfy* Minimal Agreement*. Then $\succcurlyeq_o$ satisfies $\mathcal{R}$-*Almost-objective Pareto *if and only if u_o is weakly utilitarian.*

[13]One way reconcile the ex ante Pareto axiom with some form of social SEU maximization in an environment with heterogeneous beliefs is to introduce a state-dependent *social welfare function*; see e.g. Mongin (1998, Prop.6), Chambers and Hayashi (2006, Thm.1), Desai et al. (2018, Thm.4), Sprumont (2019), and Mongin and Pivato (2020, Thm.1). But the issue under discussion here is state-dependent *individual* utility, not state-dependent *social* utility.

Intrinsic consilience. A possible criticism of Theorem 1 and Corollaries 1 and 2 is that $\mathcal{R}$-Almost-objective Pareto involves an exogenous set $\mathcal{R}$ of probability measures. The next axiom endogenizes $\mathcal{R}$.

Almost-objective Pareto*. For all $j \in \mathcal{J}$, let $\succcurlyeq_j$ be a preference order on $\mathcal{A}$ with a GH representation (3) given by some set $\mathcal{P}_j \subseteq \Delta(\mathcal{S})$. Let $\mathcal{R} := \bigcup_{j \in \mathcal{J}} \mathcal{P}_j$.

If $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are compatible $\mathcal{R}$-almost-objective acts, and $\boldsymbol{\alpha} \succ_i^\infty \boldsymbol{\beta}$ for all $i \in \mathcal{I}$, then $\boldsymbol{\alpha} \not\prec_o^\infty \boldsymbol{\beta}$.

Combining Proposition 2 with Theorem 1 yields the following result:

Corollary 3 *For all $j \in \mathcal{J}$, suppose $\succcurlyeq_j$ has a compact, contiguous, nonatomic GH representation (3), and suppose $\{u_i\}_{i \in \mathcal{I}}$ satisfy* Minimal Agreement*. Then $\succcurlyeq_o$ satisfies* Almost-objective Pareto* *if and only if u_o is weakly utilitarian.*

(One can likewise obtain versions of Corollaries 1 and 2 using Almost-objective Pareto*.) The advantage of Corollary 3 over Theorem 1 is that the relevant Pareto axiom is defined "by the agents themselves", via their belief sets $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$. The *disadvantage* is that, to verify Almost-objective Pareto*, one must exactly identify the sets $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$. In contrast, to apply Theorem 1, one need only know that these sets are all contained in *some* consilient set $\mathcal{R}$.

Proof sketch. Recall that Harsanyi's (1955) original result involved expected-utility preferences over *objective* lotteries. In that setting, if $\succcurlyeq_o$ is *not* weakly utilitarian, then the Separating Hyperplane Theorem can be used to construct a pair of lotteries that violate Ex ante Pareto. By restricting the Pareto axiom to asymptotic preferences between almost-objective acts, we have restricted it to a domain where agents' preferences are *almost* described by such objective expected utilities. This is expressed precisely by the next result, which is also of independent interest.

Proposition 6 *Let $\mathcal{R}$ be a consilient set of probability measures on $\mathcal{S}$. Let $K \in \mathbb{N}$, let $\mathbf{q} \in \Delta^K$, let $\mathbf{x} \in \mathcal{X}^K$, and let $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ be an $\mathcal{R}$-almost-objective act subordinate to $(\mathbf{q}, \mathbf{x})$. Let V be a compact GH representation (3) with $\mathcal{P} \subseteq \mathcal{R}$. Then*

$$\lim_{n \to \infty} V(\alpha^n) \quad = \quad \sum_{k=1}^K q_k \, u(x_k). \tag{9}$$

By virtue of Proposition 6, a separating hyperplane argument can be applied to prove Theorem 1. Proposition 6 also has another important consequence: when considering an agent's asymptotic preferences over almost-objective acts, *all information about that agent's beliefs is effaced.* This explains why $\mathcal{R}$-Almost-objective Pareto *cannot* entail any link between individual beliefs and collective beliefs. We will turn to this question in the next section.

5 Collective beliefs

In this section, we shall assume *Minimal Agreement on Outcomes* (MAO): there exist $x, y \in \mathcal{X}$ such that $x \succ_j y$ for all $j \in \mathcal{J}$. Let us call the pair (x, y) a *dichotomy*. Let $\alpha : \mathcal{S} \longrightarrow \mathcal{X}$ be an act. Say that α is a *dichotomous act* if there is a dichotomy (x, y) such that $\alpha(s) \in \{x, y\}$ for all $s \in \mathcal{S}$. Two dichotomous acts α and β are *compatible* if they range over the same dichotomy $\{x, y\}$. Consider the following axiom:

Dichotomous Pareto. For any compatible dichotomous acts $\alpha, \beta \in \mathcal{A}$, if $\alpha \succcurlyeq_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \succcurlyeq_o \beta$.

The next result is derived from a result of Mongin (1995).

Proposition 7 *Suppose the preferences $\{\succcurlyeq_j\}_{j \in \mathcal{J}}$ all have SEU representations with nonatomic beliefs $\{\rho_j\}_{j \in \mathcal{J}}$, and they satisfy MAO. Then $\succcurlyeq_o$ satisfies* Dichotomous Pareto *if and only if ρ_o is a weighted average of $\{\rho_i\}_{i \in \mathcal{I}}$.*

Consistent with the philosophy of this paper, Proposition 7 decouples the problem of belief aggregation from that of utility aggregation: it determines the collective beliefs but says nothing about social welfare. But it only applies when all agents are SEU maximizers. Are there similar results for other ambiguity attitudes? In uncertain decision environments where all agents have the *same* utility function, the social aggregation of beliefs has been studied by Crès et al. (2011), Nascimento (2012), Gajdos and Vergnaud (2013) and Stanca (2021) for various ambiguity attitudes including maximin expected utility and second-order subjective expected utility. By restricting to dichotomous acts, Dichotomous Pareto simulates a world where all agents have the same utility function, so Proposition 7 is comparable to this literature. This raises the question of whether there is a version of Proposition 7 for GH preferences.

Unfortunately, the class of GH preferences is too broad to admit a result analogous to Proposition 7. Even if we restrict to compatible dichotomous acts, the inequalities (3) give agents too much freedom in how they derive their preferences from their beliefs. Thus, Dichotomous Pareto cannot forge a tight connection between the social belief set $\mathcal{P}_o$ and the individual belief sets $\{\mathcal{P}_i\}_{i \in \mathcal{I}}$. To characterize belief aggregation, we must switch from GH preferences to a kind of preference which, while less decisive, entails a much closer link between preferences and beliefs.

Bewley preferences. Let $\trianglerighteq$ be a preorder on $\mathcal{A}$ —that is, a transitive, reflexive (but possibly incomplete) binary relation. A *Bewley representation* for $\trianglerighteq$ is a pair $(\mathcal{P}, u)$, where $\mathcal{P} \subset \Delta(\mathcal{S})$ and $u : \mathcal{X} \longrightarrow \mathbb{R}$, such that for all $\alpha, \beta \in \mathcal{A}$,

$$\Big(\alpha \ \trianglerighteq \ \beta\Big) \iff \left(\int_{\mathcal{S}} u \circ \alpha \ \mathrm{d}\rho \geqslant \int_{\mathcal{S}} u \circ \beta \ \mathrm{d}\rho \text{ for all } \rho \in \mathcal{P}\right). \qquad (10)$$

If $\trianglerighteq$ has such a representation, then we shall call it a *Bewley preference*. As we explain at the end of this section, these are closely related to GH preferences. When restricted to *constant* acts, a Bewley preference defines a complete order on $\mathcal{X}$. So the property of Minimal Agreement on Outcomes, the definition of dichotomous acts, and the Dichotomous Pareto axiom are all meaningful for Bewley preferences.

Theorem 2 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be strongly consilient. For all $j \in \mathcal{J}$, suppose $\trianglerighteq_j$ has a Bewley representation (10) given by a compact subset $\mathcal{P}_j \subseteq \mathcal{R}$, and suppose these preferences satisfy MAO. Let $\overline{\mathcal{P}}$ be the closed convex hull of $\bigcup_{i \in \mathcal{I}} \mathcal{P}_i$. Then $\trianglerighteq_o$ satisfies* Dichotomous Pareto *if and only if $\mathcal{P}_o \subseteq \overline{\mathcal{P}}$.*

In the special case when all agents have SEU preferences, we have $\mathcal{P}_j = \{\rho_j\}$ for all $j \in \mathcal{J}$, so that Theorem 2 reduces to Proposition 7. Although it is not obvious from the statement, the proof of Theorem 2 depends heavily on almost-objective uncertainty; see Appendix C.

Bayesian social aggregation of Bewley preferences has previously been analysed by Danan et al. (2016). In particular, Danan et al.'s Theorem 2 shows that a certain Pareto axiom implies that $\mathcal{P}_o \subseteq \overline{\mathcal{P}}$. However, like Gilboa et al. (2004), the results of Danan et al. simultaneously characterize belief aggregation *and* utility aggregation, whereas we separate these problems. By combining $\mathcal{R}$-Almost-objective Pareto and Dichotomous Pareto, we can characterize both the social welfare function and social belief set using Theorems 1 and 2. But we can also choose to impose only one or the other of these axioms, thereby constraining *either* the social welfare function *or* the social belief set, while leaving the other unconstrained.

The link between Bewley and GH. Suppose that $\mathcal{X}$ is a convex space, as in the Anscombe-Aumann framework. For any $\alpha, \beta \in \mathcal{A}$ and $q \in [0,1]$, define $\alpha \oplus_q \beta \in \mathcal{A}$ by setting $(\alpha \oplus_q \beta)(s) := q\,\alpha(s) + (1-q)\,\beta(s)$ for all $s \in \mathcal{S}$. Let $\succcurlyeq$ be a preference relation on $\mathcal{A}$. The ***unambiguous preference*** induced by $\succcurlyeq$ is the binary relation $\trianglerighteq$ on $\mathcal{A}$ defined:

$$\Big(\alpha \;\trianglerighteq\; \beta\Big) \iff \Big(\alpha \oplus_q \gamma \succcurlyeq \beta \oplus_q \gamma, \text{ for all } \gamma \in \mathcal{A} \text{ and all } q \in (0,1]\Big).$$

If $\succcurlyeq$ is an *invariant biseparable* preference, then Ghirardato et al. (2004, Propositions 5 and 7) showed that there is a weak* compact, convex set $\mathcal{P} \subseteq \Delta(\mathcal{S})$ and a utility function $u : \mathcal{X} \longrightarrow \mathbb{R}$ that yield *both* a Bewley representation (10) for $\trianglerighteq$ *and* a generalized Hurwicz representation (3) for $\succcurlyeq$. This result was later generalized to MBA preferences by Cerreia-Vioglio et al. (2011, Propositions 2 and Corollary 3).

In fact, the Anscombe-Aumann framework is not necessary for these results. Let $\mathcal{X}$ be an abstract set of outcomes, and let $u : \mathcal{X} \longrightarrow \mathbb{R}$ be a utility function arising from some representation of $\succcurlyeq$. A function $\oplus : \mathcal{X} \times \mathcal{X} \times [0,1] \longrightarrow \mathcal{X}$ is a u-***subjective mixture*** operation if, for any $x, y \in \mathcal{X}$ and $q \in [0,1]$, we have $u(x \oplus_q y) = q\,u(x) + (1-q)\,u(y)$. If $\mathcal{X}$ was a convex set and u was an affine utility function, then one could just define $x \oplus_q y := q\,x + (1-q)\,y$. But even when $\mathcal{X}$ is arbitrary, one can define subjective mixtures for a large class of preference relations, as follows. Let $V : \mathcal{A} \longrightarrow \mathbb{R}$ be a representation of $\succcurlyeq$, and define $u : \mathcal{X} \longrightarrow \mathbb{R}$ by restricting V to constant functions. For any event $\mathcal{E} \subseteq \mathcal{S}$ and outcomes $x, y \in \mathcal{X}$, let $x\mathcal{E}y$ be the dichotomous act that delivers x for all states in $\mathcal{E}$ and y

for all states outside $\mathcal{E}$. Say that $\succcurlyeq$ is ***locally biseparable*** at $\mathcal{E}$ if there is some $p \in (0,1)$ such that for any $x, y \in \mathcal{X}$ with $x \succcurlyeq y$, we have $V(x\mathcal{E}y) = p\,u(x) + (1-p)\,u(y)$. In other words, when restricted to bets on $\mathcal{E}$, V behaves like expected utility. Generalizing the work of Ghirardato et al. (2003), Ghirardato and Pennesi (2020) have shown that if $\succcurlyeq$ has even one locally biseparable event, then one can define a u-subjective mixture operation on $\mathcal{X}$. As noted by Ghirardato and Pennesi (2020, Remark 1), the results of Ghirardato et al. (2004) and Cerreia-Vioglio et al. (2011) described in the previous paragraph can then be extended to any monotone, locally biseparable preference using this subjective mixture operation, yielding combined GH/Bewley representations for $\succcurlyeq$ and $\trianglerighteq$.[14]

Thus, there is a close connection between Bewley and GH representations in the realm of MBA preferences and monotone, locally biseparable preferences. More generally, given *any* preference $\trianglerighteq$ with a Bewley representation (10), Danan et al. (2016, Prop.2) show that any transitive, Archimedean completion of $\trianglerighteq$ has a GH representation (3) using the same set $\mathcal{P}$ of beliefs.[15] In light of these results, one can interpret Theorem 2 as describing the aggregation of the subjective beliefs in the GH representations that appeared in earlier parts of this paper.

6 Discussion

We have considered a decision environment of radical uncertainty, in which the ex ante preferences of each agent admit generalized Hurwicz representation. We have introduced a very weak Pareto axiom, which applies only to asymptotic preferences along a sequence of acts for which all possible probabilistic beliefs entertained by all agents converge to the same limit. We have shown that social preferences satisfy this weak Pareto axiom if and only if the ex post social welfare function is a weighted sum of the ex post utility functions of the individuals. In other words, social preferences must be ex post utilitarian. A different Pareto axiom characterizes the formation of collective beliefs. Importantly, these results separate utility aggregation from belief aggregation, and they do not impose any relationship between collective ambiguity attitudes and individual ambiguity attitudes. As explained in Section 1, we see this as an advantage. We will now discuss the relationship between our results and the watershed paper of Gilboa et al. (2004).

For all $j \in \mathcal{J}$, suppose that $\succcurlyeq_j$ has a GH representation (3) with belief set $\mathcal{P}_j$. Let $\mathfrak{G} = \{\mathcal{G}_1, \ldots, \mathcal{G}_K\}$ be a partition of $\mathcal{S}$. Let us say that $\mathfrak{G}$ is a ***consensus partition*** if there is some $\mathbf{q} \in \Delta^K$ such that $\rho(\mathcal{G}_k) = q_k$ for all $k \in [1 \ldots K]$, all $\rho \in \mathcal{P}_j$, and all $j \in \mathcal{J}$ —in other words, all agents *exactly agree* on the probabilities of all elements of $\mathfrak{G}$. Gilboa et al. (2004) proposed a version of the following axiom:

Restricted Pareto. Let $\alpha, \beta \in \mathcal{A}$ be measurable with respect to a consensus partition $\mathfrak{G}$. If $\alpha \approx_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \approx_o \beta$. If $\alpha \succ_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \succ_o \beta$.[16]

[14] However, different agents generally have *different* subjective mixture operations. So unlike almost-objective uncertainty, subjective mixtures cannot be used for Bayesian social aggregation.

[15] Danan et al. (2016) refer to GH preferences as *variable caution rules*.

[16] Gilboa et al. used only the "indifference" part of the axiom. Also, they assumed all agents had SEU

This seems quite similar to Almost-objective Pareto. Indeed, if $\mathfrak{G}$ is a consensus partition, and we define $\mathfrak{G}^n := \mathfrak{G}$ for all $n \in \mathbb{N}$, then the sequence $(\mathfrak{G}^n)_{n=1}^{\infty}$ is trivially an "almost-objective" sequence with respect to the family $\{\rho_j\}_{j \in \mathcal{J}}$. Thus, if α and β are measurable with respect to $\mathfrak{G}$, and we define $\alpha^n := \alpha$ and $\beta^n := \beta$ for all $n \in \mathbb{N}$, then the sequences $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^{\infty}$ and $\boldsymbol{\beta} := (\beta^n)_{n=1}^{\infty}$ are compatible almost-objective acts. Thus, any unanimous preference which is admissible as input to Restricted Pareto is also admissible to Almost-objective Pareto, except that our axiom accepts a larger variety of inputs, and yields a weaker conclusion.

Let us say that a GH representation (3) is ***polytopic*** if the set $\mathcal{P}$ is a *polytope* —i.e. the convex hull of a *finite* subset of $\Delta(\mathcal{S})$. Gilboa et al. (2004) worked with SEU preferences. But their result has the following generalization to GH representations.

Theorem 3 *For all $j \in \mathcal{J}$, suppose that $\succcurlyeq_j$ has a nonatomic, polytopic GH representation (3) with belief set $\mathcal{P}_j$. Let $\overline{\mathcal{P}}$ be the convex hull of $\bigcup_{i \in \mathcal{I}} \mathcal{P}_i$. Then $\succcurlyeq_o$ satisfies* Restricted Pareto *if and only if u_o is weakly utilitarian and $\mathcal{P}_o \subseteq \overline{\mathcal{P}}$.*

In light of Theorem 3, it might seem that this paper has just deployed a lot of machinery to obtain a variation of a result that Gilboa et al. (2004) already achieved by much simpler means. But there are three important differences between $\mathcal{R}$-Almost objective Pareto and Restricted Pareto. First, our Pareto axiom allows us to decouple the ethical problem of specifying the SWF from the epistemic problem of collective belief formation (as in Theorems 1 and 2), whereas in Theorem 3 they are inextricably combined. Second, Restricted Pareto suffers from the same weakness as Almost-objective Pareto*, as remarked after Corollary 3: to apply Restricted Pareto in a particular situation, we must be able to recognize consensus partitions, which requires precise knowledge of the sets $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$ —something which may be difficult to achieve in practice. In contrast, to apply $\mathcal{R}$-Almost-objective Pareto, we need only know that $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$ are contained in some broad family of probability measures. It is possible to determine whether a partition sequence is $\mathcal{R}$-almost-objectively uncertain without knowing anything about $\{\mathcal{P}_j\}_{j \in \mathcal{J}}$, and also possible to construct such partition sequences on demand (e.g. using the methods of Appendix A).

Third, as agents acquire more information and Bayes-update their beliefs, *different* partitions of $\mathcal{S}$ will become consensus partitions. Thus, the scope of application of Restricted Pareto will shift as the information available to the agents changes. Mongin and Pivato (2020, §6, p. 649) show that this makes Restricted Pareto vulnerable to a kind of "spurious unanimity" phenomenon: different agents might "spuriously" assign the same probabilities to the cells of a partition because they receive different information. This can lead Restricted Pareto to make recommendations which are obviously incorrect in light of the aggregate information of the entire group. Mongin and Pivato refer to this as *complementary ignorance*.

$\mathcal{R}$-Almost-objective Pareto is much less vulnerable to complementary ignorance. To see this, suppose $\mathcal{R}$ is strongly consilient, and $\succcurlyeq$ has a GH representation V with utility

representations, so $\mathcal{P}_j$ was a singleton for all $j \in \mathcal{J}$. Hence, their definition of "consensus partition" is simpler.

function u and belief set $\mathcal{P} \subseteq \mathcal{R}$. Let $\mathcal{E} \subseteq \mathcal{S}$ be an event which gets positive probability from all elements of $\mathcal{P}$, and let $\mathcal{P}'$ be obtained by Bayes-updating every element of $\mathcal{P}$ by $\mathcal{E}$. Suppose $\succcurlyeq'$ is another preference, having a GH representation V' with the utility function u and belief set $\mathcal{P}'$; this could be the updated preferences of the $\succcurlyeq$-agent upon learning $\mathcal{E}$.[17] If $\boldsymbol{\alpha}$ is any almost-objective act, then Proposition 6 says $\lim_{n\to\infty} V(\alpha^n) = \lim_{n\to\infty} V'(\alpha^n)$. Thus, $\succcurlyeq$ and $\succcurlyeq'$ have *exactly the same asymptotic preferences* over almost-objective acts.

