

Savage’s theorem under changing awareness

Franz Dietrich¹

This version: May 2016

Abstract

This paper proposes a simple unified framework of changing awareness, addressing both *outcome* and (*nature*) *state* awareness, and both how *fine* and how *exhaustive* the awareness is. Six axioms characterize an (essentially unique) expected-utility representation of preferences, in which utilities and probabilities are revised systematically under changes in awareness. Revision is governed by three well-defined rules: (R1) certain utilities are transformed affinely, (R2) certain probabilities are transformed proportionally, and (R3) certain (‘objective’) probabilities are preserved. Rule R2 parallels Karni and Viero’s (2013) ‘reverse Bayesianism’ and Ahn and Ergin’s (2010) ‘partition-dependence’. Savage’s (1954) theorem emerges in the special case of fixed awareness. The theorem draws mathematically on Kopylov (2007), Niiniluoto (1972) and Wakker (1981).

Keywords: Decision under uncertainty, outcome unawareness versus state unawareness, non-refinement versus non-exhaustiveness, utility revision versus probability revision

1 Introduction

Savage’s (1954) expected-utility framework is the cornerstone of modern decision theory. A widely recognized problem is that Savage relies on ready-made and fixed concepts of outcomes and (*nature*) states. These concepts are taken to be stable, as well as highly sophisticated: ideally, *outcomes* capture everything that matters ultimately, and *states* everything that influences outcomes of actions. This ideal translates partly into Savage’s axioms, which imply high ‘state sophistication’ (i.e., infinitely many states), while permitting low ‘outcome sophistication’ (i.e., possibly just two outcomes). In sum, Savage’s theory is committed to stable outcome/state awareness and sophisticated state awareness.

A real agent’s awareness can be limited on two levels in two ways. It can be limited at the *outcome* and *state* level, and it can be *non-fine* (coarse) and *non-exhaustive* (domain-restricted). Consider a social planner deciding where to

¹Paris School of Economics & CNRS; fd@franzdietrich.net; www.franzdietrich.net. Address: Centre d’Economie de la Sorbonne, 106-112 Boulevard de l’Hôpital, 75013 Paris, France.

build a new nuclear power plant on his island. He has a non-exhaustive state concept if he fails to foresee some contingencies such as a Tsunami. He has a non-fine state concept if he conceives a Tsunami as a *primitive* possibility rather than decomposing it into the (sub)possibilities of a Tsunami from the east, west, north, or south. These are examples of *state* unawareness; analogous examples exist for *outcome* unawareness. Figure 1 gives a formal illustration with four ‘objective’




Figure 1: An act f for non-fine and non-exhaustive concepts of states and outcomes

states resp. outcomes from an external perspective, but only two subjectively conceived states resp. outcomes. The concepts are non-fine: s and t are lumped into the same state, and y and z into the same outcome. The concepts are also non-exhaustive: q and w are ignored, i.e., excluded by all conceived states resp. outcomes. State/outcome unawareness translates into act unawareness: if as in Figure 1 only two states resp. outcomes are conceived, then only $2^2 = 4$ acts (functions from states to outcomes) are conceived.

There is a clear need for a generalization of Savage’s expected-utility theory to cope with changes in awareness of the various sorts. If such a generalization has not yet been offered, it is possibly because of an obstacle: Savage’s high demands of ‘state sophistication’ conflict with (state) unawareness. Overcoming this obstacle, I offer a Savagean expected-utility theory under changing awareness, involving ‘rational’ revision rules. Future research might move towards *non*-expected-utility representations and/or ‘*boundedly* rational’ revision rules. But since such issues are orthogonal to the issue of awareness change, good scientific practice tells us to first develop a general understanding of ‘rational’ decision and ‘rational’ revision under changing awareness, thereby creating a solid starting point for future relaxations.