Now suppose we have a collection $\{\succcurlyeq_j\}_{j\in\mathcal{J}}$ of GH preferences and a collection $\{\mathcal{E}_j\}_{j\in\mathcal{J}}$ of events. For all $j \in \mathcal{J}$, let $\succcurlyeq'_j$ be a GH preference obtained by Bayes-updating $\succcurlyeq_j$ with $\mathcal{E}_j$, as in the previous paragraph. Since the asymptotic preferences of each agent are unchanged by these updates, it follows that $\mathcal{R}$-Almost-objective Pareto will apply to $\{\succcurlyeq'_j\}_{j\in\mathcal{J}}$ in exactly the same situations as it applies to $\{\succcurlyeq_j\}_{j\in\mathcal{J}}$. In other words, unlike Restricted Pareto, it is impossible to induce "spurious" instances of Almost-objective Pareto by exposing different agents to different information.

The distinction between Gilboa et al. (2004) and the present paper is analogous to the distinction between universal and existential quantifiers.[18] The Restricted Pareto axiom says that for *any* source of uncertainty, if all agents happen to share the same beliefs about that source (for whatever reason), then the ex ante Pareto axiom should apply to preferences over acts contingent on that source. But to achieve utilitarian aggregation *à la* Harsanyi, we don't need to quantify over *every* source of such "common-belief uncertainty". It suffices to apply the ex ante Pareto axiom to *one* source of common-belief uncertainty.

In the models of Mongin and Pivato (2020) and Pivato (2022), this source of common-belief uncertainty is either exogenous, or the asymptotic outcome of a learning process. In the first paper, there is an exogenous distinction between two sources of uncertainty: one "subjective" and one "objective". If social preferences satisfy ex ante Pareto *only* for the objective source, then the social utility function is utilitarian and all agents have SEU preferences with the same beliefs about the objective source, but there is no relationship between their beliefs regarding the subjective source. In the second paper, all agents have SEU preferences, and there is an infinite stream of information arriving over time, from which all agents update their beliefs, and hence their preferences over acts. If social preferences satisfy ex ante Pareto *only* for unanimous preferences which persist in the long term under this learning process, then the social utility function must be utilitarian, but no relationship is required between the original beliefs of the agents, except for a weak condition called *concordance* (roughly: the supports of their beliefs must have a common overlap).

In the present paper, the source of common-belief uncertainty is the almost-objective uncertainty introduced in Section 3. Unlike Mongin and Pivato (2020), this source is not exogenous. Unlike Pivato (2022), it does not arise from a dynamical process, and does not require any compatibility between the beliefs of different agents (their beliefs could even have pairwise disjoint support). But like these two papers, and unlike Gilboa et al. (2004), this focus on a unique source of common-belief uncertainty allows us to cleanly separate utility-aggregation from belief-aggregation.

[17] Note that we do not impose any other relationship between V and V'.

[18] We thank a referee for this apt comparison.

The simplest way to introduce common-belief uncertainty is to switch from the Savage framework of the present paper (and the three papers just discussed) to the Anscombe-Aumann framework, in which acts are functions from $\mathcal{S}$ into $\Delta(\mathcal{X})$, the space of *objective lotteries* over $\mathcal{X}$. For all $j \in \mathcal{J}$, suppose that $\succcurlyeq_j$ has a GH representation with an affine utility function $u_j : \Delta(\mathcal{X}) \longrightarrow \mathbb{R}$ (as in Proposition 5 of Cerreia-Vioglio et al., 2011). Let us call this a *von Neumann-Morgenstern* (vNM) GH representation. Let $\mathcal{A}_c$ be the set of all *constant* acts. Consider the following axiom

Lottery Pareto. For any $\alpha, \beta \in \mathcal{A}_c$, if $\alpha \approx_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \approx_o \beta$. If $\alpha \succ_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \succ_o \beta$.

The next result follows immediately from Harsanyi's (1955) original theorem:

Proposition 8 *Suppose $\{\succcurlyeq_j\}_{j \in \mathcal{J}}$ all have vNM GH representations. Then $\succcurlyeq_o$ satisfies* Lottery Pareto *if and only if u_o is weakly utilitarian.*

Like Mongin and Pivato (2020), Pivato (2022) and the present paper, Proposition 8 imposes no constraints on beliefs. So it seems like a simple and elegant solution. But this approach is not entirely satisfactory, because the Anscombe-Aumann framework assumes the existence of objective lotteries. Adherents of the "subjectivist" or "personalist" view of probability (e.g. de Finetti, Savage) deny that such objective lotteries even exist.[19] The canonical examples of "objective lotteries" are physical devices like roulette wheels, dice throws, coin tosses, and Galton boards. But classical physics is deterministic, so these macroscopic devices are not actually random —any *apparent* randomness comes from our imprecise information about the initial conditions of the device. Thus, it is not objective, it is *subjective*: observers with different information about the initial conditions may form different probabilistic beliefs. (And an observer with *perfect* information could simply *predict* the outcome with probability one, like Laplace's Demon.)[20] Furthermore, a rational agent could always subjectively believe that these devices were systematically biased.

Nevertheless, we intuit that these devices somehow "approximate" objective randomness. And the faster we spin the roulette wheel, the more vigorously we throw the die, or the more elaborate the array of pins in the Galton board, the better these approximations become. Making this intuition mathematically respectable was precisely Poincaré's (1912) motivation for introducing almost-objective uncertainty. Thus, Almost-objective Pareto is precisely the form that Lottery Pareto takes in a world where "objective lotteries" arise from deterministic devices as described above.

[19]See §4.9 of Pivato (2022) for further discussion of this point. Similar concerns about objective lotteries motivated Ghirardato et al. (2003) and Ghirardato and Pennesi (2020, 2022) to develop the machinery of *subjective mixtures* described at the end of Section 5, as well as motivating Gilboa (1987) and Casadesus-Masanell et al. (2000a,b) to re-characterize Choquet expected utility and maximin expected utility in the Savage framework, rather than the Anscombe-Aumann framework.

[20]*Genuinely* random events might arise from quantum phenomena. More broadly, the metaphysics of determinism vs. genuine randomness is a subject of ongoing investigation in philosophy. But the correct interpretation of quantum mechanics is still open to debate. And the resolution of problems in normative economics should not depend on foundational questions in physics or metaphysics.

A Proofs from Section 3

Proof of Proposition 2. Let $\{\mu_n\}_{n=1}^{\infty}$ be a countable dense subset of $\mathcal{R}$. Let $\mathbf{q} \in \Delta^K$. For all $n \in \mathbb{N}$, the Dubins-Spanier theorem yields a partition $\mathfrak{G}^n = \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$ such that $\mu_m(\mathcal{G}_k^n) = q_k$ for all $k \in [1 \ldots K]$ and all $m \in [1 \ldots n]$ (because $\mu_1, \ldots, \mu_n$ are all nonatomic). We claim that the sequence $(\mathfrak{G}^n)_{n=1}^{\infty}$ is almost-objectively random and subordinate to $\mathbf{q}$.

To see this, let $\rho \in \mathcal{R}$ and let $\epsilon > 0$. Since $\{\mu_n\}_{n=1}^{\infty}$ is dense in the norm topology, there exists $N \in \mathbb{N}$ such that $\|\mu_N - \rho\| < \epsilon$. Now, let $k \in [1 \ldots K]$. For any $n \geqslant N$, we have $\mu_N(\mathcal{G}_k^n) = q_k$, by the definition of $\mathfrak{G}^n$, while $|\rho(\mathcal{G}_k^n) - \mu_N(\mathcal{G}_k^n)| < \epsilon$ because $\|\mu_N - \rho\|_{\text{vr}} < \epsilon$. Thus, $|\rho(\mathcal{G}_k^n) - q_k| < \epsilon$, for all $n \geqslant N$. This works for any $\epsilon > 0$; thus $\lim_{n\to\infty} \rho(\mathcal{G}_k^n) = q_k$. This works for all $k \in [1 \ldots K]$, and all $\rho \in \mathcal{R}$. □

Proposition 3(a) follows immediately from Proposition 2 and the next lemma.

Lemma A.1 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$, If $\mathcal{R}$ is separable in the norm topology, and every element of $\mathcal{R}$ is separable, then $\langle\mathcal{R}\rangle$ is separable in the norm topology.*

Proof: Let $\{\nu_n\}_{n=1}^{\infty}$ be a countable dense subset of $\mathcal{R}$. For all $n \in \mathbb{N}$, let $\Delta(\mathcal{S}, \nu_n) := \{\phi \in \mathcal{L}^1(\mathcal{S}, \nu_n);\ \ \phi \geqslant 0 \text{ and } \int_{\mathcal{S}} \phi \, \mathrm{d}\nu_n = 1\}$; in other words, $\Delta(\mathcal{S}, \nu_n) = \{\frac{\mathrm{d}\rho}{\mathrm{d}\nu_n};\ \ \rho \in \Delta(\mathcal{S}) \text{ and } \rho \ll \mu\}$. Recall that the Banach space $\mathcal{L}^1(\mathcal{S}, \nu_n)$ is separable (because all elements of $\mathcal{R}$ are separable probability measures). Thus, the subset $\Delta(\mathcal{S}, \nu_n)$ is also separable in the L^1 norm. So let $\{\psi_n^m\}_{m=1}^{\infty}$ be a countable dense subset of $\Delta(\mathcal{S}, \nu_n)$. For all $m \in \mathbb{N}$, let $\lambda_n^m \in \Delta(\mathcal{S})$ be the probability measure such that $\frac{\mathrm{d}\lambda_n^m}{\mathrm{d}\nu_n} = \psi_n^m$. Note that λ_n^m is nonatomic because ν_n is nonatomic.

We claim that the countable set $\{\lambda_n^m\}_{m,n=1}^{\infty}$ is dense in $\langle\mathcal{R}\rangle$. in the total variation norm. To see this, let $\mu \in \langle\mathcal{R}\rangle$. Then there exists $\rho \in \mathcal{R}$ such that $\mu \ll \rho$, and if $\phi := \frac{\mathrm{d}\mu}{\mathrm{d}\rho}$, then there exists $C > 0$ such that $0 \leqslant \phi(s) < C$ for all $s \in \mathcal{S}$. Let $\epsilon > 0$. Since $\{\nu_n\}_{n=1}^{\infty}$ is dense in $\mathcal{R}$, there exists $n \in \mathbb{N}$ such that $\|\nu_n - \rho\|_{\text{vr}} < \epsilon/2C$. Automatically, $\phi \in \mathcal{L}^1(\mathcal{S}, \nu_n)$, because ϕ is bounded. Thus, there exists $m \in \mathbb{N}$ such that $\|\phi - \psi_n^m\|_{1,\nu_n} < \epsilon/2$, where this refers to the L^1 norm on $\mathcal{L}^1(\mathcal{S}, \nu_n)$. We will show that $\|\lambda_n^m - \mu\|_{\text{vr}} < \epsilon$.

To see this, let $\mathcal{H}_1, \ldots, \mathcal{H}_J \subseteq \mathcal{S}$ be disjoint and measurable. For all $j \in [1 \ldots J]$,

$$\left|\lambda_n^m(\mathcal{H}_j) - \mu(\mathcal{H}_j)\right| \underset{(*)}{=} \left|\int_{\mathcal{H}_j} \psi_n^m \, \mathrm{d}\nu_n - \int_{\mathcal{H}_j} \phi \, \mathrm{d}\rho\right| \tag{A1}$$

$$\leqslant \left|\int_{\mathcal{H}_j} \psi_n^m \, \mathrm{d}\nu_n - \int_{\mathcal{H}_j} \phi \, \mathrm{d}\nu_n\right| + \left|\int_{\mathcal{H}_j} \phi \, \mathrm{d}\nu_n - \int_{\mathcal{H}_j} \phi \, \mathrm{d}\rho\right|$$

$$= \left|\int_{\mathcal{H}_j} (\psi_n^m - \phi) \, \mathrm{d}\nu_n\right| + \left|\int_{\mathcal{H}_j} \phi \, \mathrm{d}(\nu_n - \rho)\right| \leqslant \int_{\mathcal{H}_j} \left|\psi_n^m - \phi\right| \, \mathrm{d}\nu_n + \int_{\mathcal{H}_j} |\phi| \, \mathrm{d}|\nu_n - \rho|,$$

where $(*)$ is because $\psi_n^m = \dfrac{\mathrm{d}\lambda_n^m}{\mathrm{d}\nu_n}$ and $\phi = \dfrac{\mathrm{d}\mu}{\mathrm{d}\rho}$. Thus, if $\mathcal{H} := \bigsqcup_{j=1}^J \mathcal{H}_j$, then

$$\begin{aligned}\sum_{j=1}^J |\lambda_n^m(\mathcal{H}_j) - \mu[\mathcal{H}_j]| \quad &\underset{(*)}{\leqslant} \quad \sum_{j=1}^J \int_{\mathcal{H}_j} \left|\psi_n^m - \phi\right| \, \mathrm{d}\nu_n + \sum_{j=1}^J \int_{\mathcal{H}_j} |\phi| \, \mathrm{d}|\nu_n - \rho| \\ = \quad \int_{\mathcal{H}} \left|\psi_n^m - \phi\right| \, \mathrm{d}\nu_n + \int_{\mathcal{H}} |\phi| \, \mathrm{d}|\nu_n - \rho| \quad &\leqslant \quad \int_{\mathcal{S}} \left|\psi_n^m - \phi\right| \, \mathrm{d}\nu_n + \int_{\mathcal{S}} |\phi| \, \mathrm{d}|\nu_n - \rho| \\ \leqslant \quad \|\psi_n^m - \phi\|_{1,\nu_n} + C \cdot \|\nu_n - \rho\|_{\mathrm{vr}} \quad < \quad \frac{\epsilon}{2} + C \cdot \frac{\epsilon}{2\,C} \quad &= \quad \epsilon,\end{aligned}$$

where $(*)$ is by inequality (A1). This works for any disjoint collection $\mathcal{H}_1, \ldots, \mathcal{H}_J \subseteq \mathcal{S}$, so from definition (4) we conclude that $\|\lambda_n^m - \mu\|_{\mathrm{vr}} \leqslant \epsilon$. This argument works for any $\epsilon > 0$, and any $\mu \in \langle\mathcal{R}\rangle$. Thus, $\{\lambda_n^m\}_{m,n=1}^\infty$ is dense in $\langle\mathcal{R}\rangle$. □

The proof of Proposition 3(b) is somewhat more involved, and requires an auxiliary concept and four preliminary lemmas. Recall that in Proposition 3(b), $\mathcal{S}$ was assumed to be a standard Borel space —that is, it is measurably isomorphic to a complete separable metric space endowed with its Borel sigma algebra. Therefore, without loss of generality we will sometimes assume in the following material that $\mathcal{S}$ is endowed with a metric d that makes it a complete separable metric space, and the sigma algebra on $\mathcal{S}$ is the resulting Borel sigma algebra.

For any $\mathcal{Y} \subseteq \mathcal{S}$, the ***diameter*** of $\mathcal{Y}$ is defined: $\mathrm{diam}(\mathcal{Y}) := \sup_{s,t\in\mathcal{Y}} d(s,t)$. For any $\epsilon > 0$, an ϵ-***partition*** is a collection $\mathfrak{Y} = \{\mathcal{Y}_n\}_{n=1}^N$ of disjoint measurable subsets of $\mathcal{S}$ (for some $N \in \mathbb{N} \cup \{\infty\}$) such that $\bigsqcup_{n=1}^N \mathcal{Y}_n = \mathcal{S}$, and $\mathrm{diam}(\mathcal{Y}_n) \leqslant \epsilon$ for all $n \in [1 \ldots N]$.[21]

Lemma A.2 *Let $(\mathcal{S}, d)$ be any metric space. Then $(\mathcal{S}, d)$ is separable if and only if it admits an ϵ-partition for all $\epsilon > 0$.*

Proof: "$\Longrightarrow$" Let $\{s_n\}_{n=1}^\infty$ be a countable dense subset of $\mathcal{S}$. Let $\epsilon > 0$. For all $s \in \mathcal{S}$, let $\mathcal{B}(s, \epsilon)$ be the open ball of radius $\frac{\epsilon}{2}$ around s. For all $N \in \mathbb{N}$, let $\mathcal{Y}_N := \mathcal{B}(s_N, \epsilon) \setminus \bigcup_{n=1}^{N-1} \mathcal{B}(s_n, \epsilon)$; then $\mathrm{diam}(\mathcal{Y}_N) \leqslant \epsilon$. Thus, $\{\mathcal{Y}_n\}_{n=1}^\infty$ is an ϵ-partition of $\mathcal{S}$.

"$\Longleftarrow$" For all $m \in \mathbb{N}$, let $\mathfrak{Y}^m = \{\mathcal{Y}_n^m\}_{n=1}^\infty$ be a $(\frac{1}{m})$-partition. For all $(n,m) \in \mathbb{N}^2$, let $s_{n,m} \in \mathcal{Y}_n^m$. Then $\{s_{n,m}\}_{n,m=1}^\infty$ is a countable dense subset of $\mathcal{S}$. □

Let $\mathcal{P}$ be a collection of Borel probability measures on $\mathcal{S}$, let $K \in \mathbb{N}$, and let $\mathbf{q} = (q_1, \ldots, q_K) \in \Delta^K$. A $\mathbf{q}$-***Poincaré sequence*** for $\mathcal{P}$ is a sequence $\{(\mathfrak{G}^n, \mathfrak{Y}^n, \epsilon_n)\}_{n=1}^\infty$, where for all $n \in \mathbb{N}$, $\mathfrak{G}^n = \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$ is a K-element measurable partition of $\mathcal{S}$, $\epsilon_n > 0$ and $\mathfrak{Y}^n$ is an ϵ_n-partition, such that

[21] Note that we allow these partitions to have a countably infinite number of elements. This is necessary because $\mathcal{S}$ is not necessarily compact.

- $\lim_{n\to\infty} \epsilon_n = 0$.
- For all $\rho \in \mathcal{P}$, there exists $N \in \mathbb{N}$ such that for all $n \geqslant N$, all $k \in [1\ldots K]$, and all $\mathcal{Y} \in \mathfrak{Y}^n$, $\rho[\mathcal{G}_k^n \cap \mathcal{Y}] = q_k\,\rho[\mathcal{Y}]$ (and thus, $\rho[\mathcal{G}_k^n] = q_k$).

Example. Let $\mathcal{S} := [0,1)$. Let $\mathcal{P} := \{\lambda\}$ where λ is the Lebesgue measure. Let $\mathbf{q} = (\frac{1}{2}, \frac{1}{2})$. For all $n \in \mathbb{N}$, let $\epsilon := 1/2^n$ and let $\mathfrak{Y}^n := \{\mathcal{Y}_1^n, \ldots, \mathcal{Y}_{2^n}^n\}$ where $\mathcal{Y}_k^n := [\frac{k-1}{2^n}, \frac{k}{2^n})$ for all $k \in [1\ldots 2^n]$. Finally, let $\mathfrak{G}^n := \{\mathcal{G}_1^n, \mathcal{G}_2^n\}$, where

$$\mathcal{G}_1^n \quad := \quad \bigcup_{\substack{k=1\\ k \text{ odd}}}^{2^{n+1}-1} \mathcal{Y}_k^{n+1} \quad \text{and} \quad \mathcal{G}_2^n \quad := \quad \bigcup_{\substack{k=2\\ k \text{ even}}}^{2^{n+1}} \mathcal{Y}_k^{n+1}.$$

Then $\{(\mathfrak{G}^n, \mathfrak{Y}^n, \epsilon_n)\}_{n=1}^\infty$ is a $(\frac{1}{2}, \frac{1}{2})$-Poincaré sequence for $\{\lambda\}$.