In short, I propose a simple unified model of changing awareness, capturing changes in outcome as well as state awareness, and in refinement as well as exhaustiveness. Six axioms are shown to characterize an expected-utility agent who uses three revision rules to update utilities and probabilities when his outcome/state concepts change:

- R1:** *utilities of unaffected outcomes are transformed in an increasing affine way;*
- R2:** *probabilities of unaffected events are transformed proportionally;*
- R3:** *‘objective’ probabilities (in an endogenous sense) are preserved.²*

²I.e., events that are ‘risky’ (in an endogenous sense) have description-invariant probabilities.

Probabilities are unique; so R2's coefficient of proportionality is unique. Utilities are essentially unique. Utility revision is a genuine feature: utilities cannot generally be normalised such that R1's transformation is always the identity transformation. The theorem addresses the two problems raised at the outset: it permits *unstable* and *unsophisticated* awareness, of both outcomes and states. Further, it generalizes Savage's Theorem: it reduces to it in case of stable awareness, as our axioms then reduce to Savage's axioms, while rules R1 and R2 hold trivially and R3 can be shown to reduce to Savage's *atomlessness* condition on probabilities.

To my knowledge, the current framework and theorem are new. I wish to relate the paper to two seminal contributions, the Ahn-Ergin (2010) model of framed contingencies and the Karni-Viero (2013) model of growing awareness. Ahn and Ergin (2010) assume that each of various possible 'framings' of the relevant contingencies leads to a particular partition of the objective state space (representing the agent's state concept), and to a particular preference relation over those acts which are measurable relative to that partition. Under plausible axioms on partition-dependent preferences, they derive a compact expected-utility representation with fixed utilities and partition-dependent probabilities. The systematic way in which these probabilities change with the partition implies our rule R2 (without an equivalence). Karni and Viero (2013), by contrast, model the discovery of new acts, outcomes, and act-outcome links. Given their goal, they use a non-Savagean framework (going back to Schmeidler and Wakker 1987 and Karni and Schmeidler 1991) which takes acts as primitive objects and states as functions from acts to outcomes. They characterize preference change under growing awareness, using various combinations of axioms. A key finding is that probabilities are revised in a *reverse Bayesian* way, a property once again related to our revision rule R2. The compatibility of R2 with Ahn-Ergin's and Karni-Viero's findings on belief revision confirms the robustness of their findings.

The current analysis differs strongly from Ahn-Ergin's and Karni-Viero's. I now mention some differences. I analyse awareness change on both levels (outcomes and states) and of both kinds (refinement and exhaustiveness), while Ahn-Ergin limit attention to changes in state refinement (with fixed state exhaustiveness and fixed outcome awareness), and Karni-Viero assume fixed outcome refinement.³ Ahn-Ergin and Karni-Viero find that only probabilities are revised, yet I find that also utilities are revised. Ahn-Ergin and Karni-Viero introduce lotteries as primitives (following Anscombe and Aumann 1963), while I invoke no exogenous objective probabilities (following Savage 1954). Ahn-Ergin and Karni-Viero exclude the classical base-line case of 'state sophistication' with an infinite state space, while I allow 'state sophistication' to be reached sometimes (or *always*, or *never*); this

³Karni-Viero do capture changes in outcome *exhaustiveness*, through the discovery of new outcomes. Changes in state awareness are captured indirectly: the discovery of new acts resp. new outcomes effectively *refines* states resp. renders states more *exhaustive*.

flexibility is crucial for ‘generalizing Savage’. Karni-Viero invoke different axioms for different types of awareness change (such as the discovery of new outcomes), while I use a unified set of axioms.

The theorem’s long proof, presented in different appendices, makes use of key theorems by Kopylov (2007), Niiniluoto (1972) and Wakker (1981). In the background of the paper is a vast and active literature on unawareness (e.g., Dekel, Lipman and Rustichini 1998, Halpern 2001, Halpern and Rego 2008, Hill 2010, Pivato and Vergopoulos 2015, Karni and Viero 2015). I do not attempt to review this diverse body of work, ranging from epistemic to choice-theoretic studies, from static to dynamic studies, and from decision- to game-theoretic studies.