Lemma A.3 *Let $(\mathcal{S}, d)$ be any separable metric space. Let $\mathcal{H} \subseteq \mathcal{M}(\mathcal{S})$ be a countable collection of nonatomic signed measures on $\mathcal{S}$. Let $\mathcal{F}$ be the linear subspace of $\mathcal{M}(\mathcal{S})$ consisting of all finite linear combinations of elements of $\mathcal{H}$. Let $\mathcal{P} \subseteq \mathcal{F}$ be the set of all probability measures in $\mathcal{F}$. Then for all $K \in \mathbb{N}$ and all $\mathbf{q} \in \Delta^K$, $\mathcal{P}$ has a $\mathbf{q}$-Poincaré sequence.*

Proof: Suppose that $\mathcal{H} = \{\eta_n\}_{n=1}^\infty$. For all $n \in \mathbb{N}$, the Hahn-Jordan Decomposition Theorem says that $\eta_n = \eta_n^+ - \eta_n^-$, where η_n^+ and η_n^- are either zero or positive measures. They are nonatomic because η_n is nonatomic. Thus, by replacing $\{\eta_n\}_{n=1}^\infty$ with $\{\eta_n^\pm\}_{n=1}^\infty$ if necessary, we can assume without loss of generality that all elements of $\mathcal{H}$ are positive, nonatomic measures.

Let $\{\epsilon_n\}_{n=1}^\infty$ be a positive sequence with $\lim_{n\to\infty} \epsilon_n = 0$. For all $N \in \mathbb{N}$, Lemma A.2 says $\mathcal{S}$ has an ϵ_N-partition $\mathfrak{Y}^N$.

Claim 1: *For all $N \in \mathbb{N}$, and all $\mathcal{Y} \in \mathfrak{Y}^N$, there is a measurable partition $\{\mathcal{G}_1^\mathcal{Y}, \ldots, \mathcal{G}_K^\mathcal{Y}\}$ of $\mathcal{Y}$ such that $n \in [1\ldots N]$, we have*

$$\eta_n(\mathcal{G}_k^\mathcal{Y}) \quad = \quad q_k \cdot \eta_n(\mathcal{Y}), \quad \textit{for all } k \in [1\ldots K]. \tag{A2}$$

Proof: Let $n \in [1\ldots N]$. If $\eta_n(\mathcal{Y}) = 0$, then the equations (A2) are trivially satisfied for *any* partition $\{\mathcal{G}_\mathcal{Y}^1, \ldots, \mathcal{G}_\mathcal{Y}^K\}$. So, let $\mathcal{N} := \{n \in [1\ldots N];\ \ \eta_n(\mathcal{Y}) > 0\}$; it suffices to construct a partition satisfying the equations (A2) for all $n \in \mathcal{N}$. For all $n \in \mathcal{N}$, let $\widetilde{\eta}_n$ be the nonatomic probability measure on $\mathcal{Y}$ defined by setting $\widetilde{\eta}_n(\mathcal{U}) := \eta_n(\mathcal{U})/\eta_n(\mathcal{Y})$ for all measurable $\mathcal{U} \subseteq \mathcal{Y}$. Thus $\{\widetilde{\eta}_n\}_{n\in\mathcal{N}}$ is a finite collection of nonatomic probability measures, so the Dubins-Spanier Theorem yields a partition $\{\mathcal{G}_1^\mathcal{Y}, \ldots, \mathcal{G}_K^\mathcal{Y}\}$ of $\mathcal{Y}$ such that

$$\widetilde{\eta}_n(\mathcal{G}_k^\mathcal{Y}) \quad = \quad q_k \quad \text{for all } k \in [1\ldots K] \text{ and } n \in \mathcal{N}. \tag{A3}$$

(Aliprantis and Border, 2006, Theorem 13.34, p.478). For all $n \in \mathcal{N}$, multiply both sides of equation (A3) by $\eta_n(\mathcal{Y})$ to obtain equation (A2). ◇ `Claim 1`

Fix $N \in \mathbb{N}$, and apply Claim 1 to all $\mathcal{Y} \in \mathfrak{Y}^N$. Observe that the sets in the collection $\{\mathcal{G}_k^{\mathcal{Y}};\ \mathcal{Y} \in \mathfrak{Y}^N \text{ and } k \in [1 \ldots K]\}$ are all disjoint. For all $k \in [1 \ldots K]$, define

$$\mathcal{G}_k^N \quad := \quad \bigsqcup_{\mathcal{Y} \in \mathfrak{Y}^N} \mathcal{G}_k^{\mathcal{Y}}. \tag{A4}$$

Then $\{\mathcal{G}_1^N, \ldots, \mathcal{G}_K^N\}$ is a measurable partition of $\mathcal{S}$: these sets are are disjoint, and

$$\bigsqcup_{k=1}^{K} \mathcal{G}_k^N \quad = \quad \bigsqcup_{k=1}^{K} \left(\bigsqcup_{\mathcal{Y} \in \mathfrak{Y}^N} \mathcal{G}_k^{\mathcal{Y}} \right) \quad = \quad \bigsqcup_{\mathcal{Y} \in \mathfrak{Y}^N} \left(\bigsqcup_{k=1}^{K} \mathcal{G}_k^{\mathcal{Y}} \right) \quad = \quad \bigsqcup_{\mathcal{Y} \in \mathfrak{Y}^N} \mathcal{Y} \quad = \quad \mathcal{S}.$$

Furthermore, for all $\mathcal{Y} \in \mathfrak{Y}^N$, we have $\mathcal{G}_k^N \cap \mathcal{Y} = \mathcal{G}_k^{\mathcal{Y}}$ for all $k \in [1 \ldots K]$; thus, for all $n \in [1 \ldots N]$,

$$\eta_n(\mathcal{G}_k^N \cap \mathcal{Y}) \quad = \quad \eta_n(\mathcal{G}_k^{\mathcal{Y}}) \underset{(*)}{=} q_k\, \eta_n(\mathcal{Y}), \tag{A5}$$

where $(*)$ is by equation (A2).

Now, let $\rho \in \mathcal{P}$. Then there exists some $N \in \mathbb{N}$ such that ρ is a linear combination of $\eta_1, \ldots, \eta_N$. Thus, for any $n \geqslant N$, ρ is also a linear combination of $\eta_1, \ldots, \eta_n$ (with zero coefficients for $\eta_{N+1}, \ldots, \eta_n$). Thus, for all $\mathcal{Y} \in \mathfrak{Y}^n$ and all $k \in [1 \ldots K]$, equation (A5) yields $\rho[\mathcal{G}_k^n \cap \mathcal{Y}] = q_k\, \rho[\mathcal{Y}]$, as desired. □

Lemma A.4 *Suppose $(\mathcal{S}, d)$ is a complete, separable metric space. Let $K \in \mathbb{N}$, let $\mathbf{q} \in \Delta^K$, Let $\mathcal{P}$ be a collection of probability measures on $\mathcal{S}$, and let $\{(\mathfrak{G}^n, \mathfrak{Y}^n, \epsilon_n)\}_{n=1}^{\infty}$ be a $\mathbf{q}$-Poincaré sequence for $\mathcal{P}$. Let $\mathcal{L} = \langle \mathcal{P} \rangle$. Then $(\mathfrak{G}^n)_{n=1}^{\infty}$ is $\mathcal{L}$-almost-objectively uncertain and subordinate to $\mathbf{q}$.*

Proof: Let $\lambda \in \mathcal{L}$ and let $k \in [1 \ldots K]$. We will show that

$$\lim_{n \to \infty} \lambda(\mathcal{G}_k^n) \quad = \quad q_k. \tag{A6}$$

There exists $\rho \in \mathcal{P}$ such that $\lambda \ll \rho$. Let $\phi := \dfrac{\mathrm{d}\lambda}{\mathrm{d}\rho}$ and $C := \sup_{s \in \mathcal{S}} \phi(s)$. Then $C < \infty$. Fix $\epsilon > 0$. Since $\mathcal{S}$ is complete and separable, it is Polish, so Lusin's Theorem yields a compact subset $\mathcal{K} \subseteq \mathcal{S}$ such that $\phi_{\upharpoonright \mathcal{K}}$ is uniformly continuous on $\mathcal{K}$ and

$$\rho(\mathcal{K}^{\complement}) \quad < \quad \frac{\epsilon}{8\,C}. \tag{A7}$$

(Aliprantis and Border, 2006, Theorem 12.8, p.438). It follows that

$$\lambda[\mathcal{K}^{\complement}] \quad = \quad \int_{\mathcal{K}^{\complement}} \phi \ \mathrm{d}\rho \underset{(*)}{\leqslant} C \cdot \rho[\mathcal{K}^{\complement}] \underset{(\dagger)}{\leqslant} C \cdot \frac{\epsilon}{8\,C} \quad = \quad \frac{\epsilon}{8}, \tag{A8}$$

where $(*)$ is because $0 \leqslant \phi(s) \leqslant C$ for all $s \in \mathcal{S}$, and $(\dagger)$ is by inequality (A7). Since $\{(\mathfrak{G}^n, \mathfrak{Y}^n, \epsilon_n)\}_{n=1}^{\infty}$ is a Poincaré sequence for $\mathcal{P}$, there is some $N_1 \in \mathbb{N}$ such that for all $n \geqslant N_1$ and all $\mathcal{Y} \in \mathfrak{Y}^n$,

$$\rho[\mathcal{G}_k^n \cap \mathcal{Y}] \quad = \quad q_k\, \rho[\mathcal{Y}]. \tag{A9}$$

Claim 1: *For all $n \geqslant N_1$,* $\displaystyle\sum_{\mathcal{Y} \in \mathfrak{Y}^n} \Big| \rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - q_k\, \rho[\mathcal{Y} \cap \mathcal{K}] \Big| \leqslant \frac{\epsilon}{4\,C}.$

Proof: Let $n \geqslant N$. For all $\mathcal{Y} \in \mathfrak{Y}^n$,

$$\begin{aligned}
&\Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - q_k\,\rho[\mathcal{Y} \cap \mathcal{K}]\Big| \\
&= \Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - \rho[\mathcal{G}_k^n \cap \mathcal{Y}] + \rho[\mathcal{G}_k^n \cap \mathcal{Y}] - q_k\,\rho[\mathcal{Y} \cap \mathcal{K}]\Big| \\
&\underset{(*)}{=} \Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - \rho[\mathcal{G}_k^n \cap \mathcal{Y}] + q_k\,\rho[\mathcal{Y}] - q_k\,\rho[\mathcal{Y} \cap \mathcal{K}]\Big| \\
&= \Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - \rho[\mathcal{G}_k^n \cap \mathcal{Y}] + q_k\,\Big(\rho[\mathcal{Y}] - \rho[\mathcal{Y} \cap \mathcal{K}]\Big)\Big| \\
&\leqslant \Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y}] - \rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}]\Big| + q_k\,\Big|\rho[\mathcal{Y}] - \rho[\mathcal{Y} \cap \mathcal{K}]\Big| \\
&= \rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}^\complement] + q_k\,\rho[\mathcal{Y} \cap \mathcal{K}^\complement]. \qquad \text{(A10)}
\end{aligned}$$

Here, $(*)$ is by equation (A9). Thus,

$$\begin{aligned}
\sum_{\mathcal{Y}\in\mathfrak{Y}^n}\Big|\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}] - q_k\,\rho[\mathcal{Y} \cap \mathcal{K}]\Big| &\underset{(\dagger)}{\leqslant} \sum_{\mathcal{Y}\in\mathfrak{Y}^n}\Big(\rho[\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}^\complement] + q_k\,\rho[\mathcal{Y} \cap \mathcal{K}^\complement]\Big) \\
&= \rho\left[\bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n}(\mathcal{G}_k^n \cap \mathcal{Y} \cap \mathcal{K}^\complement)\right] + q_k\,\rho\left[\bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n}(\mathcal{Y} \cap \mathcal{K}^\complement)\right] \\
&= \rho\left[\mathcal{G}_k^n \cap \mathcal{K}^\complement \cap \bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n}\mathcal{Y}\right] + q_k\,\rho\left[\mathcal{K}^\complement \cap \bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n}\mathcal{Y}\right] \\
&\underset{(*)}{=} \rho\left[\mathcal{G}_k^n \cap \mathcal{K}^\complement\right] + q_k\,\rho\left[\mathcal{K}^\complement\right] \underset{(\diamond)}{\leqslant} \frac{\epsilon}{8\,C} + \frac{\epsilon}{8\,C} = \frac{\epsilon}{4\,C},
\end{aligned}$$

as claimed. Here, $(\dagger)$ is by applying inequality (A10) to each $\mathcal{Y} \in \mathfrak{Y}^n$, $(*)$ is because $\bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n}\mathcal{Y} = \mathcal{S}$, and $(\diamond)$ is by inequality (A7). $\diamond$ `Claim 1`

Recall that $\phi_{\upharpoonright\mathcal{K}}$ is uniformly continuous on $\mathcal{K}$. Thus, there exists some $\delta > 0$ such that, for all $s_1, s_2 \in \mathcal{K}$, if $d(s_1, s_2) \leqslant \delta$, then $|\phi(s_1) - \phi(s_2)| < \frac{\epsilon}{4}$. Find $N_2 \in \mathbb{N}$ such that $\epsilon_n \leqslant \delta$ for all $n \geqslant N_2$. Thus, if $n \geqslant N_2$ and $\mathcal{Y} \in \mathfrak{Y}^n$, then $\text{diam}(\mathcal{Y}) \leqslant \epsilon_n \leqslant \delta$, so that for all $y_1, y_2 \in \mathcal{Y} \cap \mathcal{K}$ we have $|\phi(y_1) - \phi(y_2)| < \frac{\epsilon}{4}$. Thus, there is some $c_\mathcal{Y} \in \mathbb{R}_+$ such that $|\phi(y) - c_\mathcal{Y}| < \frac{\epsilon}{4}$ for all $y \in \mathcal{Y} \cap \mathcal{K}$. Thus, for all $n \geqslant N_2$,

$$\begin{aligned}
\Big|\lambda[\mathcal{Y} \cap \mathcal{K} \cap \mathcal{G}_k^n] - c_\mathcal{Y}\cdot\rho[\mathcal{Y} \cap \mathcal{K} \cap \mathcal{G}_k^n]\Big| &\underset{(*)}{=} \left|\int_{\mathcal{Y}\cap\mathcal{K}\cap\mathcal{G}_k^n}(\phi - c_\mathcal{Y})\ \mathrm{d}\rho\right| \\
\leqslant \int_{\mathcal{Y}\cap\mathcal{K}\cap\mathcal{G}_k^n}|\phi - c_\mathcal{Y}|\ \mathrm{d}\rho \leqslant \int_{\mathcal{Y}\cap\mathcal{K}\cap\mathcal{G}_k^n}\frac{\epsilon}{4}\,\mathrm{d}\rho &= \frac{\epsilon}{4}\cdot\rho[\mathcal{Y} \cap \mathcal{K} \cap \mathcal{G}_k^n], \qquad \text{(A11)}
\end{aligned}$$

where $(*)$ is because $\phi = \dfrac{\mathrm{d}\lambda}{\mathrm{d}\rho}$. By a very similar argument,

$$\Big|\lambda[\mathcal{Y} \cap \mathcal{K}] - c_\mathcal{Y}\rho[\mathcal{Y} \cap \mathcal{K}]\Big| \leqslant \frac{\epsilon}{4}\cdot\rho[\mathcal{Y} \cap \mathcal{K}], \qquad \text{for all } n \geqslant N_2. \qquad \text{(A12)}$$

Now, for any $n \in \mathbb{N}$,

$$\begin{aligned}
\lambda[\mathcal{G}_k^n \cap \mathcal{K}] - q_k\,\lambda[\mathcal{K}] \;&\underset{(*)}{=}\; \sum_{\mathcal{Y}\in\mathfrak{Y}} \lambda[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} \lambda[\mathcal{K}\cap\mathcal{Y}] \\
&= \sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - \sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] + \sum_{\mathcal{Y}\in\mathfrak{Y}} \lambda[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] \\
&\qquad - q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} \lambda[\mathcal{K}\cap\mathcal{Y}] + q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}}\rho[\mathcal{K}\cap\mathcal{Y}] - q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}}\rho[\mathcal{K}\cap\mathcal{Y}] \\
&= \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big(c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - q_k\,c_{\mathcal{Y}}\,\rho[\mathcal{K}\cap\mathcal{Y}]\Big) + \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big(\lambda[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}]\Big) \\
&\qquad - q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big(\lambda[\mathcal{K}\cap\mathcal{Y}] - c_{\mathcal{Y}}\rho[\mathcal{K}\cap\mathcal{Y}]\Big), \qquad \text{(A13)}
\end{aligned}$$

where $(*)$ is because $\bigsqcup_{\mathcal{Y}\in\mathfrak{Y}^n} \mathcal{Y} = \mathcal{S}$. Now let $N_\epsilon := \max\{N_1, N_2\}$. Then for all $n \geqslant N_\epsilon$,

$$\begin{aligned}
&\Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K}] - q_k\,\lambda[\mathcal{K}]\Big| \\
&\underset{(\diamond)}{\leqslant} \left|\sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}} \Big(\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - q_k\,\rho[\mathcal{K}\cap\mathcal{Y}]\Big)\right| + \left|\sum_{\mathcal{Y}\in\mathfrak{Y}} \Big(\lambda[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}]\Big)\right| \\
&\qquad + q_k \left|\sum_{\mathcal{Y}\in\mathfrak{Y}} \Big(\lambda[\mathcal{K}\cap\mathcal{Y}] - c_{\mathcal{Y}}\rho[\mathcal{K}\cap\mathcal{Y}]\Big)\right| \\
&\leqslant \sum_{\mathcal{Y}\in\mathfrak{Y}} c_{\mathcal{Y}} \Big|\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - q_k\,\rho[\mathcal{K}\cap\mathcal{Y}]\Big| + \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - c_{\mathcal{Y}}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}]\Big| \\
&\qquad + q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big|\lambda[\mathcal{K}\cap\mathcal{Y}] - c_{\mathcal{Y}}\rho[\mathcal{K}\cap\mathcal{Y}]\Big| \\
&\underset{(*)}{\leqslant} C \sum_{\mathcal{Y}\in\mathfrak{Y}} \Big|\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] - q_k\,\rho[\mathcal{K}\cap\mathcal{Y}]\Big| + \sum_{\mathcal{Y}\in\mathfrak{Y}} \frac{\epsilon}{4}\,\rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] + q_k \sum_{\mathcal{Y}\in\mathfrak{Y}} \frac{\epsilon}{4}\,\rho[\mathcal{K}\cap\mathcal{Y}] \\
&\underset{(\dagger)}{\leqslant} C\frac{\epsilon}{4C} + \frac{\epsilon}{4} \sum_{\mathcal{Y}\in\mathfrak{Y}} \rho[\mathcal{G}_k^n \cap \mathcal{K} \cap \mathcal{Y}] + \frac{\epsilon}{4} \sum_{\mathcal{Y}\in\mathfrak{Y}} \rho[\mathcal{K}\cap\mathcal{Y}] \\
&\leqslant \frac{\epsilon}{4} + \frac{\epsilon}{4}\,\rho[\mathcal{G}_k^n \cap \mathcal{K}] + \frac{\epsilon}{4}\,\rho[\mathcal{K}] \;\leqslant\; \frac{\epsilon}{4} + \frac{\epsilon}{4} + \frac{\epsilon}{4} \;=\; \frac{3\epsilon}{4}. \qquad \text{(A14)}
\end{aligned}$$

Here, $(\diamond)$ is by equation (A13), while $(*)$ is by inequalities (A11) and (A12). Finally, $(\dagger)$ is by Claim 1, and also uses the fact that $q_k \leqslant 1$. Thus, for all $n \geqslant N_\epsilon$, we have:

$$\begin{aligned}
\big|\lambda[\mathcal{G}_k^n] - q_k\big| &= \Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K}^{\complement}] + \lambda[\mathcal{G}_k^n \cap \mathcal{K}] - q_k\,\Big(\lambda[\mathcal{K}] + \lambda[\mathcal{K}^{\complement}]\Big)\Big| \\
&\leqslant \Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K}^{\complement}]\Big| + \Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K}] - q_k\,\lambda[\mathcal{K}]\Big| + \Big|\lambda[\mathcal{K}^{\complement}]\Big| \\
&\underset{(*)}{\leqslant} \frac{\epsilon}{8} + \Big|\lambda[\mathcal{G}_k^n \cap \mathcal{K}] - q_k\,\lambda[\mathcal{K}]\Big| + \frac{\epsilon}{8} \\
&\underset{(\dagger)}{\leqslant} \frac{\epsilon}{8} + \frac{3\epsilon}{4} + \frac{\epsilon}{8} \;=\; \epsilon.
\end{aligned}$$

where $(*)$ is by two applications of inequality (A8), while $(\dagger)$ is by inequality (A14).

We can construct such an N_ϵ for any $\epsilon > 0$. This proves the limit (A6). □

Lemma A.5 *Let $\mathcal{S}$ be any measurable space, and let $\mathcal{L}$ be a collection of probability measures on $\mathcal{S}$. Let $\mathcal{R}$ be the convex closure of $\mathcal{L}$ in the total variation norm. Let $\mathbf{q} \in \Delta^K$. If a partition sequence $(\mathfrak{G}^n)_{n=1}^\infty$ is $\mathcal{L}$-almost-objectively uncertain and subordinate to $\mathbf{q}$, then $(\mathfrak{G}^n)_{n=1}^\infty$ is $\mathcal{R}$-almost-objectively uncertain and subordinate to $\mathbf{q}$.*

Proof: Let $\mathcal{R}_0$ be the convex hull of $\mathcal{L}$. If $(\mathfrak{G}^n)_{n=1}^\infty$ is $\mathcal{L}$-almost-objectively uncertain and subordinate to $\mathbf{q}$, then it is easily shown that $(\mathfrak{G}^n)_{n=1}^\infty$ is also $\mathcal{R}_0$-almost-objectively uncertain subordinate to $\mathbf{q}$.

For all $n \in \mathbb{N}$, suppose $\mathfrak{G}^n = \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$. Let $\rho \in \mathcal{R}$. Then there is a sequence $\{\rho_m\}_{m=1}^\infty$ in $\mathcal{R}_0$ such that $\lim_{k\to\infty} \|\rho_m - \rho\|_{\text{vr}} = 0$. For all $k \in [1 \ldots K]$, we must show that the limit (5) holds for ρ.

Let $\epsilon > 0$. There exists $m \in \mathbb{N}$, with $\|\rho_m - \rho\|_{\text{vr}} < \frac{\epsilon}{2}$. This means that $|\rho_m(\mathcal{G}) - \rho(\mathcal{G})| < \epsilon/2$ for all measurable $\mathcal{G} \subseteq \mathcal{S}$. In particular,

$$|\rho(\mathcal{G}_k^n) - \rho_m(\mathcal{G}_k^n)| \quad < \quad \frac{\epsilon}{2}, \quad \text{for all } n \in \mathbb{N}, \text{ all } k \in [1 \ldots K]. \tag{A15}$$

The limit (5) holds for ρ_m, so there exists some $N_\epsilon \in \mathbb{N}$ such that

$$|\rho_m(\mathcal{G}_k^n) - q_k| \quad < \quad \frac{\epsilon}{2} \quad \text{for all } k \in [1 \ldots K] \text{ and all } n \geqslant N_\epsilon. \tag{A16}$$

Combining inequalities (A15) and (A16) yields $|\rho(\mathcal{G}_k^n) - q_k| < \epsilon$ for all $n \geqslant N_\epsilon$. We can obtain such an N_ϵ for any $\epsilon > 0$. Therefore, the limit (5) holds for ρ. □

Proof of Proposition 3(b) Suppose $\mathcal{S}$ is a standard Borel space. We can assume without loss of generality that there is a metric d making $(\mathcal{S}, d)$ a complete separable metric space, and the sigma algebra on $\mathcal{S}$ is the Borel sigma algebra. Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be separable and nonatomic; we must show that $\langle \mathcal{R} \rangle$ is consilient.

Let $\mathcal{N}$ be the closed subspace of $\mathcal{M}(\mathcal{S})$ spanned by $\mathcal{R}$. Then $\mathcal{N}$ is separable because $\mathcal{R}$ is separable. Thus, $\mathcal{N}$ it is spanned by a countable subset $\mathcal{H}$.[22] Since $\mathcal{R}$ (and hence $\mathcal{N}$) is nonatomic, all elements of $\mathcal{H}$ are nonatomic. Let $\mathcal{F}$ be the linear subspace of $\mathcal{M}(\mathcal{S})$ consisting of all finite linear combinations of elements from $\mathcal{H}$. Then $\mathcal{N}$ is the norm-closure of $\mathcal{F}$. Let $\mathcal{P} := \mathcal{F} \cap \Delta(\mathcal{S})$, and then let $\mathcal{L} := \langle \mathcal{P} \rangle$.

Claim 1: *$\langle \mathcal{R} \rangle$ is contained in the norm-closure of $\mathcal{L}$.*

[22] i.e., $\mathcal{N}$ is the norm-closure of the vector space of all finite linear combinations of elements of $\mathcal{H}$.

Proof: Let $\mu \in \langle\mathcal{R}\rangle$. Find $\rho \in \mathcal{R}$ such that $\mu \ll \rho$ and $\phi := \frac{\mathrm{d}\mu}{\mathrm{d}\rho}$ is bounded. Since $\mathcal{R} \subset \mathcal{N}$, and $\mathcal{N}$ is the norm-closure of $\mathcal{F}$, there exists a sequence $(\nu_n)_{n=1}^{\infty}$ in $\mathcal{F}$ converging to ρ in norm. For all $n \in \mathbb{N}$, let $\widetilde{\lambda}_n \in \mathcal{M}(\mathcal{S})$ be the measure such that $\widetilde{\lambda}_n \ll \nu_n$ and $\frac{\mathrm{d}\widetilde{\lambda}_n}{\mathrm{d}\nu_n} = \phi$. Next, let $\lambda_n := \widetilde{\lambda}_n/\ell_n$, where $\ell_n := \widetilde{\lambda}_n(\mathcal{S})$. Then $\lambda_n \in \mathcal{L}$. (*Proof:* By construction, λ_n is a probability measure, and $\lambda_n \ll \nu_n$. Let $\pi_n := \nu_n/\nu_n(\mathcal{S})$; then $\pi_n \in \mathcal{P}$, $\lambda_n \ll \pi_n$, and $\frac{\mathrm{d}\lambda_n}{\mathrm{d}\pi_n}$ is a multiple of ϕ, hence bounded.) To prove the claim, it suffices to show that the sequence $\{\lambda_n\}_{n=1}^{\infty}$ converges to μ in norm. For any $n \in \mathbb{N}$,

$$\|\mu - \lambda_n\|_{\mathrm{vr}} \quad \leqslant \quad \left\|\mu - \widetilde{\lambda}_n\right\|_{\mathrm{vr}} + \left\|\widetilde{\lambda}_n - \lambda_n\right\|_{\mathrm{vr}}. \tag{A17}$$

Now, for any measurable $\mathcal{U} \subseteq \mathcal{S}$,

$$\begin{aligned} \left|\mu(\mathcal{U}) - \widetilde{\lambda}_n(\mathcal{U})\right| \quad &\underset{(*)}{=} \quad \left|\int_{\mathcal{U}} \phi \ \mathrm{d}\rho - \int_{\mathcal{U}} \phi \ \mathrm{d}\nu_n\right| \quad = \quad \left|\int_{\mathcal{U}} \phi \ \mathrm{d}(\rho - \nu_n)\right| \\ &\leqslant \quad \|\phi\|_{\infty} \cdot |\rho(\mathcal{U}) - \nu_n(\mathcal{U})|\,, \end{aligned}$$

where $(*)$ is because $\frac{\mathrm{d}\mu}{\mathrm{d}\rho} = \phi = \frac{\mathrm{d}\widetilde{\lambda}_n}{\mathrm{d}\nu_n}$. Combining this inequality with defining formula (4), we deduce that $\left\|\mu - \widetilde{\lambda}_n\right\|_{\mathrm{vr}} \leqslant \|\phi\|_{\infty} \cdot \|\rho - \nu_n\|_{\mathrm{vr}} \xrightarrow[n\to\infty]{(\dagger)} 0$, where $(\dagger)$ is because ν_n converges to ρ in norm by hypothesis. Thus,

$$\lim_{n\to\infty} \left\|\mu - \widetilde{\lambda}_n\right\|_{\mathrm{vr}} \quad = \quad 0. \tag{A18}$$

Meanwhile,

$$\begin{aligned} \left\|\widetilde{\lambda}_n - \lambda_n\right\|_{\mathrm{vr}} \quad &= \quad \|\ell_n\lambda_n - \lambda_n\|_{\mathrm{vr}} \quad = \quad |1 - \ell_n| \cdot \|\lambda_n\|_{\mathrm{vr}} \quad = \quad |1 - \ell_n| \\ &= \quad \left|\mu(\mathcal{S}) - \widetilde{\lambda}_n(\mathcal{S})\right| \quad \underset{(*)}{=} \quad \left|\int_{\mathcal{S}} \phi \ \mathrm{d}\rho - \int_{\mathcal{S}} \phi \ \mathrm{d}\nu_n\right| \\ &= \quad \left|\int_{\mathcal{S}} \phi \ \mathrm{d}(\rho - \nu_n)\right| \quad \leqslant \quad \|\phi\|_{\infty} \cdot \|\rho - \nu_n\|_{\mathrm{vr}} \quad \xrightarrow[n\to\infty]{(\dagger)} \quad 0, \end{aligned}$$

where again, $(*)$ is because $\frac{\mathrm{d}\mu}{\mathrm{d}\rho} = \phi = \frac{\mathrm{d}\widetilde{\lambda}_n}{\mathrm{d}\nu_n}$ and $(\dagger)$ is because ν_n converges to ρ in norm. Thus,

$$\lim_{n\to\infty} \left\|\widetilde{\lambda}_n - \lambda_n\right\|_{\mathrm{vr}} \quad = \quad 0. \tag{A19}$$

Equations (A17), (A18) and (A19) yield $\lim_{n\to\infty} \|\mu - \lambda_n\|_{\mathrm{vr}} = 0$, as desired. ◇ `Claim 1`

Let $\mathbf{q} \in \Delta^K$. Since $\mathcal{S}$ is separable, Lemma A.3 says that $\mathcal{P}$ has a $\mathbf{q}$-Poincaré sequence $\{(\mathfrak{G}^n, \mathfrak{Y}^n, \epsilon_n)\}_{n=1}^{\infty}$. Then Lemma A.4 says that $(\mathfrak{G}^n)_{n=1}^{\infty}$ is $\mathcal{L}$-almost-objectively uncertain, subordinate to $\mathbf{q}$. Then Lemma A.5 and Claim 1 says that $(\mathfrak{G}^n)_{n=1}^{\infty}$ is $\langle\mathcal{R}\rangle$-almost-objectively uncertain, subordinate to $\mathbf{q}$. □

Proof of Proposition 4. Let $K \in \mathbb{N}$ and let $\mathbf{q} \in \Delta^K$. By hypothesis, there is an $\mathcal{R}$-almost-objectively uncertain sequence of partitions $(\mathfrak{G}^n)_{n=1}^{\infty}$ of $\mathcal{S}$ that is subordinate to $\mathbf{q}$. For all $n \in \mathbb{N}$, suppose $\mathfrak{G}^n = \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$. For all $k \in [1 \ldots K]$, let $\widehat{\mathcal{G}}_1^n := \phi^{-1}(\mathcal{G}_1^n)$. Then $\widehat{\mathfrak{G}}^n := \{\widehat{\mathcal{G}}_1^n, \ldots, \widehat{\mathcal{G}}_K^n\}$ is a measurable partition of $\widehat{\mathcal{S}}$ (because ϕ is measurable). This yields a partition sequence $(\widehat{\mathfrak{G}}^n)_{n=1}^{\infty}$ of $\widehat{\mathcal{S}}$. We will show that it is $\widehat{\mathcal{R}}$-almost-objectively uncertain and subordinate to $\mathbf{q}$.

To see this, let $\widehat{\rho} \in \widehat{\mathcal{R}}$. Let $\rho := \phi_*(\widehat{\rho})$. Then $\rho \in \mathcal{R}$. For all $k \in [1 \ldots K]$, we have $\widehat{\rho}(\widehat{\mathcal{G}}_k^n) = \rho(\mathcal{G}_k^n)$ for all $n \in \mathbb{N}$, so $\lim_{n\to\infty} \widehat{\rho}(\widehat{\mathcal{G}}_k^n) = \lim_{n\to\infty} \rho(\mathcal{G}_k^n) = q_k$, as desired.

To prove the claim for *strong* consilience, it suffices to show that $\langle\widehat{\mathcal{R}}\rangle \subseteq (\phi_*)^{-1}(\langle\mathcal{R}\rangle)$. To see this, let $\widehat{\mu} \in \langle\widehat{\mathcal{R}}\rangle$. Then there exists $\widehat{\rho} \in \widehat{\mathcal{R}}$ such that $\widehat{\mu} \ll \widehat{\rho}$ and such that $\widehat{\psi} := \frac{\mathrm{d}\widehat{\mu}}{\mathrm{d}\widehat{\rho}}$ is bounded. Let $\mu := \phi_*(\widehat{\mu})$ and $\rho := \phi_*(\widehat{\rho})$. Then $\mu \ll \rho$ and $\rho \in \mathcal{R}$. Furthermore, if $\psi := \frac{\mathrm{d}\mu}{\mathrm{d}\rho}$, then $\psi \circ \phi = \widehat{\psi}$. Thus, ψ is also bounded. Thus, $\mu \in \langle\mathcal{R}\rangle$. Thus, $\widehat{\mu} \in (\phi_*)^{-1}(\langle\mathcal{R}\rangle)$. □

Proof of Proposition 5. If $(\mathcal{S}, \mu, \phi)$ is mixing, then it is ergodic, and hence μ is nonatomic. Let $\mathfrak{G} = (\mathcal{G}_1, \ldots, \mathcal{G}_K)$ be a measurable partition such that $\mu[\mathcal{G}_k] = q_k$ for all $k \in [1 \ldots K]$; this exists because μ is nonatomic. Now, for all $n \in \mathbb{N}$, let $\mathfrak{G}^n := (\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n)$, where $\mathcal{G}_k^n := \phi^{-n}(\mathcal{G}_k)$ for all $k \in [1 \ldots K]$. We shall show that the sequence $(\mathfrak{G}^n)_{n=1}^{\infty}$ is $\mathcal{R}$-almost-objectively uncertain and subordinate to $\mathbf{q}$.

Let $\rho \in \mathcal{R}$; then $\rho \ll \mu$. Let $\psi := \frac{\mathrm{d}\rho}{\mathrm{d}\mu}$, then $\psi \in \mathcal{L}^2(\mathcal{S}, \mu)$ by hypothesis. For any measurable $\mathcal{G} \subseteq \mathcal{S}$, let $\mathbf{1}_{\mathcal{G}}$ be its indicator function. Then $\mathbf{1}_{\mathcal{G}} \in \mathcal{L}^2(\mathcal{S}, \mu)$, and

$$\lim_{n\to\infty} \int_{\mathcal{G}} \psi \circ \phi^n \, \mathrm{d}\mu \;=\; \lim_{n\to\infty} \langle \mathbf{1}_{\mathcal{G}},\ \psi \circ \phi^n \rangle \;\underset{(*)}{=}\; \int_{\mathcal{S}} \mathbf{1}_{\mathcal{G}} \, \mathrm{d}\mu \cdot \int_{\mathcal{S}} \psi \, \mathrm{d}\mu \;=\; \mu[\mathcal{G}] \cdot \rho[\mathcal{S}] \;=\; \mu[\mathcal{G}], \tag{A20}$$

where $(*)$ is a standard property of mixing MPDS (Walters 1982, Theorem 1.23(iii.2) on p.45 of §1.7; Fremlin 2006b, Proposition 372Q(iv), p.195). By applying change of variables, (A20) becomes

$$\lim_{n\to\infty} \int_{\phi^{-n}(\mathcal{G})} \psi \, \mathrm{d}\mu \;=\; \mu[\mathcal{G}]. \tag{A21}$$

In particular, we can apply (A21) to all $\mathcal{G}_k$ for all $k \in [1 \ldots K]$ to conclude that

$$\lim_{n\to\infty} \rho[\mathcal{G}_k^n] \;=\; \lim_{n\to\infty} \int_{\mathcal{G}_k^n} \psi \, \mathrm{d}\mu \;=\; \lim_{n\to\infty} \int_{\phi^{-n}(\mathcal{G}_k)} \psi \, \mathrm{d}\mu \;\underset{(*)}{=}\; \mu[\mathcal{G}_k] \;=\; q_k,$$

as desired. Here $(*)$ is by (A21).

This proves that $\mathcal{R}$ is consilient. It is *strongly* consilient because $\langle\mathcal{R}\rangle = \mathcal{R}$. To see this, suppose $\nu \in \langle\mathcal{R}\rangle$. Then $\nu \ll \rho$ for some $\rho \in \mathcal{R}$, and $\phi := \frac{\mathrm{d}\nu}{\mathrm{d}\rho}$ is bounded. By the definition of $\mathcal{R}$, $\rho \ll \mu$ and $\psi := \frac{\mathrm{d}\rho}{\mathrm{d}\mu} \in \mathcal{L}^2(\mathcal{S}, \mu)$. Thus, $\nu \ll \mu$ and $\frac{\mathrm{d}\nu}{\mathrm{d}\mu} = \phi \cdot \psi$ is also in $\mathcal{L}^2(\mathcal{S}, \mu)$ (because $\|\phi \cdot \psi\|_2 \leqslant \|\phi\|_\infty \cdot \|\psi\|_2$). Thus $\nu \in \mathcal{R}$. □

B Proofs from Section 4

The proof of Theorem 1 uses Proposition 6, so we will prove that first. The proof of Propositions 6, in turn, uses the following result, which can be seen as the special case of Proposition 6 for SEU representations.