2 A unified model of changing awareness

2.1 The variable Savage framework

Before introducing our own primitives, I recall Savage’s original primitives:

Definition 1 *A **Savage framework** is a triple (X, S, \succsim) of a non-empty finite⁴ set X (of **outcomes** or **consequences**), a non-empty set S (of **states**), and a (**preference**) relation \succsim on the set of functions from S to X (**acts**).*

I replace Savage’s fixed outcome/state spaces by context-dependent ones. This leads to a family of Savage frameworks $(X_\alpha, S_\alpha, \succsim_\alpha)$ where α ranges over a set of contexts. I take each X_α to partition (coarsen) some underlying set of ‘objective’ outcomes, and each S_α to partition (coarsen) some underlying set of ‘objective’ states. This captures changing awareness of the ‘objective’ world.⁵

Definition 2 *A **variable Savage framework** is a family of Savage frameworks $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ indexed by some non-empty set Γ (of **contexts**), where*

- *each X_α is a partition of some set (of **objective outcomes encompassed in context** α),*
- *each S_α is a partition of some set (of **objective states encompassed in context** α).*

*An **objective outcome** resp. **state simpliciter** is an objective outcome resp. state encompassed in at least one context.*

⁴Savage in fact did not impose finiteness. I add finiteness for simplicity.

⁵A **partition** of a set is a set of *non-empty*, pairwise exclusive and exhaustive subsets.

From now on, let $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ be a given variable Savage framework. Let

$$\begin{aligned}
F_\alpha &:= X_\alpha^{S_\alpha} && (\text{set of acts conceived in context } \alpha \in \Gamma), \\
X_\alpha &:= \cup_{x \in X_\alpha} x && (\text{set of objective outcomes encompassed in context } \alpha \in \Gamma) \\
S_\alpha &:= \cup_{s \in S_\alpha} s && (\text{set of objective states encompassed in context } \alpha \in \Gamma) \\
\mathbf{X} &:= \cup_{\alpha \in \Gamma} X_\alpha && (\text{set of objective outcomes}) \\
\mathbf{S} &:= \cup_{\alpha \in \Gamma} S_\alpha && (\text{set of objective states}) \\
\mathbf{F} &:= \mathbf{X}^\mathbf{S} && (\text{set of objective acts}).
\end{aligned}$$

The spaces X_α and S_α could represent the *framing* of options in context α , e.g., the *mode of presentation* or *level of descriptive detail* (see Section 2.2). Here is a two-context example: let $\Gamma = \{\alpha, \beta\}$, where $X_\alpha = \{\{x\}, \{y, z\}\}$, $S_\alpha = \{\{r\}, \{s, t\}\}$, $X_\beta = \{\{x, y\}, \{z, w\}\}$, and $S_\beta = \{\{q\}, \{r, s, t\}\}$; so $\mathbf{X}_\alpha = \{x, y, z\}$, $\mathbf{S}_\alpha = \{r, s, t\}$, $\mathbf{X} = \mathbf{X}_\beta = \{x, y, z, w\}$, and $\mathbf{S} = \mathbf{S}_\beta = \{q, r, s, t\}$. The context α was illustrated in Figure 1. The outcome/state spaces are non-fine in both contexts, and non-exhaustive only in context α . In general, the smaller the sets in X_α and S_α are, the finer the agent's outcome/state concepts are, up to the point of singleton sets (full refinement). The larger the sets \mathbf{X}_α and \mathbf{S}_α are, the more exhaustive these concepts are, up to the point of the entire sets \mathbf{X} and \mathbf{S} (full exhaustiveness).

When does the agent have full awareness of *some* type at *some* level?

Definition 3 *The variable Savage framework has*

- (a) **exhaustive outcomes** if $X_\alpha = \mathbf{X}$ in all contexts $\alpha \in \Gamma$,
- (b) **exhaustive states** if $S_\alpha = \mathbf{S}$ in all contexts $\alpha \in \Gamma$,
- (c) **fine outcomes** if all outcomes $x \in \mathbf{X}_\alpha$ are singleton in all contexts $\alpha \in \Gamma$,
- (d) **fine states** if all states $s \in \mathbf{S}_\alpha$ are singleton in all contexts $\alpha \in \Gamma$.