Lemma B.1 *Let $\mathcal{R}$, $\mathbf{q} \in \Delta^K$, $\mathbf{x} \in \mathcal{X}^K$, and $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ be as in Proposition 6. For any $\rho \in \mathcal{R}$, and any measurable $u : \mathcal{X} \longrightarrow \mathbb{R}$,*

$$\lim_{n\to\infty} \int_{\mathcal{S}} u \circ \alpha^n \, \mathrm{d}\rho \quad = \quad \sum_{k=1}^K q_k \, u(x_k).$$

Proof: By hypothesis, there is an $\mathcal{R}$-almost-objectively uncertain partition sequence $\boldsymbol{\mathcal{G}} = (\mathfrak{G}^n)_{n=1}^\infty$ subordinate to the probability vector $\mathbf{q}$, and for all $n \in \mathbb{N}$, the act α^n is $\mathfrak{G}^n$-measurable. Suppose $\mathbf{q} = (q_1, \ldots, q_K) \in \Delta^K$. For all $n \in \mathbb{N}$, write $\mathfrak{G}^n := \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$, such that the limit equations (5) hold. By hypothesis, there is a K-tuple $\mathbf{x} \in \mathcal{X}^K$ such that for all $n \in \mathbb{N}$, all $k \in [1 \ldots K]$, and all $s \in \mathcal{G}_k^n$, we have $\alpha^n(s) = x_k$. Thus, for any $\rho \in \mathcal{R}$,

$$\int_{\mathcal{S}} u \circ \alpha^n \, \mathrm{d}\rho \quad = \quad \sum_{k=1}^K u(x_k) \, \rho(\mathcal{G}_k^n).$$

$$\text{Thus,} \quad \lim_{n\to\infty} \int_{\mathcal{S}} u \circ \alpha^n \, \mathrm{d}\rho \quad = \quad \lim_{n\to\infty} \sum_{k=1}^K u(x_k) \, \rho(\mathcal{G}_k^n) \quad = \quad \sum_{k=1}^K u(x_k) \, \lim_{n\to\infty} \rho(\mathcal{G}_k^n) \quad \underset{(*)}{=} \quad \sum_{k=1}^K u(x_k) \, q_k,$$

where $(*)$ is by the limit equations (5). □

Proof of Proposition 6. Recall the notation of equation (3). We will first show that the limit equation (9) holds for $\underline{V}$ and $\overline{V}$, and then show that it holds for V itself.

Claim 1: $\displaystyle \lim_{n\to\infty} \underline{V}(\alpha^n) = \sum_{k=1}^K q_k \, u(x_k).$

Proof: Let $B := \|u\|_\infty$. Then $B < \infty$, and the sequence $\{\underline{V}(\alpha^n)\}_{n=1}^\infty$ is bounded in the interval $[-B, B]$, so it has convergent subsequences. To prove the claim, it suffices to show that *every* convergent subsequence of $\{\underline{V}(\alpha^n)\}_{n=1}^\infty$ converges to $\sum_{k=1}^K q_k\, u(x_k)$.

So, let $\{n(\ell)\}_{\ell=1}^\infty$ be an increasing sequence in $\mathbb{N}$ such that the subsequence $\{\underline{V}(\alpha^{n(\ell)})\}_{\ell=1}^\infty$ converges to some limit V^*. We must show that $V^* = \sum_{k=1}^K q_k\, u(x_k)$. For all $\ell \in \mathbb{N}$, define the linear function $v_\ell : \Delta(\mathcal{S}) \longrightarrow \mathbb{R}$ by

$$v_\ell(\rho) \quad := \quad \int_\mathcal{S} u \circ \alpha^{n(\ell)} \, \mathrm{d}\rho, \qquad \text{for all } \rho \in \Delta(\mathcal{S}). \tag{B1}$$

This function is continuous in the norm topology, while $\mathcal{P}$ is closed in this topology. Thus,

$$\underline{V}(\alpha^{n(\ell)}) \quad = \quad \min_{\rho \in \mathcal{P}} v_\ell(\rho) \quad = \quad v_\ell(\rho_\ell), \tag{B2}$$

for some $\rho_\ell \in \mathcal{P}$. Furthermore, $\mathcal{P}$ is norm-compact. Thus, the sequence $\{\rho_\ell\}_{\ell=1}^\infty$ has a subsequence $\{\rho_{\ell_m}\}_{m=1}^\infty$ that converges to some limit point $\rho_* \in \mathcal{P}$ in the norm topology.

Let $\epsilon > 0$. There exists $M_1 \in \mathbb{N}$ such that, for all $m \geqslant M_1$, $\|\rho_{\ell_m} - \rho_*\|_{\mathrm{vr}} < \frac{\epsilon}{3\,B}$. Thus, for all $n \in \mathbb{N}$ and all $m \geqslant M_1$,

$$\left| \int_\mathcal{S} u \circ \alpha^n \, \mathrm{d}\rho_{\ell_m} - \int_\mathcal{S} u \circ \alpha^n \, \mathrm{d}\rho_* \right| \quad = \quad \left| \int_\mathcal{S} u \circ \alpha^n \, \mathrm{d}(\rho_{\ell_m} - \rho_*) \right| \\ \leqslant \quad \|u \circ \alpha^n\|_\infty \cdot \|\rho_{\ell_m} - \rho_*\|_{\mathrm{vr}} \quad < \quad B \cdot \frac{\epsilon}{3\,B} \quad = \quad \frac{\epsilon}{3}. \tag{B3}$$

In particular, setting $n := n(\ell_m)$ in (B3) and invoking equation (B1) yields

$$\left| v_{\ell_m}(\rho_{\ell_m}) - v_{\ell_m}(\rho_*) \right| \quad < \quad \frac{\epsilon}{3}. \tag{B4}$$

Next, substituting equation (B2) into inequality (B4) yields

$$\left| \underline{V}(\alpha^{n(\ell_m)}) - v_{\ell_m}(\rho_*) \right| \quad < \quad \frac{\epsilon}{3}. \tag{B5}$$

Meanwhile, $\rho_* \in \mathcal{R}$, so Lemma B.1 yields some $N \in \mathbb{N}$ such that,

$$\left| \int_\mathcal{S} u \circ \alpha^n \, \mathrm{d}\rho_* - \sum_{k=1}^K q_k\, u(x_k) \right| \quad < \quad \frac{\epsilon}{3} \qquad \text{for all } n \geqslant N. \tag{B6}$$

Since the sequence $\{n(\ell_m)\}_{m=1}^\infty$ is strictly increasing, there is some $M_2 \in \mathbb{N}$ such that $n(\ell_m) > N$ for all $m \geqslant M_2$. From this and inequality (B6), it follows that

$$\left| \int_\mathcal{S} u \circ \alpha^{n(\ell_m)} \, \mathrm{d}\rho_* - \sum_{k=1}^K q_k\, u(x_k) \right| \quad < \quad \frac{\epsilon}{3}, \qquad \text{for all } m \geqslant M_2. \tag{B7}$$

Using the defining equation (B1), we can rewrite inequality (B7) as follows:

$$\left| v_{\ell_m}(\rho_*) - \sum_{k=1}^{K} q_k \, u(x_k) \right| \quad < \quad \frac{\epsilon}{3}, \qquad \text{for all } m \geqslant M_2. \tag{B8}$$

Finally, by hypothesis, $\lim_{\ell \to \infty} \underline{V}(\alpha^{n(\ell)}) = V^*$. So there is some $L \in \mathbb{N}$ such that

$$\left| V^* - \underline{V}(\alpha^{n(\ell)}) \right| \quad < \quad \frac{\epsilon}{3}, \qquad \text{for all } \ell \geqslant L. \tag{B9}$$

Since the sequence $\{\ell_m\}_{m=1}^{\infty}$ is strictly increasing, there is some $M_3 \in \mathbb{N}$ such that $\ell_m > L$ for all $m \geqslant M_3$. From this and inequality (B9), it follows that

$$\left| V^* - \underline{V}(\alpha^{n(\ell_m)}) \right| \quad < \quad \frac{\epsilon}{3}, \qquad \text{for all } m \geqslant M_3. \tag{B10}$$

Now let $M_\epsilon := \max\{M_1, M_2, M_3\}$. Then for all $m \geqslant M_\epsilon$, we have

$$\begin{aligned} &\left| V^* - \sum_{k=1}^{K} q_k \, u(x_k) \right| \\ \leqslant \quad & \left| V^* - \underline{V}(\alpha^{n(\ell_m)}) \right| + \left| \underline{V}(\alpha^{n(\ell_m)}) - v_{\ell_m}(\rho_*) \right| + \left| v_{\ell_m}(\rho_*) - \sum_{k=1}^{K} q_k \, u(x_k) \right| \\ \underset{(*)}{<} \quad & \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \quad = \quad \epsilon, \end{aligned}$$

where $(*)$ is by inequalities (B5), (B8), and (B10).

This argument works for any $\epsilon > 0$. Thus, $V^* = \sum_{k=1}^{K} q_k \, u(x_k)$. ◇ Claim 1

By an argument similar to Claim 1 (replacing min with max), we can show that

$$\lim_{n \to \infty} \overline{V}(\alpha^n) \quad = \quad \sum_{k=1}^{K} q_k \, u(x_k). \tag{B11}$$

Combining inequality (3) with Claim 1 and equation (B11) yields equation (9), proving the theorem. □

Proposition 6 yields a convenient condition for asymptotic preferences.

Lemma B.2 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be consilient. Suppose $\succcurlyeq$ has a compact, contiguous GH representation (3) V with $\mathcal{P} \subseteq \mathcal{R}$. Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be almost-objective acts. Then $\boldsymbol{\alpha} \succ^{\infty} \boldsymbol{\beta}$ if and only if there exist $N \in \mathbb{N}$ and $\epsilon > 0$ such that $V(\alpha^n) > V(\beta^n) + \epsilon$ for all $n \geqslant N$.*

Proof: "$\Longrightarrow$" If $\boldsymbol{\alpha} \succ^\infty \boldsymbol{\beta}$, then there exist $\alpha', \beta' \in \mathcal{A}$ and $N \in \mathbb{N}$ such that for all $n \geqslant N$, we have $V(\alpha^n) > V(\alpha') > V(\beta') > V(\beta^n)$, and thus $V(\alpha^n) - V(\beta^n) > V(\alpha') - V(\beta') > 0$. So, let $\epsilon := V(\alpha') - V(\beta')$. Then $\epsilon > 0$, and $V(\alpha^n) > V(\beta^n) + \epsilon$ for all $n \geqslant N$.

"$\Longleftarrow$" Let $\mathbf{q} \in \Delta^K$ and $\mathbf{x} \in \mathcal{X}^K$ (for some $K \in \mathbb{N}$) and suppose that $\boldsymbol{\alpha}$ is subordinate to the lottery $(\mathbf{q}, \mathbf{x})$. Let $\mathbf{p} \in \Delta^L$ and $\mathbf{y} \in \mathcal{X}^L$ (for some $L \in \mathbb{N}$) and suppose that $\boldsymbol{\beta}$ is subordinate to the lottery $(\mathbf{p}, \mathbf{y})$. Let $A := \sum_{k=1}^{K} q_k\, u(x_k)$ and $B := \sum_{\ell=1}^{L} p_\ell\, u(y_\ell)$. Then Proposition 6 says that

$$\lim_{n\to\infty} V(\alpha^n) \quad = \quad A \qquad \text{and} \qquad \lim_{n\to\infty} V(\beta^n) \quad = \quad B. \tag{B12}$$

If $V(\alpha^n) > V(\beta^n) + \epsilon$ for all $n \geqslant N$, then the limits (B12) imply that $A \geqslant B + \epsilon$. Thus, $A - \frac{\epsilon}{3} > B + \frac{\epsilon}{3}$. The limits (B12) yield $M \in \mathbb{N}$ such that $V(\alpha^m) > A - \frac{\epsilon}{3}$ and $V(\beta^m) < B + \frac{\epsilon}{3}$ for all $m \geqslant M$. Since V is contiguous, its image $V(\mathcal{A})$ is a dense subset of an interval in $\mathbb{R}$. By prior observations, this interval must contain the subinterval $[B + \frac{\epsilon}{3}, A - \frac{\epsilon}{3}]$. So there exist $a, b \in V(\mathcal{A})$ such that $A - \frac{\epsilon}{3} > a > b > B + \frac{\epsilon}{3}$. Then for all $m \geqslant M$,

$$V(\alpha^m) \;>\; A - \frac{\epsilon}{3} \;>\; a \;>\; b \;>\; B + \frac{\epsilon}{3} \;>\; V(\beta^m). \tag{B13}$$

Let $\alpha', \beta' \in \mathcal{A}$ be such that $V(\alpha') = a$ and $V(\beta') = b$. Then for all $m \geqslant M$, the inequalities (B13) imply that $\alpha^m \succ \alpha' \succ \beta' \succ \beta^m$, as desired. □

Let $\mathcal{U}$ be the Banach space of bounded, measurable, real-valued functions on $\mathcal{X}$, endowed with the norm $\|\cdot\|_\infty$ defined by $\|u\|_\infty := \sup_{x\in\mathcal{X}} |u(x)|$ for all $u \in \mathcal{U}$. We shall use the following straightforward consequence of the Separating Hyperplane Theorem.

Lemma B.3 *Let $\{u_j\}_{j\in\mathcal{J}} \subset \mathcal{U}$, and suppose $\{u_i\}_{i\in\mathcal{I}}$ satisfy* Minimal Agreement. *Suppose there exists $z \in \mathcal{X}$ such that $u_j(z) = 0$ for all $j \in \mathcal{J}$. Let $\mathcal{C}$ be the convex cone in $\mathcal{U}$ spanned by $\{u_i\}_{i\in\mathcal{I}}$ and 0. If $u_o \notin \mathcal{C}$, then there exist finitely additive probability measures ν_1 and ν_2 on $\mathcal{X}$ such that*

$$\int_{\mathcal{X}} u_o \,\mathrm{d}\nu_1 \;<\; \int_{\mathcal{X}} u_o \,\mathrm{d}\nu_2, \quad \textit{while} \quad \int_{\mathcal{X}} u_i \,\mathrm{d}\nu_1 \;>\; \int_{\mathcal{X}} u_i \,\mathrm{d}\nu_2 \quad \textit{for all } i \in \mathcal{I}. \tag{B14}$$

Proof: (Pivato, 2022, Lemma A.2). □

Proof of Theorem 1. "$\Longrightarrow$" (by contradiction) Suppose $\succcurlyeq_o$ satisfies Almost-objective Pareto, but u_o is *not* weakly utilitarian. Let $z \in \mathcal{X}$. We can assume without loss of generality that $u_j(z) = 0$ for all $j \in \mathcal{J}$. To see this, let $c_j := u_j(z)$, and then define $\widetilde{u}_j(x) := u_j(x) - c_j$ for all $x \in \mathcal{X}$. If $\succcurlyeq_j$ has a GH representation (3), then $\succcurlyeq_j$ also admits a GH representation where u_j is replaced by $\widetilde{u}_j$.

Now let $\mathcal{C}$ be the closed, convex cone in $\mathcal{U}$ spanned by $\{u_i\}_{i\in\mathcal{I}}$ and 0. Then u_o is weakly utilitarian if and only if $u_o \in \mathcal{C}$. Thus, if u_o is *not* weakly utilitarian, then $u_o \notin \mathcal{C}$, in which case Lemma B.3 yields finitely additive probability measures ν_1 and ν_2 on $\mathcal{X}$ satisfying the inequalities (B14). For all $j \in \mathcal{J}$, let $\epsilon_j := \left|\int_{\mathcal{X}} u_j \, \mathrm{d}\nu_1 - \int_{\mathcal{X}} u_j \, \mathrm{d}\nu_2\right|$. Let

$$\epsilon \quad := \quad \frac{1}{5} \min_{j \in \mathcal{J}} \epsilon_j. \tag{B15}$$

Then $\epsilon > 0$. Inequalities (B14) and definition (B15) yield

$$\int_{\mathcal{X}} u_o \, \mathrm{d}\nu_2 - \int_{\mathcal{X}} u_o \, \mathrm{d}\nu_1 \quad > \quad 5\,\epsilon, \tag{B16}$$

$$\text{while} \quad \int_{\mathcal{X}} u_i \, \mathrm{d}\nu_1 - \int_{\mathcal{X}} u_i \, \mathrm{d}\nu_2 \quad > \quad 5\,\epsilon, \quad \text{for all } i \in \mathcal{I}. \tag{B17}$$

Let $R := \max\{\|u_j\|_\infty\}_{j \in \mathcal{J}}$; this value is finite because $\{u_j\}_{j \in \mathcal{J}}$ are bounded. Let $N := \lceil R/\epsilon \rceil + 1$; then $N\epsilon > R$, so the interval $[-N\,\epsilon,\ N\,\epsilon)$ contains the ranges of $\{u_j\}_{j \in \mathcal{J}}$. For all $j \in \mathcal{J}$ and all $n \in [-N \dots N]$, let $\mathcal{Y}_n^j := (u_j)^{-1}[n\,\epsilon, (n+1)\,\epsilon)$. Then $\mathfrak{Y}^j := \{\mathcal{Y}_n^j\}_{n=-N}^N$ is a measurable partition of $\mathcal{X}$. Let $\mathfrak{Y}$ be the common refining partition of $\{\mathfrak{Y}^j\}_{j \in \mathcal{J}}$. This is a measurable partition of $\mathcal{X}$. Suppose it has K cells, and write $\mathfrak{Y} = \{\mathcal{Y}_k\}_{k=1}^K$. For all $k \in [1..K]$, let $p_k^1 := \nu_1(\mathcal{Y}_k)$ and $p_k^2 := \nu_2(\mathcal{Y}_k)$. Then $\mathbf{p}^1 := (p_k^1)_{k=1}^K$ and $\mathbf{p}^2 := (p_k^2)_{k=1}^K$ are K-dimensional probability vectors. For all $k \in [1 \dots K]$, let $x_k \in \mathcal{Y}_k$.

Claim 1: *For all $j \in \mathcal{J}$,*

$$\left|\sum_{k=1}^K p_k^1\, u_j(x_k) - \int_{\mathcal{X}} u_j \, \mathrm{d}\nu_1\right| \ < \ \epsilon \quad \textit{and} \quad \left|\sum_{k=1}^K p_k^2\, u_j(x_k) - \int_{\mathcal{X}} u_j \, \mathrm{d}\nu_2\right| \ < \ \epsilon.$$

Proof: To prove the first inequality, note that

$$\left|\sum_{k=1}^K p_k^1\, u_j(x_k) - \int_{\mathcal{X}} u_j \, \mathrm{d}\nu_1\right| \quad = \quad \left|\sum_{k=1}^K \nu_1(\mathcal{Y}_k)\, u_j(x_k) - \sum_{k=1}^K \int_{\mathcal{Y}_k} u_j \, \mathrm{d}\nu_1\right|$$
$$= \quad \left|\sum_{k=1}^K \left(\int_{\mathcal{Y}_k} u_j(x_k) \, \mathrm{d}\nu_1 - \int_{\mathcal{Y}_k} u_j \, \mathrm{d}\nu_1\right)\right| \quad = \quad \left|\sum_{k=1}^K \left(\int_{\mathcal{Y}_k} u_j(x_k) - u_j(y) \, \mathrm{d}\nu_1[y]\right)\right|$$
$$\leqslant \quad \sum_{k=1}^K \int_{\mathcal{Y}_k} |u_j(x_k) - u_j(y)| \, \mathrm{d}\nu_1[y] \quad \underset{(*)}{<} \quad \sum_{k=1}^K \int_{\mathcal{Y}_k} \epsilon \, \mathrm{d}\nu_1 \quad = \quad \sum_{k=1}^K \epsilon\, \nu_1(\mathcal{Y}_k) \quad = \quad \epsilon,$$

as claimed. Here $(*)$ is because for all $k \in [1 \dots K]$, we have $x_k \in \mathcal{Y}_k$ while $n\,\epsilon \leqslant u_j(y) < (n+1)\,\epsilon)$ for all $y \in \mathcal{Y}_k$, so that $|u_j(x_k) - u_j(y)| < \epsilon$ for all $y \in \mathcal{Y}_k$. The proof of the second inequality is similar. ◇ Claim 1

Combining inequalities (B16) and (B17) with Claim 1 yields

$$\sum_{k=1}^K p_k^2\, u_o(x_k) - \sum_{k=1}^K p_k^1\, u_o(x_k) \quad > \quad 3\,\epsilon, \tag{B18}$$

$$\text{while} \quad \sum_{k=1}^K p_k^1\, u_i(x_k) - \sum_{k=1}^K p_k^2\, u_i(x_k) \quad > \quad 3\,\epsilon, \quad \text{for all } i \in \mathcal{I}. \tag{B19}$$

Let $\mathbf{q} \in \Delta^{K\times K}$ be the probability vector defined by $q_{k,\ell} := p^1_k\, p^2_\ell$ for all $k,\ell \in [1\ldots K]$. Since $\mathcal{R}$ is consilient, there is an $\mathcal{R}$-almost-objectively uncertain partition sequence $(\mathfrak{G}^n)_{n=1}^\infty$ subordinate to $\mathbf{q}$. For all $n \in \mathbb{N}$, write $\mathfrak{G}^n = \{\mathcal{G}^n_{k,\ell}\}_{k,\ell=1}^K$, with

$$\lim_{n\to\infty} \rho(\mathcal{G}^n_{k,\ell}) \;=\; q_{k,\ell}, \quad \text{for all } \rho \in \mathcal{R} \text{ and } k,\ell \in [1\ldots K]. \tag{B20}$$

For all $n \in \mathbb{N}$, and $\ell, k \in [1\ldots K]$, define $\mathcal{G}^n_{k,*} := \mathcal{G}^n_{k,1} \cup \mathcal{G}^n_{k,2} \cup \cdots \cup \mathcal{G}^n_{k,K}$ and $\mathcal{G}^n_{*,\ell} := \mathcal{G}^n_{1,\ell} \cup \mathcal{G}^n_{2,\ell} \cup \cdots \cup \mathcal{G}^n_{K,\ell}$. Then the equation (B20) yields

$$\lim_{n\to\infty} \rho(\mathcal{G}^n_{k,*}) \;=\; p^1_k \quad \text{and} \quad \lim_{n\to\infty} \rho(\mathcal{G}^n_{*,\ell}) \;=\; p^2_\ell, \quad \text{for all } \rho \in \mathcal{R}. \tag{B21}$$

For all $n \in \mathbb{N}$, define acts $\alpha^n, \beta^n : \mathcal{S} \longrightarrow \mathcal{X}$ as follows.