Our theorem will simplify under exhaustive states, and simplify *differently* under fine states. Examples demonstrate the generality and flexibility of our model:

- **Example 1: Savage.** Γ contains a single context γ . Our variable framework reduces to a classic Savage framework $(X, S, \succsim) := (X_\gamma, S_\gamma, \succsim_\gamma)$. Objective outcomes and states are not needed: w.l.o.g. we can, like Savage, let X and S be primitive sets, rather than partitions.
- **Example 2: stable outcome awareness.** All contexts α lead to the same outcome space $X_\alpha = X$, which we may take as a primitive set, not a partition.
- **Example 3: stable state awareness.** All contexts α lead to the same state space $S_\alpha = S$, which we may take to be a primitive set, not a partition.
- **Example 4: fully variable awareness.** All logically possible awareness states occur: for all partitions X and S of \mathbf{X} resp. \mathbf{S} (or of non-empty subsets of \mathbf{X} resp. \mathbf{S} , to allow non-exhaustive awareness), where $|X| < \infty$, there is a context $\alpha \in \Gamma$ in which $X_\alpha = X$ and $S_\alpha = S$. This permits arbitrary ways to conceive the world.

- **Example 5: finite awareness.** All spaces X_α and S_α , and so all act sets F_α , are finite. The agent can only conceive finitely many things at a time.
- **Example 6: contexts as awareness states.** Let each context *be*, not just *induce*, a pair of an outcome space and a state space. Formally, Γ is a set of pairs of partitions (X, S) (the ‘possible’ awareness states). The framework $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ can then be abbreviated as $(\succsim_\alpha)_{\alpha \in \Gamma}$, as each context $\alpha = (X, S) \in \Gamma$ already contains the information of the spaces $X_\alpha := X$ and $S_\alpha := S$. Such a ‘compact framework’ $(\succsim_\alpha)_{\alpha \in \Gamma}$ is the slimmest point of departure for studying the effect of awareness on preference. It uses no independent ‘context’ notion, be this an advantage or a loss.

Throughout I assume *independence between outcome and state awareness*: the agent’s outcome awareness and state awareness do not constrain one another. Formally, any occurring outcome and state spaces X_α and S_β ($\alpha, \beta \in \Gamma$) can occur *jointly*, i.e., some context $\gamma \in \Gamma$ has $X_\gamma = X_\alpha$ and $S_\gamma = S_\beta$.⁶

2.2 Four interpretive remarks

1. One might compare objective and subjective states with Savage’s (1954) *grand-world* resp. *small-world* states, although he takes both types of states to be fixed.
2. The spaces X_α and S_α ($\alpha \in \Gamma$) represent the awareness (concepts) *ascribed* to the agent *by the observer*.⁷ The ascription could be based on the framing effects which are at work in a context and render certain outcome/state concepts salient, perhaps through an explicit *mode of presentation*, following Ahn and Ergin (2010) and extending their idea also to outcomes. If the agent is presented car insurance policies in terms of their net benefit as a function of the number (up to 10) of accidents, then S_α contains the 11 events ‘ n accidents’ for $n = 0, 1, \dots, 10$, and X_α contains the 11 net-benefit outcomes; another context β with a different mode of presentation will induce different spaces S_β and X_β . Framing effects are important, but by far not the only possible basis for ascribing spaces X_α and S_α ($\alpha \in \Gamma$) to the agent.⁸ Moreover, one could take the spaces X_α and S_α to represent the agent’s *real*

⁶This excludes that the agent conceives the outcome ‘I am popular’ only when conceiving the state ‘I win in the lottery’, or that he conceives fine outcomes only when conceiving coarse states.

⁷So X_α and S_α reflect how we take him to perceive or describe the world in context α . They embody our hypothesis (or theory, stipulation, conjecture etc.) about the agent’s awareness.

⁸At least in principle, the ascription could also be based on (i) common sense and intuition; or (ii) neurophysiological evidence about how the context affects the cognitive system; or (iii) the sort of options that are *feasible* in the context (here X_α and S_α are constructed such that all feasible options become representable as subjective acts, in a sense made precise in Section 2.4); or (iv) patterns of *observed choice* that are taken to *reveal* the agent’s awareness, in a sense that can be made precise (here X_α and S_α are constructed so as to be fine enough to distinguish between those *objective* acts between which observed behaviour distinguishes).

rather than *ascribed* awareness in context α , adopting a first-person rather than third-person perspective. The same two interpretations are also commonly applied to a standard Savage framework (X, S, \succsim) : its spaces X and S could represent the agent's *ascribed* or *real* outcome/state concepts. Savage himself had the second interpretation in mind: he focused on the notion of rationality rather than on an observer's third-person perspective.

3. By modelling outcomes and states as sets of objective outcomes resp. states, I by no means suggest that the agent subjectively conceives outcomes and states *in terms of* (complex) sets. He may conceive them as indecomposable primitives. He may conceive the outcome 'having close friends' in complete unawareness of the huge (infinite) set of underlying objective outcomes.

4. Crucially, the agent may in one context α conceive an event $E \subseteq S_\alpha$ and in another context β conceive a different event $E' \subseteq S_\beta$, where E and E' represent the same *objective* event, and yet the agent attaches a different probability to E (in context α) than to E' (in context β). The idea is that belief is description-sensitive: it depends on how objective events are perceived subjectively. Imagine that in context α the agent conceives (fine) states $\{s\}$ and $\{t\}$ (where $s, t \in \mathbf{S}$) and hence the event $E = \{\{s\}, \{t\}\}$, while in context β he conceives the (coarser) state $\{s, t\}$ and hence the event $E' = \{\{s, t\}\}$. Although E and E' represent the same objective event $\{s, t\}$, the agent might in context α find E unlikely on the grounds that $\{s\}$ and $\{t\}$ each appear implausible, while in context β finding E' likely because he fails to analyse this event in terms of its implausible subcases.⁹

2.3 Terminology and notation

The objective/subjective terminology: I carefully distinguish between the two levels of description (often dropping the adjective 'subjective' for brevity):

- An **objective outcome, state, act** resp. **event** is a member of $\mathbf{X}, \mathbf{S}, \mathbf{F}$ ($= \mathbf{X}^{\mathbf{S}}$) resp. $2^{\mathbf{S}}$.
- A **(subjective) outcome, state, act** resp. **event conceived in context** α ($\in \Gamma$) is a member of $X_\alpha, S_\alpha, F_\alpha$ ($= X_\alpha^{S_\alpha}$) resp. 2^{S_α} ; the **(subjective) state space** resp. **outcome space in context** α is S_α resp. X_α .
- A **(subjective) outcome, state, act** or **event** simpliciter (without reference to a context) is one that is conceived in *some* context, i.e., a member of *some* $X_\alpha, S_\alpha, F_\alpha$ resp. 2^{S_α} ($\alpha \in \Gamma$); a **(subjective) outcome space** resp. **state space** simpliciter is *some* X_α resp. S_α ($\alpha \in \Gamma$).

⁹Concretely, s could stand for country S attacking country T, and t for T attacking S. In context α the agent finds event $E = \{\{s\}, \{t\}\}$ unlikely: he reasons that $\{s\}$ and $\{t\}$ are each implausible, as S most probably won't attack T, and vice versa. In context β , he finds event E' likely on unsophisticated grounds: he treats E' as a primitive scenario of 'war', which seems likely to him, failing to realise that a war requires an (unlikely) attack by either country.