- For all $k \in [1\ldots K]$, let $\alpha^n(s) := x_k$ for all $s \in \mathcal{G}^n_{k,*}$.
- For all $\ell \in [1\ldots K]$, let $\beta^n(s) := x_\ell$ for all $s \in \mathcal{G}^n_{*,\ell}$.

Thus, $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ and $\boldsymbol{\beta} = (\beta^n)_{n=1}^\infty$ are $\mathcal{R}$-almost-objectively uncertain acts. They are compatible because for all $n \in \mathbb{N}$, α^n and β^n are both $\mathfrak{G}^n$-measurable. By construction and equations (B21), $\boldsymbol{\alpha}$ is subordinate to $(\mathbf{p}^1, \mathbf{x})$, while $\boldsymbol{\beta}$ is subordinate to $(\mathbf{p}^2, \mathbf{x})$.

Claim 2: $\boldsymbol{\alpha} \succ^\infty_i \boldsymbol{\beta}$ *for all* $i \in \mathcal{I}$.

Proof: For all $i \in \mathcal{I}$, let $V_i : \mathcal{A} \longrightarrow \mathbb{R}$ be a GH representation for $\succcurlyeq_i$ in which $\mathcal{P}_i \subseteq \mathcal{R}$ is norm-compact. Proposition 6 says that

$$\lim_{n\to\infty} V_i(\alpha^n) \;=\; \sum_{k=1}^K p^1_k\, u_i(x_k) \quad \text{and} \quad \lim_{n\to\infty} V_i(\beta^n) \;=\; \sum_{k=1}^K p^2_k\, u_i(x_k).$$

Thus, there exists $N \in \mathbb{N}$ such that

$$\left| V_i(\alpha^n) - \sum_{k=1}^K p^1_k\, u_i(x_k) \right| \;<\; \epsilon \quad \text{and} \quad \left| V_i(\beta^n) - \sum_{k=1}^K p^2_k\, u_i(x_k) \right| \;<\; \epsilon, \quad \text{for all } n \geqslant N. \tag{B22}$$

Combining inequalities (B19) and (B22), we obtain $V_i(\alpha^n) - V_i(\beta^n) > \epsilon$, for all $n \geqslant N$. Thus, $\boldsymbol{\alpha} \succ^\infty_i \boldsymbol{\beta}$ by Lemma B.2. ◇ Claim 2

By an argument identical to Claim 2, but using inequality (B18) rather than (B19), it is easy to prove that $\boldsymbol{\alpha} \prec^\infty_o \boldsymbol{\beta}$. This, together with Claim 2, is a violation of Almost-objective Pareto. Contradiction. To avoid this contradiction, u_o must be weakly utilitarian.

"$\Longleftarrow$" (by contradiction) Suppose u_o is weakly utilitarian; thus, $u_o = \sum_{i\in\mathcal{I}} c_i\, u_i$ for some constants $c_i \geqslant 0$. Suppose Almost-objective Pareto is violated. Then there exist compatible almost-objective acts $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that $\boldsymbol{\alpha} \succ^\infty_i \boldsymbol{\beta}$ for all $i \in \mathcal{I}$, while $\boldsymbol{\alpha} \prec^\infty_o \boldsymbol{\beta}$. Thus, for all $i \in \mathcal{I}$, Lemma B.2 yields $\epsilon_i > 0$ and $N_i \in \mathbb{N}$ such that

$$V_i(\alpha^n) - V_i(\beta^n) \quad > \quad 2\,\epsilon_i, \quad \text{for all } n \geqslant N_i, \tag{B23}$$

whereas there is some $\epsilon_o > 0$ and some $N_o \in \mathbb{N}$ such that

$$V_o(\beta^n) - V_o(\alpha^n) \quad > \quad 2\,\epsilon_o, \qquad \text{for all } n \geqslant N_o. \tag{B24}$$

There exist $K \in \mathbb{N}$, $\mathbf{p} \in \Delta^K$, and $\mathbf{x} \in \mathcal{X}^K$ such that $\boldsymbol{\alpha}$ is subordinate to $(\mathbf{p}, \mathbf{x})$. Likewise, there exist $L \in \mathbb{N}$, $\mathbf{q} \in \Delta^L$, and $\mathbf{y} \in \mathcal{X}^L$ such that $\boldsymbol{\beta}$ is subordinate to $(\mathbf{q}, \mathbf{y})$.

Claim 3: *For all $i \in \mathcal{I}$,* $\displaystyle\sum_{k=1}^{K} p_k\, u_i(x_k) \;-\; \sum_{\ell=1}^{L} q_\ell\, u_i(y_\ell) \;>\; 0.$

Proof: For all $i \in \mathcal{I}$, let $V_i : \mathcal{A} \longrightarrow \mathbb{R}$ be a GH representation for $\succcurlyeq_i$ in which $\mathcal{P}_i \subseteq \mathcal{R}$ is norm-compact. Now follow the argument from the proof of Claim 2 to obtain $M_i \in \mathbb{N}$ such that

$$\left| V_i(\alpha^m) - \sum_{k=1}^{K} p_k\, u_i(x_k) \right| \;<\; \epsilon_i \;\text{ and }\; \left| V_i(\beta^m) - \sum_{\ell=1}^{L} q_\ell\, u_i(y_\ell) \right| \;<\; \epsilon_i, \quad \text{for all } m \geqslant M_i. \tag{B25}$$

Now let $n \geqslant \max\{N_i, M_i\}$, and combine (B23) and (B25) to get the claimed inequality. $\diamond$ `Claim 3`

By an argument similar to Claim 3, but using inequality (B24) rather than (B23), one can show that

$$\sum_{k=1}^{K} p_k\, u_o(x_k) \;-\; \sum_{\ell=1}^{L} q_\ell\, u_o(y_\ell) \quad < \quad 0. \tag{B26}$$

Now, $u_o = \sum_{i \in \mathcal{I}} c_i\, u_i$. Thus,

$$\begin{aligned} \sum_{k=1}^{K} p_k\, u_o(x_k) \;-\; \sum_{\ell=1}^{L} q_\ell\, u_o(y_\ell) \;&=\; \sum_{k=1}^{K} p_k \sum_{i \in \mathcal{I}} c_i\, u_i(x_k) \;-\; \sum_{\ell=1}^{L} q_\ell \sum_{i \in \mathcal{I}} c_i\, u_i(y_\ell) \\ &=\; \sum_{i \in \mathcal{I}} c_i \left(\sum_{k=1}^{K} p_k\, u_i(x_k) \;-\; \sum_{\ell=1}^{L} q_\ell\, u_i(y_\ell) \right). \end{aligned} \tag{B27}$$

But $c_i \geqslant 0$ for all $i \in \mathcal{I}$, so equation (B27), inequality (B26) and Claim 3 are logically inconsistent. To avoid this contradiction, Almost-objective Pareto must be satisfied. □

Proof of Corollary 2. For all $j \in \mathcal{J}$, let $W_j := \int_{\mathcal{S}} w_j \, \mathrm{d}\rho_j$, and then define a new probability measure $\widetilde{\rho}_j \in \mathcal{M}(\mathcal{S})$ such that $\frac{\mathrm{d}\widetilde{\rho}_j}{\mathrm{d}\rho_j} = \frac{w_j}{W_j}$. Observe that $\widetilde{\rho}_j \in \langle \mathcal{R} \rangle$ because $\rho_j \in \mathcal{R}$, $\widetilde{\rho}_j \ll \rho_j$, and $\frac{w_j}{W_j}$ is bounded. It is easily verified that the state-dependent SEU representation (8) of $\succcurlyeq_j$ in terms of u_j, w_j and ρ_j is equivalent to a state-*independent* SEU representation (2) in terms of u_j and $\widetilde{\rho}_j$. Now apply Corollary 1 to the SEU representations $\{(u_j, \widetilde{\rho}_j)\}_{j \in \mathcal{J}}$ to prove the result. □

Proof of Corollary 3. For all $j \in \mathcal{J}$, the preference $\succcurlyeq_j$ has a GH representation induced by a compact set $\mathcal{P}_j \subseteq \Delta(\mathcal{S})$ of nonatomic probability measures. Let $\mathcal{R} := \bigcup_{j \in \mathcal{J}} \mathcal{P}_j$. Then $\mathcal{R}$ is compact (because $\mathcal{J}$ is finite), hence a separable subset of $\Delta(\mathcal{S})$. Thus, Proposition 2 say that $\mathcal{R}$ is consilient. By definition, $\succcurlyeq_o$ satisfies Almost-objective Pareto* if and only if it satisfies $\mathcal{R}$-Almost-objective Pareto, which (by Theorem 1) is the case if and only if u_o is weakly utilitarian. □

C Proof of results from Section 5

Proof of Proposition 7. The following axiom about beliefs is due to Mongin (1995):

C1. For all $\mathcal{A}, \mathcal{B} \subseteq \mathcal{S}$, if $\rho_i(\mathcal{A}) \geqslant \rho_i(\mathcal{B})$ for all $i \in \mathcal{I}$, then $\rho_o(\mathcal{A}) \geqslant \rho_o(\mathcal{B})$.

Claim 1: $\succeq_o$ *satisfies* Dichotomous Pareto *if and only if* ρ_o *satisfies C1.*

Proof: Let α, β be compatible dichotomous acts, ranging over a dichotomy $\{x, y\}$. Then there exist measurable subsets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{S}$ such that for all $s \in \mathcal{S}$, we have

$$\alpha(s) = \begin{cases} x & \text{if } s \in \mathcal{A}; \\ y & \text{otherwise.} \end{cases} \qquad \beta(s) = \begin{cases} x & \text{if } s \in \mathcal{B}; \\ y & \text{otherwise.} \end{cases} \tag{C1}$$

Thus, for all $j \in \mathcal{J}$, we have $\alpha \succeq_j \beta$ if and only if $\rho_j(\mathcal{A}) \geqslant \rho_j(\mathcal{B})$. It follows that Dichotomous Pareto (for α vs. β) is equivalent to C1 (for $\mathcal{A}$ vs. $\mathcal{B}$). We can make this argument for any compatible pair of dichotomous acts α and β. Conversely, for any measurable $\mathcal{A}, \mathcal{B} \subseteq \mathcal{A}$, we can construct compatible dichotomous acts α and β satisfying statement (C1). ◇ Claim 1

"$\Longrightarrow$" If $\succeq_o$ satisfies Dichotomous Pareto, then Claim 1 says that ρ_o satisfies C1. Thus, ρ_o is a weighted average of $\{\rho_i\}_{i \in \mathcal{I}}$ by Proposition 2 (p.321) of Mongin (1995).

"$\Longleftarrow$" If ρ_o is a weighted average of $\{\rho_i\}_{i \in \mathcal{I}}$, then it clearly satisfies C1. Thus, $\succeq_o$ satisfies Dichotomous Pareto, by Claim 1. □

Theorem 2 is a consequence of a more general result. Let $\trianglerighteq$ be a preorder on $\mathcal{A}$ (e.g. a Bewley preference). We will write $\boldsymbol{\alpha} \trianglerighteq^\omega \boldsymbol{\beta}$ if there exists $N \in \mathbb{N}$ such that $\alpha_n \trianglerighteq \beta_n$ for all $n \geqslant N$.

Now let $\{\trianglerighteq_j\}_{j \in \mathcal{J}}$ be a family of Bewley preferences satisfying MAO, and consider a sequence of acts $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$. We shall say that $\boldsymbol{\alpha}$ is *dichotomous* if there is some dichotomy (x, y) such that α^n ranges over $\{x, y\}$ for all $n \in \mathbb{N}$. Suppose that $\mathcal{R} \subseteq \Delta(\mathcal{S})$ is consilient. We shall say that $\boldsymbol{\alpha}$ is $\mathcal{R}$-*piecewise almost-objective* if there is a measurable partition $\mathfrak{H} = \{\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_J\}$ of $\mathcal{S}$ and a family of $\mathcal{R}$-almost-objective acts $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \ldots, \boldsymbol{\alpha}_J$ such that for all $n \in \mathbb{N}$, and all $j \in [1 \ldots J]$, we have

$$\alpha^n(s) = \alpha_j^n(s) \quad \text{for all } s \in \mathcal{H}_j. \tag{C2}$$

In other words, $\boldsymbol{\alpha}$ is achieved by "patching together" $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_J$ according to the partition $\mathfrak{H}$. Any almost-objective act is piecewise almost-objective (via the trivial partition). Consider the following axiom:

$\mathcal{R}$-Dichotomous piecewise almost-objective Pareto. Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be two dichotomous $\mathcal{R}$-piecewise almost-objective acts. If $\boldsymbol{\alpha} \trianglerighteq_i^\omega \boldsymbol{\beta}$ for all $i \in \mathcal{I}$, then $\boldsymbol{\alpha} \trianglerighteq_o^\omega \boldsymbol{\beta}$.

Compared to $\mathcal{R}$-almost-objective Pareto, this new axiom is broader in one way (it applies to *piecewise* almost-objective acts), but narrower in another way (it applies only to *dichotomous* almost-objective acts). It also differs from $\mathcal{R}$-almost-objective Pareto in that it involves the (possibly incomplete) Bewley preferences $\{\trianglerighteq_j\}_{j\in\mathcal{J}}$ instead of the weak orders $\{\succcurlyeq_j\}_{j\in\mathcal{J}}$, and it requires the planner's asymptotic preferences to actually *agree* with those of the individuals, rather than simply not disagree. Theorem 2 is an immediate consequence of the following more general result.

Theorem C.1 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be strongly consilient. For all $j \in \mathcal{J}$, let $\trianglerighteq_j$ be a Bewley preferences induced by a compact subset $\mathcal{P}_j \subseteq \mathcal{R}$ and utility function $u_j : \mathcal{X} \longrightarrow \mathbb{R}$. Suppose $\{\trianglerighteq_j\}_{j\in\mathcal{J}}$ satisfy MAO. Let $\overline{\mathcal{P}}$ be the $\mathfrak{T}$-closed convex hull of $\bigcup_{i\in\mathcal{I}} \mathcal{P}_i$. The following are equivalent:*

- **(a)** $\trianglerighteq_o$ *satisfies* Dichotomous piecewise $\mathcal{R}$-almost-objective Pareto.
- **(b)** $\trianglerighteq_o$ *satisfies* Dichotomous Pareto.
- **(c)** $\mathcal{P}_o \subseteq \overline{\mathcal{P}}$.

The proof of Theorem C.1 requires some preliminaries. A measurable function $\phi : \mathcal{S} \longrightarrow \mathbb{R}$ is *simple* if it takes only a finite number of values. For any simple function ϕ, define $\phi^* : \mathcal{M}(\mathcal{S}) \longrightarrow \mathbb{R}$ by setting by $\phi^*(\mu) := \int_{\mathcal{S}} \phi \, \mathrm{d}\mu$ for all $\mu \in \mathcal{M}(\mathcal{S})$. Then ϕ^* is a linear functional and continuous in the norm topology.

Lemma C.2 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be strongly consilient. Let $\{\trianglerighteq_j\}_{j\in\mathcal{J}}$ be Bewley preferences on $\mathcal{A}$ that satisfy MAO, and let (x, y) be a dichotomy for $\{\trianglerighteq_j\}_{j\in\mathcal{J}}$. Suppose their Bewley representations (10) have belief sets contained in $\mathcal{R}$, and utility functions $\{u_j\}_{j\in\mathcal{J}}$ that are renormalized such that that $u_j(x) = 1$ and $u_j(y) = 0$ for all $j \in \mathcal{J}$. Let $\phi : \mathcal{S} \longrightarrow \mathbb{R}$ be a simple function, and consider the functional $\phi^* : \mathcal{M}(\mathcal{S}) \longrightarrow \mathbb{R}$. There exists a dichotomous piecewise $\mathcal{R}$-almost-objective act $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ such that*

$$\lim_{n\to\infty} \int_{\mathcal{S}} u_j \circ \alpha^n \, \mathrm{d}\rho_j \quad = \quad \phi^*(\rho_j), \quad \textit{for all } j \in \mathcal{J} \textit{ and all } \rho_j \in \mathcal{P}_j.$$

Proof: By hypothesis, there exists a measurable partition $\{\mathcal{G}_1, \ldots, \mathcal{G}_L\}$ of $\mathcal{S}$ and some $r_1, \ldots, r_L \in \mathbb{R}$ such that $\phi = \sum_{\ell=1}^{L} r_\ell \, \mathbf{1}_{\mathcal{G}_\ell}$, where $\mathbf{1}_{\mathcal{G}_\ell}$ is the indicator function of $\mathcal{G}_\ell$.

Claim 1: *For all $\ell \in [1 \ldots L]$, there exists a sequence $(\mathcal{F}_\ell^n)_{n=1}^\infty$ of subsets of $\mathcal{G}^\ell$ such that $\lim_{n\to\infty} \rho(\mathcal{F}_\ell^n) = r_\ell \cdot \rho(\mathcal{G}_\ell)$ for all $\rho \in \mathcal{R}$.*

Proof: For all $\rho \in \mathcal{R}$, let $\rho_{\mathcal{G}_\ell} \in \Delta(\mathcal{S})$ be the measure obtained by Bayes-updating ρ on $\mathcal{G}_\ell$. Then $\rho_{\mathcal{G}_\ell} \in \langle\mathcal{R}\rangle$, because $\rho_{\mathcal{G}_\ell} \ll \rho$ and $\frac{\mathrm{d}\rho_{\mathcal{G}_\ell}}{\mathrm{d}\rho} = \mathbf{1}_{\mathcal{G}}/\rho(\mathcal{G}_\ell)$ is bounded.