Translating between ‘subjective’ and ‘objective’: Given a context $\alpha \in \Gamma$,

- any objective outcome $x \in \mathbf{X}_\alpha$ has a **subjectivization** $x_\alpha \ni x$ in X_α ,
- any objective state $s \in \mathbf{S}_\alpha$ has a **subjectivization** $s_\alpha \ni s$ in S_α ,
- any subjective event $E \subseteq S_\alpha$ induces (i.e., partitions) an objective one denoted $E^* := \cup_{s \in E} s \subseteq \mathbf{S}_\alpha$; E and E^* are said to **correspond** to each other;
- any subjective act $f \in F_\alpha$ induces a function on \mathbf{S}_α denoted f^* and given by $f^*(s) := f(s_\alpha)$; f and f^* are said to **correspond** to each other.

Standard notation: Let f_E be the restriction of function f to its subdomain E . For objective or subjective outcomes x and sets S , let x_S be the function on S with constant value x . For functions f and g on disjoint domains, fg is the function on the union of domains matching f on f ’s domain and g on g ’s domain. Examples are ‘mixed’ acts $f_E g_{S_\alpha \setminus E} \in F_\alpha$, where $f, g \in F_\alpha$ and $E \subseteq S_\alpha$ ($\alpha \in \Gamma$).

2.4 Excursion: awareness and choice behaviour

The setting is easily connected to choice behaviour. Assume the agent finds himself in a context $\alpha \in \Gamma$ and faces a choice between some concrete (pre-theoretic) options, such as meals or holiday destinations. The modeller faces two possibilities: he could model options *either* as subjective acts in F_α *or* as objective acts in \mathbf{F} . Neither possibility is generally superior: all depends on the intended level of description. In the first case, the feasible set is a subset of F_α , and the prediction is simply that a most \succsim_α -preferred member is chosen. For the rest of this subsection, I assume the second case: let options be objective acts. So the feasible set is a subset of \mathbf{F} , not F_α . Which choice does \succsim_α predict? It predicts that the agent chooses a feasible objective act whose *subjective representation* in F_α is most \succsim_α -preferred. I now spell this out formally.

Definition 4 *In a context $\alpha \in \Gamma$, an act in F_α is the **(subjective) representation** of the objective act $f \in \mathbf{F}$, denoted f_α , if it agrees with f ‘modulo subjectivization’: for all $s \in \mathbf{S}$ and $s' \in S_\alpha$, if $s \in s'$ then $f(s) \in f_\alpha(s')$. An $f \in \mathbf{F}$ is **(subjectively) representable** in context α if its representation $f_\alpha \in F_\alpha$ exists.*

Remark 1 *In a context $\alpha \in \Gamma$, an act in F_α is the subjective representation of $f \in \mathbf{F}$, denoted f_α , if and only if $f(\mathbf{S}_\alpha) \subseteq \mathbf{X}_\alpha$ and $f_\alpha(s_\alpha) = [f(s)]_\alpha$ for all $s \in \mathbf{S}_\alpha$ (so the diagram in Figure 2 commutes). The condition simplifies under exhaustive states and outcomes: $f_\alpha(s_\alpha) = [f(s)]_\alpha$ for all $s \in \mathbf{S}$.*

Remark 2 (uniqueness) *Any objective act $f \in \mathbf{F}$ has at most one representation in a context.*




Figure 2: An objective act $f : \mathbf{S} \rightarrow \mathbf{X}$ which is representable in context α (so maps \mathbf{S}_α into \mathbf{X}_α by Remark 1), and the subjective representation $f_\alpha : S_\alpha \rightarrow X_\alpha$

Remark 3 (existence condition) *In a context $\alpha \in \Gamma$, an objective act $f \in \mathbf{F}$ is representable if and only if $f(\mathbf{S}_\alpha) \subseteq \mathbf{X}_\alpha$ and $f_{\mathbf{S}_\alpha}$ is (S_α, X_α) -measurable.¹⁰ The condition simplifies under exhaustive states and outcomes: f is (S_α, X_α) -measurable.*

As an illustration, consider an objective act f that makes the agent rich if a coin lands heads (and poor otherwise), and that might also do many other things, such as making him sick in the event of cold weather. In context α the agent conceives only ‘wealth outcomes’ and ‘coin states’: $X_\alpha = \{r, p\}$ and $S_\alpha = \{h, t\}$, where r and p are the outcomes (sets of objective outcomes) in which he is rich resp. poor, and h and t are the states (sets of objective states) in which the coin lands heads resp. tails. Then f is represented by the subjective act f_α that maps h to r and t to p . But if instead $X_\alpha = \{r, p\}$ and $S_\alpha = \{\mathbf{S}\}$, the state concept no longer captures the coin toss, and f is no longer representable.

I can now define choice predictions: our framework predicts that whenever in a context $\alpha \in \Gamma$ the agent has to choose from a set $A \subseteq \mathbf{F}$ of representable objective acts, then he chooses an $f \in A$ such that $f_\alpha \succsim_\alpha g_\alpha$ for all $g \in A$. (This may lead to choice reversals as the context changes; see Section 4.) No prediction is made about choice from *non-representable* objective acts: the model is silent on such choices. Does the model thereby miss out on many choice situations? Perhaps not, because the mental process of forming outcome/state concepts might (consciously or ‘automatically’) adapt these concepts to the feasible options, to ensure representability. I call the agent – or more exactly his awareness, i.e., the spaces $(X_\alpha, S_\alpha)_{\alpha \in \Gamma}$ – **adaptive (to feasible options)** if whenever in a context $\alpha \in \Gamma$ an objective act $f \in \mathbf{F}$ is feasible, then f is representable in context α .¹¹ The idea is that the agent forms awareness of a coin toss *when and because* feasible

¹⁰ (S_α, X_α) -measurability means that members of the same $s \in S_\alpha$ are mapped into the same $x \in X_\alpha$, or equivalently, that the inverse image of any $x \in X_\alpha$ is a union of zero or more $s \in S_\alpha$.

¹¹A full-fledged definition could state as follows. Let *choice situations* be pairs (A, α) of a non-empty menu $A \subseteq \mathbf{F}$ of (feasible) objective acts and a context or ‘frame’ $\alpha \in \Gamma$ (in which the choice from A is made). Some choice situations occur, others do not. Let \mathcal{CS} be the set of *occurring* (or *feasible*) choice situations. *Adaptiveness (to feasible options)* means that for all $(A, \alpha) \in \mathcal{CS}$ each $f \in A$ is representable in context α .

objective acts depend on it. Forming awareness is a costly mental activity, which is likely to be guided by the needs of real choice situations, including the need to represent feasible options. Adaptiveness can thus be viewed as a rationality requirement on the agent’s awareness.¹²

Is there any way to predict choices *even* when some feasible options are non-representable, i.e., even without adaptiveness? There is indeed, if one is ready to make one of two auxiliary assumptions: one could take non-representable options to be *ignored* (*‘not perceived’*), or rather to be *misrepresented* (*‘misperceived’*).¹³

3 Six axioms

Sections 3–5 temporarily assume *exhaustive states* (see Definition 3). In fact, each axiom, theorem or proposition, and most definitions and remarks, will already be stated in their general form, for possibly non-exhaustive states. For transparency, the three exceptions – two definitions and one remark – will be marked by ‘exh’. So ‘Definition 13_{exh}’ applies only under exhaustive states, but ‘Definition 5’ applies generally. For each exception (identified by ‘exh’), a general re-statement is given in Section 6 where I lift the restriction to exhaustive states.

The current section states six axioms. They are equivalent to Savage’s axioms in the single-context case. I begin with the analogue of Savage’s first axiom:

Axiom 1 (weak order): For all contexts $\alpha \in \Gamma$, \succsim_α is a transitive and complete relation (on F_α).

Savage’s sure-thing principle can be rendered in two ways in our setting, by applying sure-thing reasoning either *within* each context, or even *across* contexts:

Axiom 2* (sure-thing principle, local version): For all contexts $\alpha \in \Gamma$, acts $f, g, f', g' \in F_\alpha$, and events $E \subseteq S_\alpha$, if $f_E = f'_E$, $g_E = g'_E$, $f_{S_\alpha \setminus E} = g_{S_\alpha \setminus E}$ and $f'_{S_\alpha \setminus E} = g'_{S_\alpha \setminus E}$, then $f \succsim_\alpha g \Leftrightarrow f' \succsim_\alpha g'$.