By strong consilience, there is a sequence of measurable subsets $(\mathcal{E}^n)_{n=1}^\infty$ in $\mathcal{S}$ such that $\lim_{n\to\infty} \mu(\mathcal{E}^n) = r_\ell$ for all $\mu \in \langle\mathcal{R}\rangle$. Thus, $\lim_{n\to\infty} \rho_{\mathcal{G}_\ell}(\mathcal{E}^n) = r_\ell$ for all $\rho \in \mathcal{R}$, by the previous paragraph. For all $n \in \mathbb{N}$, let $\mathcal{F}_\ell^n := \mathcal{E}^n \cap \mathcal{G}_\ell$. Then $\mathcal{F}_\ell^n \subseteq \mathcal{G}_\ell$. For all $\rho \in \mathcal{R}$, we have $\rho(\mathcal{F}_\ell^n) = \rho_{\mathcal{G}_\ell}(\mathcal{F}_\ell^n) \cdot \rho(\mathcal{G}_\ell)$ and $\rho_{\mathcal{G}_\ell}(\mathcal{F}_\ell^n) = \rho_{\mathcal{G}_\ell}(\mathcal{E}^n)$. Thus, $\lim_{n\to\infty} \rho_{\mathcal{G}_\ell}(\mathcal{F}_\ell^n) = r_\ell$, and hence $\lim_{n\to\infty} \rho(\mathcal{F}_\ell^n) = r_\ell \cdot \rho(\mathcal{G}_\ell)$. ◇ Claim 1

Now, for all $n \in \mathbb{N}$, let $\mathcal{F}_1^n, \mathcal{F}_2^n, \ldots, \mathcal{F}_L^n$ be as in Claim 1; these sets are disjoint because $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_L$ are disjoint. Let $\mathcal{F}^n := \bigsqcup_{\ell=1}^{L} \mathcal{F}_\ell^N$, and then define $\alpha^n \in \mathcal{A}$ by:

$$\text{for all } s \in \mathcal{S}, \qquad \alpha^n(s) \quad := \quad \begin{cases} x & \text{if } s \in \mathcal{F}^n; \\ y & \text{otherwise.} \end{cases}$$

The sequence $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ is clearly dichotomous, and is piecewise $\mathcal{R}$-almost objective (with respect to the original partition $\mathfrak{G}$). For all $j \in \mathcal{J}$, we have $u_j \circ \alpha^n = \mathbf{1}_{\mathcal{F}^n}$. Thus, for any $\rho \in \mathcal{R}$,

$$\int_{\mathcal{S}} u_j \circ \alpha^n \, \mathrm{d}\rho \;=\; \int_{\mathcal{S}} \mathbf{1}_{\mathcal{F}^n} \, \mathrm{d}\rho \;=\; \rho[\mathcal{F}^n] \;=\; \rho\left[\bigsqcup_{\ell=1}^{L} \mathcal{F}_\ell^N\right] \;=\; \sum_{\ell=1}^{L} \rho[\mathcal{F}_\ell^n]. \tag{C3}$$

Thus,

$$\lim_{n\to\infty} \int_{\mathcal{S}} u_j \circ \alpha^n \, \mathrm{d}\rho \;\underset{(*)}{=}\; \lim_{n\to\infty} \sum_{\ell=1}^{L} \rho[\mathcal{F}_\ell^n] \;=\; \sum_{\ell=1}^{L} \lim_{n\to\infty} \rho[\mathcal{F}_\ell^n]$$
$$\underset{(\dagger)}{=}\; \sum_{\ell=1}^{L} r_\ell \cdot \rho(\mathcal{G}_\ell) \;=\; \int_{\mathcal{S}} \sum_{\ell=1}^{L} r_\ell \, \mathbf{1}_{\mathcal{G}_\ell} \, \mathrm{d}\rho \;=\; \int_{\mathcal{S}} \phi \, \mathrm{d}\rho \;=\; \phi^*(\rho),$$

as desired. Here, $(*)$ is by equation (C3), and $(\dagger)$ is by Claim 1. □

Let $\mathfrak{T}$ be the weak topology on $\mathcal{M}(\mathcal{S})$ induced by the family $\{\phi^*;\ \phi : \mathcal{S}\longrightarrow\mathcal{R}$ a simple function$\}$. The total variation norm topology on $\mathcal{M}(\mathcal{S})$ is finer than $\mathfrak{T}$. Thus, if a subset $\mathcal{P} \subset \mathcal{M}(\mathcal{S})$ is compact in the total variation norm topology, then $\mathcal{P}$ is compact in $\mathfrak{T}$. For any measurable $\mathcal{B} \subseteq \mathcal{M}(\mathcal{S})$, define $\eta_{\mathcal{B}} : \mathcal{M}(\mathcal{S})\longrightarrow\mathbb{R}$ by setting $\eta_{\mathcal{B}}(\mu) := \mu[\mathcal{B}]$ for all $\mu \in \mathcal{M}(\mathcal{S})$. For any simple function $\phi : \mathcal{S}\longrightarrow\mathbb{R}$, with corresponding linear functional $\phi^* : \mathcal{M}(\mathcal{S})\longrightarrow\mathbb{R}$, if $\phi = \sum_{\ell=1}^{L} r_\ell \, \mathbf{1}_{\mathcal{G}_\ell}$, then $\phi^* := \sum_{\ell=1}^{L} r_\ell \, \eta_{\mathcal{G}_\ell}$.

Proof of Theorem C.1. "(b) $\Longrightarrow$ (a)" Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be dichotomous $\mathcal{R}$-piecewise almost-objective acts, and suppose that $\boldsymbol{\alpha} \trianglerighteq_i^\omega \boldsymbol{\beta}$ for all $i \in \mathcal{I}$. Thus, for all $i \in \mathcal{I}$ there is some $N_i \in \mathbb{N}$ such that $\alpha^n \trianglerighteq_i \beta^n$ for all $n \geqslant N_i$. Let $N := \max\{N_i\}_{i\in\mathcal{I}}$. Then $\alpha^n \trianglerighteq_o \beta^n$ for all $n \geqslant N$, by Dichotomous Pareto. Thus, $\boldsymbol{\alpha} \trianglerighteq_o^\omega \boldsymbol{\beta}$, as desired.

"(c) $\Longrightarrow$ (b)" Let (x, y) be a dichotomy for $\{\trianglerighteq_j\}_{j\in\mathcal{J}}$. Define $v : \{x, y\}\longrightarrow\mathbb{R}$ by $v(x) = 1$ and $v(y) = 0$. For all $j \in \mathcal{J}$, suppose $\trianglerighteq_j$ has a Bewley representation $(u_j, \mathcal{P}_j)$ for some $u_j : \mathcal{X}\longrightarrow\mathbb{R}$. By applying positive affine transformations to $\{u_j\}_{j\in\mathcal{J}}$ if necessary, we can assume without loss of generality that u_j agrees with v on $\{x, y\}$, for all $j \in \mathcal{J}$.

Let $\alpha, \beta \in \mathcal{A}$ be compatible dichotomous acts ranging over $\{x, y\}$. Then $u_j \circ \alpha = v \circ \alpha$ and $u_j \circ \beta = v \circ \beta$ for all $j \in \mathcal{J}$. Suppose $\alpha \trianglerighteq_i \beta$ for all $i \in \mathcal{I}$. Then for all $i \in \mathcal{I}$, we have $\int_{\mathcal{S}} u_i \circ \alpha \, d\rho \geqslant \int_{\mathcal{S}} u_i \circ \beta \, d\rho$ for all $\rho \in \mathcal{P}_i$. Using the above identities, we can rewrite this $\int_{\mathcal{S}} v \circ \alpha \, d\rho \geqslant \int_{\mathcal{S}} v \circ \beta \, d\rho$ for all $\rho \in \mathcal{P}_i$ and all $i \in \mathcal{I}$. Convex combinations of probability measures preserve weak inequalities of expected values, so this inequality also holds for all ρ in the convex hull of $\bigcup_{i\in\mathcal{I}} \mathcal{P}_i$. Furthermore, $v \circ \alpha$ and $v \circ \beta$ are simple functions, and $\mathfrak{T}$-limits preserve weak inequalities of expected values for simple functions (because $\mathfrak{T}$ is the weak topology generated by simple functions). Thus, we deduce that $\int_{\mathcal{S}} v \circ \alpha \, d\rho \geqslant \int_{\mathcal{S}} v \circ \beta \, d\rho$ for all $\rho \in \overline{\mathcal{P}}$. Since $\mathcal{P}_o \subseteq \overline{\mathcal{P}}$, this implies that $\int_{\mathcal{S}} v \circ \alpha \, d\rho \geqslant \int_{\mathcal{S}} v \circ \beta \, d\rho$ for all $\rho \in \mathcal{P}_o$. In other words, $\int_{\mathcal{S}} u_o \circ \alpha \, d\rho \geqslant \int_{\mathcal{S}} u_o \circ \beta \, d\rho$ for all $\rho \in \mathcal{P}_o$. Thus, $\alpha \trianglerighteq_o \beta$, as desired.[23]

"(a) $\Longrightarrow$ (c)" (by contrapositive) Suppose $\mathcal{P}_o \nsubseteq \overline{\mathcal{P}}$. Let $\mathcal{P}_*$ be a nonempty norm-compact, convex subset of $\mathcal{P}_o$ that is disjoint from $\overline{\mathcal{P}}$. (For example, let $\mathcal{P}_* := \{\rho_o\}$, for any $\rho_o \in \mathcal{P}_o \setminus \overline{\mathcal{P}}$.) Then $\mathcal{P}_*$ is also $\mathfrak{T}$-compact, as explained above. In the $\mathfrak{T}$ topology, $\mathcal{M}(\mathcal{S})$ is a locally convex topological vector space, and $\mathcal{P}_*$ and $\overline{\mathcal{P}}$ are disjoint, closed convex subsets, one of which is compact. So the Strong Separating Hyperplane Theorem (Aliprantis and Border, Thm. 5.79, p.207) yields a $\mathfrak{T}$-continuous linear functional $\varphi : \mathcal{M}(\mathcal{S})\longrightarrow\mathbb{R}$ and $r_1 < r_2 \in \mathbb{R}$ such that

$$\varphi(\mu) \;<\; r_1 \;<\; r_2 \;<\; \varphi(\rho), \qquad \text{for all } \mu \in \mathcal{P}_* \text{ and } \rho \in \overline{\mathcal{P}}. \tag{C4}$$

Let $r := (r_1 + r_2)/2$ and let $\epsilon := (r_1 - r_2)/6$; then $r_1 = r - 3\,\epsilon$ and $r_2 = r + 3\,\epsilon$. Consider the $\mathfrak{T}$-continuous linear functional $\eta_{\mathcal{S}} : \mathcal{M}(\mathcal{S})\longrightarrow\mathbb{R}$ defined by $\eta_{\mathcal{S}}(\mu) := \mu[\mathcal{S}]$ for all $\mu \in \mathcal{M}(\mathcal{S})$. Let $\varphi' := \varphi - r \cdot \eta_{\mathcal{S}}$. Then φ' is also a $\mathfrak{T}$-continuous linear functional, and inequality (C4) yields:

$$\varphi'(\mu) \;<\; -3\,\epsilon \;<\; 0 \;<\; 3\,\epsilon \;<\; \varphi'(\rho), \qquad \text{for all } \mu \in \mathcal{P}_* \text{ and } \rho \in \overline{\mathcal{P}}. \tag{C5}$$

Any $\mathfrak{T}$-linear functional on $\mathcal{M}(\mathcal{S})$ has the form ϕ^* for some simple function $\phi : \mathcal{S}\longrightarrow\mathbb{R}$, because $\mathfrak{T}$ is the weak topology on $\mathcal{M}(\mathcal{S})$ generated by the vector space of simple functions (Aliprantis and Border, 2006, Theorem 5.93, p. 212). Thus, $\varphi' = \sum_{\ell=1}^{L} r_\ell \, \eta_{\mathcal{G}_\ell}$ for some disjoint measurable subsets $\mathcal{G}_1, \ldots, \mathcal{G}_L \subseteq \mathcal{S}$ and some $r_1, \ldots, r_L \in \mathbb{R}$. By rearranging $\mathcal{G}_1, \ldots, \mathcal{G}_L$ if necessary, we can assume that $r_1, \ldots, r_J < 0$ and $r_{J+1}, \ldots, r_L > 0$ for

[23] This proof does not use consilience. So in fact it works for any $\mathcal{R} \subseteq \Delta(\mathcal{S})$.

some $J \in \mathbb{N}$. Let $\varphi_- := -\sum_{j=1}^{J} r_j\, \eta_{\mathcal{G}_j}$ and $\varphi_+ := \sum_{\ell=J+1}^{L} r_\ell\, \eta_{\mathcal{G}_\ell}$. Then $\varphi' = \varphi_+ - \varphi_-$, so we can rewrite inequality (C5) as

$$\varphi_+(\mu) - \varphi_-(\mu) \;<\; -3\,\epsilon \;<\; 0 \;<\; 3\,\epsilon \;<\; \varphi_+(\rho) - \varphi_-(\rho) \qquad \text{for all } \mu \in \mathcal{P}_* \text{ and } \rho \in \overline{\mathcal{P}}.$$

In other words,

$$\varphi_+(\mu) \;<\; \varphi_-(\mu) - 3\,\epsilon \text{ for all } \mu \in \mathcal{P}_*, \text{ whereas } \varphi_+(\rho) \;>\; \varphi_-(\rho) + 3\,\epsilon \quad \text{for all } \rho \in \overline{\mathcal{P}}. \tag{C6}$$

Now, let $\{x, y\}$ be a dichotomy, and assume without loss of generality that $u_j(x) = 1$ and $u_j(y) = 0$ for all $j \in \mathcal{J}$, as in the proof of "(c)$\Longrightarrow$(b)". Lemma C.2 yields piecewise $\mathcal{R}$-almost-objective dichotomous acts $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ and $\boldsymbol{\beta} = (\beta^n)_{n=1}^\infty$ such that for all $j \in \mathcal{J}$, and all $\rho_j \in \mathcal{P}_j$,

$$\lim_{n\to\infty} \int_{\mathcal{S}} u_j \circ \alpha^n \, \mathrm{d}\rho_j \;=\; \varphi_+(\rho_j) \quad \text{and} \quad \lim_{n\to\infty} \int_{\mathcal{S}} u_j \circ \beta^n \, \mathrm{d}\rho_j \;=\; \varphi_-(\rho_j). \tag{C7}$$

Now let $i \in \mathcal{I}$. Since $\mathcal{P}_i$ is compact in the total variation norm, there is a finite subset $\{\lambda_i^\ell\}_{\ell=1}^{L_i} \subset \mathcal{P}_i$ that is ϵ-***dense*** in $\mathcal{P}_i$, in the sense that for any $\rho \in \mathcal{P}_i$, we have $\|\rho - \lambda_i^\ell\|_{\mathrm{var}} < \epsilon$ for some $\ell \in [1 \ldots L_i]$. For all $\ell \in [1 \ldots L_i]$, the right inequality in statement (C6) applies to λ_i^ℓ, because $\mathcal{P}_i \subseteq \overline{\mathcal{P}}$. Combining this inequality with the limit equations (C7) yield some $N_i^\ell \in \mathbb{N}$ such that

$$\int_{\mathcal{S}} u_i \circ \alpha^n \, \mathrm{d}\lambda_i^\ell \;>\; 2\,\epsilon + \int_{\mathcal{S}} u_i \circ \beta^n \, \mathrm{d}\lambda_i^\ell, \quad \text{for all } n \geqslant N_i^\ell. \tag{C8}$$

Let $N_i := \max\{N_i^\ell\}_{\ell=1}^{L_i}$. For any $n \geqslant N_i$, the inequality (C8) holds for all $\ell \in [1 \ldots L_i]$. Now let $\rho \in \mathcal{P}_i$ be arbitrary. By construction, there is some $\ell \in [1 \ldots L_i]$ such that $\|\rho - \lambda_i^\ell\|_{\mathrm{var}} < \epsilon$. Thus, for any $n \geqslant N_i$,

$$\left| \int_{\mathcal{S}} u_i \circ \alpha^n \, \mathrm{d}\rho - \int_{\mathcal{S}} u_i \circ \alpha^n \, \mathrm{d}\lambda_i^\ell \right| \;\leqslant\; \|u_i \circ \alpha_n\|_\infty \cdot \|\rho - \lambda_i^\ell\|_{\mathrm{var}} \;<\; \epsilon,$$

$$\text{and likewise,} \quad \left| \int_{\mathcal{S}} u_i \circ \beta^n \, \mathrm{d}\rho - \int_{\mathcal{S}} u_i \circ \beta^n \, \mathrm{d}\lambda_i^\ell \right| \;<\; \epsilon, \tag{C9}$$

where we use the fact that $\|u_i \circ \alpha_n\|_\infty = \|u_i \circ \beta_n\|_\infty = 1$ because $\alpha_n(\mathcal{S}) = \beta_n(\mathcal{S}) = \{x, y\}$ and $u_i(\{x, y\}) = \{0, 1\}$. Combining inequalities (C8) and (C9), we get

$$\int_{\mathcal{S}} u_i \circ \alpha^n \, \mathrm{d}\rho \;>\; \int_{\mathcal{S}} u_i \circ \beta^n \, \mathrm{d}\rho, \quad \text{for all } \rho \in \mathcal{P}_i, \tag{C10}$$

and thus $\alpha^n \rhd_i \beta^n$. This holds for all $n \geqslant N_i$, so $\boldsymbol{\alpha} \rhd_i^\omega \boldsymbol{\beta}$. This holds for all $i \in \mathcal{I}$.

Now let $\rho_o \in \mathcal{P}_*$ be arbitrary. The limit equations (C7) and the left inequality in statement (C6) yield some $N \in \mathbb{N}$ such that

$$\int_{\mathcal{S}} u_o \circ \alpha^n \, \mathrm{d}\rho_o \;<\; \int_{\mathcal{S}} u_o \circ \beta^n \, \mathrm{d}\rho_o \quad \text{for all } \rho_o \in \mathcal{P}_* \text{ and } n \geqslant N. \tag{C11}$$

Since $\mathcal{P}_* \subseteq \mathcal{P}_o$, this means it is impossible that $\alpha_n \trianglerighteq \beta_n$. This holds for all $n \geqslant N$; thus, it is not the case that $\boldsymbol{\alpha} \trianglerighteq_o^\omega \boldsymbol{\beta}$. This contradicts $\mathcal{R}$-Dichotomous piecewise almost-objective Pareto. □

D Proofs of other results

This appendix contains proofs of additional statements made in the text, regarding the uniqueness of GH representations, the logical relationship between utilitarianism and weak utilitarianism, and observations made in Section 6. The first two proofs uses Propositions 2 and 6. But the other proofs are logically independent from the rest of the paper.

Uniqueness. Proposition 1 is actually a consequence of a more general result.

Proposition D.1 *Let $\mathcal{R} \subseteq \Delta(\mathcal{S})$ be consilient. Let $\succcurlyeq$ be a preference order on $\mathcal{A}$, and let $V_1, V_2 : \mathcal{A} \longrightarrow \mathbb{R}$ be two compact GH representations of $\succcurlyeq$ with utility functions $u_1, u_2 : \mathcal{X} \longrightarrow \mathbb{R}$ and belief sets $\mathcal{P}_1, \mathcal{P}_2 \subseteq \mathcal{R}$. Then*

(a) *There exist constants $a > 0$ and $b \in \mathbb{R}$ such that $u_1 = a\, u_2 + b$.*

(b) *If $\succcurlyeq$ satisfies* Certainty equivalence*, then also $V_1 = a\, V_2 + b$.*

(c) *If, furthermore, $\mathcal{P}_1$ and $\mathcal{P}_2$ are minimal, then $\mathcal{P}_1 = \mathcal{P}_2$.*

Proof: Part (c) follows from (a) and (b). To prove part (a), recall that for all $\alpha \in \mathcal{A}$,

$$\inf_{\rho \in \mathcal{P}_1} \int_{\mathcal{S}} u_1 \circ \alpha \; \mathrm{d}\rho \;\leqslant\; V_1(\alpha) \;\leqslant\; \sup_{\rho \in \mathcal{P}_1} \int_{\mathcal{S}} u_1 \circ \alpha \; \mathrm{d}\rho, \quad \text{and} \tag{D1}$$

$$\inf_{\rho \in \mathcal{P}_2} \int_{\mathcal{S}} u_2 \circ \alpha \; \mathrm{d}\rho \;\leqslant\; V_2(\alpha) \;\leqslant\; \sup_{\rho \in \mathcal{P}_2} \int_{\mathcal{S}} u_2 \circ \alpha \; \mathrm{d}\rho. \tag{D2}$$

Let $\boldsymbol{\alpha} = (\alpha^n)_{n=1}^\infty$ and $\boldsymbol{\beta} = (\beta^n)_{n=1}^\infty$ be compatible $\mathcal{R}$-almost-objective acts, and suppose that $\boldsymbol{\alpha} \succ^\infty \boldsymbol{\beta}$. Then Lemma B.2 yields $\epsilon_1, \epsilon_2 > 0$ such that for all sufficiently large $n \in \mathbb{N}$, we have $V_1(\alpha^n) > V_1(\beta^n) + \epsilon_1$ and $V_2(\alpha^n) > V_2(\beta^n) + \epsilon_2$.

Suppose $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are subordinate to the almost-objectively uncertain partition sequence $\boldsymbol{\mathcal{G}} = (\mathfrak{G}^n)_{n=1}^\infty$, where $\mathfrak{G}^n = \{\mathcal{G}_1^n, \ldots, \mathcal{G}_K^n\}$ for all $n \in \mathbb{N}$, and suppose $\boldsymbol{\mathcal{G}}$ is subordinate to the probability vector $\mathbf{q} = (q_1, \ldots, q_K)$. Suppose $\boldsymbol{\alpha}$ is subordinate to the K-tuple $(x_1, \ldots, x_K) \in \mathcal{X}^K$, while $\boldsymbol{\beta}$ is subordinate to the K-tuple $(y_1, \ldots, y_K)$. Then Proposition 6 says that $\lim_{n\to\infty} V_1(\beta^n) = \sum_{k=1}^K q_k\, u_1(y_k)$ and $\lim_{n\to\infty} V_2(\beta^n) = \sum_{k=1}^K q_k\, u_2(y_k)$.

Thus, since $V_1(\alpha^n) > V_1(\beta^n) + \epsilon_1$ and also $V_2(\alpha^n) > V_2(\beta^n) + \epsilon_2$ for all sufficiently large $n \in \mathbb{N}$, we conclude that

$$\sum_{k=1}^K q_k\, u_1(x_k) \;\geqslant\; \sum_{k=1}^K q_k\, u_1(y_k) + \epsilon_1 \quad \text{and} \quad \sum_{k=1}^K q_k\, u_2(x_k) \;\geqslant\; \sum_{k=1}^K q_k\, u_2(y_k) + \epsilon_2. \tag{D3}$$

Now, by a suitable choice of almost-objective acts $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, we can achieve versions of (D3) for any $\epsilon_1, \epsilon_2 > 0$ and $K \in \mathbb{N}$, any probability vector $\mathbf{q} \in \Delta^K$ and any K-tuples of outcomes $(x_1, \ldots, x_K)$ and $(y_1, \ldots, y_K)$. We conclude that for all $K \in \mathbb{N}$, all $\mathbf{q} \in \Delta^K$ and all $(x_1, \ldots, x_K)$ and $(y_1, \ldots, y_K)$ in $\mathcal{X}^K$,

$$\left(\sum_{k=1}^{K} q_k\, u_1(x_k) > \sum_{k=1}^{K} q_k\, u_1(y_k)\right) \iff \left(\sum_{k=1}^{K} q_k\, u_2(x_k) > \sum_{k=1}^{K} q_k\, u_2(y_k)\right). \tag{D4}$$

By standard uniqueness theorems for SEU representations, it follows from (D4) that u_1 is a positive affine transformation of u_2 —in other words, there exist $a > 0$ and $b \in \mathbb{R}$ such that $u_1 = a\, u_2 + b$. This proves (a).

To prove (b), suppose that $\succcurlyeq$ satisfies Certainty equivalence. Let $\mathcal{V} := V_2(\mathcal{A}) \subseteq \mathbb{R}$. V_1 and V_2 both represent $\succcurlyeq$, so there is an increasing function $\phi : \mathcal{V} \longrightarrow \mathbb{R}$ such that $V_1 = \phi \circ V_2$. We must show that $\phi(x) = a\, v + b$ for all $v \in \mathcal{V}$.

For any $v \in \mathcal{V}$, there is some $\alpha \in \mathcal{A}$ such that $v = V_2(\alpha)$. By Certainty equivalence, there is some constant act κ such that $\alpha \approx \kappa$. Thus, $V_2(\kappa) = V_2(\alpha)$. If κ has the constant value x, then the inequalities (D2) force $V_2(\kappa) = u_2(x)$. Thus, $u_2(x) = v$.

By a similar argument $V_1(\alpha) = u_1(x) = a\, u_2(x) + b = a\, v + b$. But we also have $\phi \circ V_2(\alpha) = V_1(\alpha)$. Thus, we get: $\phi(v) = a\, v + b$, as desired. This argument works for any $v \in \mathcal{V}$. We conclude that $V_1 = a\, V_2 + b$. □

Proof of Proposition 1. Let $\mathcal{P}_*$ and $\mathcal{P}'_*$ be the minimal belief sets for the GH representations V and V'. Let $\mathcal{R} = \mathcal{P}_* \cup \mathcal{P}'_*$. Then $\mathcal{R}$ is nonatomic and separable, so Proposition 2 says that $\mathcal{R}$ is consilient. Thus, Proposition D.1 says that $\mathcal{P}_* = \mathcal{P}'_*$ and $V' = a\, V + b$ (because $\succcurlyeq$ satisfies Certainty equivalents). □

Utilitarianism vs. weak utilitarianism. To explain the logical relationship between these two concepts, we need two hypotheses: *Ex post Pareto* and *Independent prospects.* The social preference $\succcurlyeq_o$ satisfies the Ex post Pareto axiom with respect to $\{\succcurlyeq_i\}_{i \in \mathcal{I}}$ if, for any constant acts $\alpha, \beta \in \mathcal{A}$,

- If $\alpha \succcurlyeq_i \beta$ for all $i \in \mathcal{I}$, then $\alpha \succcurlyeq_o \beta$.
- If, in addition, $\alpha \succ_i \beta$ for some $i \in \mathcal{I}$, then $\alpha \succ_o \beta$.

Now suppose that each of the preference orders $\{\succcurlyeq_j\}_{j \in \mathcal{J}}$ has a GH representation (3) with an associated utility function $u_j : \mathcal{X} \longrightarrow \mathbb{R}$. We shall say that the collection $\{u_i\}_{i \in \mathcal{I}}$ satisfies *Independent Prospects* if, for all $j \in \mathcal{J}$, there exist outcomes $x, y \in \mathcal{X}$ such that $u_j(x) > u_j(y)$ whereas $u_i(x) = u_i(y)$ for all $i \in \mathcal{I} \setminus \{j\}$.

Proposition D.2 *Suppose $\{u_i\}_{i \in \mathcal{I}}$ satisfy* Independent Prospects. *Then u_o is utilitarian if and only if it is weakly utilitarian and $\succcurlyeq_o$ satisfies* Ex post Pareto *for $\{\succcurlyeq_i\}_{i \in \mathcal{I}}$.*

Proof: By definition, if u_o is utilitarian, then it is weakly utilitarian. We will just show that ex post Pareto is satisfied. Let α and β be two constant acts such that $\alpha \succcurlyeq_i \beta$ for all i. Assume that $\alpha(s) = x$ and $\beta(s) = y$ for all states $s \in \mathcal{S}$. We will have $V_i(\alpha) = u_i(x)$ and $V_i(\beta) = u_i(y)$, for all $i \in \mathcal{J}$. Thus, with $u_i(x) \geqslant u_i(y)$ for all $i \in \mathcal{I}$ and $u_o = b + \sum_{i\in\mathcal{I}} c_i u_i$ we have $u_o(x) \geqslant u_o(y)$. Furthermore, if there is $i \in \mathcal{I}$ such that $u_i(x) > u_i(y)$, since $c_i > 0$, we will obviously have $u_o(x) > u_o(y)$.

Conversly, if u_o is weakly utilitarian, then for all $i \in \mathcal{I}$, there is $c_i \geqslant 0$ such that $u_o = b + \sum_{i\in\mathcal{I}} c_i\, u_i$. Let $i \in \mathcal{I}$. To show that $c_i > 0$ let $x_i, y_i \in \mathcal{X}$ such that $u_i(x_i) > u_i(y_i)$ and $u_j(x_i) = u_j(y_i)$ for $j \neq i$; this exists by the hypothesis of Independent Prospects. Considering the constant acts $\alpha_i(s) = x_i$ and $\beta_i(s) = y_i$, we have $V_j(\alpha_i) \geqslant V_j(\beta_i)$ for all $j \in \mathcal{I}$ and $V_i(\alpha_i) > V_i(\beta_i)$. By Ex post Pareto, we have $V_o(\alpha_i) > V_o(\beta_i)$. Thus, $u_o(x_i) - u_o(y_i) = c_i(u_i(x_i) - u_i(y_i)) > 0$. But since $(u_i(x_i) - u_i(y_i)) > 0$, we get $c_i > 0$. □

Proof of Theorem 3. Following the terminology of Gilboa et al. (2004), let us say that an act is a ***lottery*** if it is measurable with respect to some consensus partition.

For all $j \in \mathcal{J}$, suppose $\mathcal{P}_j$ is the convex hull of some finite collection $\mathcal{R}_j := \{\rho_j^1, \ldots, \rho_j^{N_j}\}$ of nonatomic probability measures. For all $j \in \mathcal{J}$ and all $n \in [1 \ldots N_j]$, let $\succcurlyeq_j^n$ be the preference order on $\mathcal{A}$ defined by the SEU representation with utility function u_j and probability measure ρ_j^n. Clearly, a partition of $\mathcal{S}$ is a consensus partition for the original agents $\{\succcurlyeq_j\}_{j\in\mathcal{J}}$ if and only if it is a consensus partition for the new agents $\{\succcurlyeq_j^n\}_{j\in\mathcal{J}}^{n\in[1\ldots N_j]}$; thus, an act is a lottery for the former group of agents if and only if it is a lottery for the latter group. Thus, the scope of the Restricted Pareto axiom for the former group is exactly the same as the scope of this axiom for the latter group.

Furthermore, if $\alpha \in \mathcal{A}$ is a lottery, then for all $j \in \mathcal{J}$, it is easily checked that

$$V_j(\alpha) \;=\; \int_{\mathcal{S}} u_j \circ \alpha \; \mathrm{d}\rho_j^1 \;=\; \int_{\mathcal{S}} u_j \circ \alpha \; \mathrm{d}\rho_j^2 \;=\; \cdots \;=\; \int_{\mathcal{S}} u_j \circ \alpha \; \mathrm{d}\rho_j^{N_j}.$$

In other words, the agents $\{\succcurlyeq_j^1, \ldots, \succcurlyeq_j^{N_j}\}$ all have the same preferences over lotteries as the agent $\succcurlyeq_j$. It follows that for any lotteries α and β,

$$\Big(\alpha \succcurlyeq_i \beta \text{ for all } i \in \mathcal{I}\Big) \iff \Big(\alpha \succcurlyeq_i^n \beta \text{ for all } n \in [1 \ldots N_i] \text{ and } i \in \mathcal{I}\Big).$$

and likewise,

$$\Big(\alpha \succcurlyeq_o \beta \text{ for all } i \in \mathcal{I}\Big) \iff \Big(\alpha \succcurlyeq_o^m \beta \text{ for all } m \in [1 \ldots N_o] \text{ and } o \in \mathcal{I}\Big).$$

Thus, $\succcurlyeq_o$ satisfies Restricted Pareto with respect to $\{\succcurlyeq_i\}_{i\in\mathcal{I}}$ if and only if, for *all* $m \in [1 \ldots N_o]$, $\succcurlyeq_o^m$ satisfies Restricted Pareto with respect to $\{\succcurlyeq_i\}_{i\in\mathcal{I}}^{n\in[1\ldots N_i]}$. Thus, the theorem of Gilboa et al. (2004) says that $\succcurlyeq_o$ satisfies Restricted Pareto (the "indifference" part) with respect to $\{\succcurlyeq_i\}_{i \in \mathcal{I}}$ if and only if (1) u_o is a linear combination of $\{u_i\}_{i \in \mathcal{I}}$ and (2) for all $m \in [1 \ldots N_o]$, ρ_o^m is a linear combination of the elements of $\bigcup_{j \in \mathcal{J}} \mathcal{R}_j$. The "strict preference" part of Restricted Pareto ensures that the coefficients in these linear combinations are nonnegative; this means that u_o is weakly utilitarian, and for all $m \in [1 \ldots N_o]$, ρ_o^m is in the convex hull of $\bigcup_{j \in \mathcal{J}} \mathcal{P}_j$. □

References

Aliprantis, C. D., Border, K. C., 2006. Infinite dimensional analysis: A hitchhiker's guide, 3rd Edition. Springer, Berlin.

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

Baccelli, J., 2017. Do bets reveal beliefs? Synthese 194 (9), 3393–3419.

Bewley, T. F., 2002. Knightian decision theory. Part I. Decisions in economics and finance 25 (2), 79–110.

Billot, A., Qu, X., 2021. Utilitarian aggregation with heterogeneous beliefs. American Economic Journal: Microeconomics 13 (3), 112–23.

Brandl, F., 2021. Belief-averaging and relative utilitarianism. Journal of Economic Theory 198, 105368.

Casadesus-Masanell, R., Klibanoff, P., Ozdenoren, E., 2000a. Maxmin expected utility over Savage acts with a set of priors. Journal of Economic Theory 92 (1), 35–65.

Casadesus-Masanell, R., Klibanoff, P., Ozdenoren, E., 2000b. Maxmin expected utility through statewise combinations. Economics Letters 66 (1), 49–54.

Ceron, F., Vergopoulos, V., 2019. Aggregation of Bayesian preferences: unanimity vs monotonicity. Social Choice and Welfare 52 (3), 419–451.

Cerreia-Vioglio, S., Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2011. Rational preferences under ambiguity. Economic Theory 48 (2-3), 341–375.

Chambers, C., Hayashi, T., 2006. Preference aggregation under uncertainty: Savage vs. Pareto. Games Econom. Behav. 54, 430–440.

Chateauneuf, A., Ventura, C., Vergopoulos, V., 2020. A simple characterization of the hurwicz criterium under uncertainty. Revue economique 71 (2), 331–336.

Crès, H., Gilboa, I., Vieille, N., 2011. Aggregation of multiple prior opinions. J. Econom. Theory 146, 2563–2582.

Danan, E., Gajdos, T., Hill, B., Tallon, J.-M., 2016. Robust social decisions. Am. Econ. Rev. 106 (9), 2407–2425.

Desai, N., Critch, A., Russell, S. J., 2018. Negotiable reinforcement learning for Pareto optimal sequential decision-making. In: Advances in Neural Information Processing Systems. pp. 4712–4720.

Dietrich, F., 2021. Fully Bayesian aggregation. Journal of Economic Theory 194, 105255.

Drèze, J., 1987. Essays on economic decisions under uncertainty. Cambridge University Press, Cambridge.

Duffie, D., 2014. Challenges to a policy treatment of speculative trading motivated by differences in beliefs. The Journal of Legal Studies 43 (S2), S173–S182.

Fleurbaey, M., 2018. Welfare economics, risk and uncertainty. Canadian Journal of Economics 51 (1), 5–40.

Fremlin, D., 2006a. Measure Theory, Volume 4: Topological Measure Spaces (Part I). Torres Fremlin, Colchester.

Fremlin, D., 2006b. Measure Theory, Volume 4: Topological Measure Spaces (Part II). Torres Fremlin, Colchester.

Gajdos, T., Tallon, J.-M., Vergnaud, J.-C., 2008. Representation and aggregation of preferences under uncertainty. J. Econom. Theory 141, 68–99.

Gajdos, T., Vergnaud, J.-C., 2013. Decisions with conflicting and imprecise information. Social Choice and Welfare 41 (2), 427–452.

Ghirardato, P., Maccheroni, F., Marinacci, M., 2004. Differentiating ambiguity and ambiguity attitude. Journal of Economic Theory 118 (2), 133–173.

Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2003. A subjective spin on roulette wheels. Econometrica 71 (6), 1897–1908.

Ghirardato, P., Pennesi, D., 2020. A general theory of subjective mixtures. Journal of Economic Theory 188, 105056.

Ghirardato, P., Pennesi, D., 2022. Randomizing without randomness. Economic Theory, 1–29.

Gilboa, I., 1987. Expected utility with purely subjective non-additive probabilities. Journal of Mathematical Economics 16 (1), 65–88.

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

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

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

Hartmann, L., 2022. An axiomatization of α-maxmin expected utility in the Anscombe-Aumann framework, (preprint).
URL http://dx.doi.org/10.2139/ssrn.3866453

Hayashi, T., 2021. Collective decision under ignorance. Social Choice and Welfare 57 (2), 347–359.

Hayashi, T., Lombardi, M., 2019. Fair social decision under uncertainty and responsibility for beliefs. Econ. Theor 67 (4), 775–816.

Hurwicz, L., 1951. Optimality criteria for decision making under ignorance. Cowles commission papers 370.

Karni, E., 1996. Probabilities and beliefs. Journal of Risk and Uncertainty 13 (3), 249–262.

Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. Econometrica 73 (6), 1849–1892.

Machina, M. J., 2004. Almost-objective uncertainty. Economic Theory 24 (1), 1–54.

Machina, M. J., 2005. 'Expected utility/subjective probability' analysis without the sure-thing principle or probabilistic sophistication. Economic Theory 26 (1), 1–62.

Mongin, P., 1995. Consistent Bayesian aggregation. Journal of Economic Theory 66, 313–351.

Mongin, P., 1997. Spurious unanimity and the Pareto principle. Tech. rep., THEMA, Université de Cergy-Pontoise, published as Mongin (2016).

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

Mongin, P., 2016. Spurious unanimity and the Pareto principle. Economics and Philosophy 32, 511–532, (earlier circulated as Mongin (1997)).

Mongin, P., Pivato, M., 2015. Ranking multidimensional alternatives and uncertain prospects. Journal of Economic Theory 157, 146–171.

Mongin, P., Pivato, M., 2016. Social evaluation under risk and uncertainty. In: Adler, M. D., Fleurbaey, M. (Eds.), Handbook of Well-being and Public Policy. Oxford University Press, Ch. 24, pp. 711–745.

Mongin, P., Pivato, M., 2020. Social preference under twofold uncertainty. Economic Theory 70 (3), 633–663.

Nascimento, L., 2012. The ex-ante aggregation of opinions under uncertainty. Theoretical Economics 7 (3), 535–570.

Pivato, M., 2022. Bayesian social aggregation with accumulating evidence. Journal of Economic Theory 200, 105399.

Poincaré, H., 1912. Calcul des probabilités, 2nd Edition. Gauthier-Villars, Paris.

Qu, X., 2017. Separate aggregation of beliefs and values under ambiguity. Economic Theory 63 (2), 503–519.

Savage, L. J., 1954. The Foundations of Statistics. John Wiley & Sons, New York.

Schervish, M. J., Seidenfeld, T., Kadane, J. B., 1990. State-dependent utilities. Journal of the American Statistical Association 85 (411), 840–847.

Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica, 571–587.

Sprumont, Y., 2018. Belief-weighted Nash aggregation of Savage preferences. Journal of Economic Theory 178, 222–245.

Sprumont, Y., 2019. Relative utilitarianism under uncertainty. Social Choice and Welfare 53 (4), 621–639.

Stanca, L., 2021. Smooth aggregation of bayesian experts. Journal of Economic Theory 196, 105308.

Walters, P., 1982. An Introduction to Ergodic Theory. Springer-Verlag, New York.

Zuber, S., 2016. Harsanyi's theorem without the sure-thing principle: On the consistent aggregation of Monotonic Bernoullian and Archimedean preferences. Journal of Mathematical Economics 63, 78–83.