Axiom 2 (sure-thing principle, global version): For all contexts $\alpha, \alpha' \in \Gamma$,

¹²The agent’s awareness (his spaces X_α and S_α) can be ‘irrational’ in two distinct ways, the second way being non-adaptiveness. (1) Outcomes may be too coarse to incorporate all relevant features of objective outcomes that the agent would care about had he considered them (in the above example, health features are absent from $X_\alpha = \{r, p\}$, though presumably relevant). (2) States may be too coarse (given how outcome are specified) for all feasible objective acts to be representable (in the above example, f is not representable if $S_\alpha = \{\mathbf{S}\}$, given that $X_\alpha = \{r, p\}$). In (1) and (2) I assumed exhaustive states and outcomes, but the idea can be generalized.

¹³Under the first hypothesis, the agent considers not the full feasible set, but only the subset of *representable* feasible options (among which he picks an option whose representation is most \succsim_α -preferred). Under the second hypothesis, a non-representable feasible option f in \mathbf{F} is not ignored, but (mis)perceived as some subjective act in F_α which fails to properly represent f . Which is this subjective act? Here one would need to develop a theory of misrepresentation.

acts $f, g \in F_\alpha$ and $f', g' \in F_{\alpha'}$, and events conceived in both contexts $E \subseteq S_\alpha \cap S_{\alpha'}$, if $f_E = f'_E$, $g_E = g'_E$, $f_{S_\alpha \setminus E} = g_{S_\alpha \setminus E}$ and $f'_{S_{\alpha'} \setminus E} = g'_{S_{\alpha'} \setminus E}$, then $f \succsim_\alpha g \Leftrightarrow f' \succsim_{\alpha'} g'$.

Remark 4 *Axiom 2* is the restriction of Axiom 2 to the case that $\alpha = \alpha'$.*

I will employ Axiom 2 rather than 2*. Axiom 2 renders sure-thing reasoning in a particularly rigorous way, applying it all the way through, regardless of irrelevant barriers of context.¹⁴ I now extend four familiar Savagean notions to our setting:

Definition 5 (preferences over outcomes) *In a context $\alpha \in \Gamma$, an outcome $x \in X_\alpha$ is weakly preferred to another $y \in X_\alpha$ – written $x \succsim_\alpha y$ – if $x_{S_\alpha} \succsim_\alpha y_{S_\alpha}$.*

Definition 6 (conditional preferences) *In a context $\alpha \in \Gamma$, an act $f \in F_\alpha$ is weakly preferred to another $g \in F_\alpha$ **given** an event $E \subseteq S_\alpha$ – written $f \succsim_{\alpha, E} g$ – if $f' \succsim_\alpha g'$ for some (hence under Axiom 2 any) acts $f', g' \in F_\alpha$ which agree with f resp. g on E and with each other on $S_\alpha \setminus E$.*

Definition 7 (conditional preferences over outcomes) *In a context $\alpha \in \Gamma$, an outcome $x \in X_\alpha$ is **weakly preferred** to another $y \in F_\alpha$ **given** an event $E \subseteq S_\alpha$ – written $x \succsim_{\alpha, E} y$ – if $x_{S_\alpha} \succsim_{\alpha, E} y_{S_\alpha}$.*

Definition 8 (null events) *In a context $\alpha \in \Gamma$, an event $E \subseteq S_\alpha$ is **null** if it does not affect preferences, i.e., $f \sim_\alpha g$ whenever acts $f, g \in F_\alpha$ agree outside E .*

I am ready to state the analogue of Savage's third axiom:

Axiom 3 (state independence): For all contexts $\alpha \in \Gamma$, outcomes $x, y \in X_\alpha$, and non-null events $E \subseteq S_\alpha$, $x \succsim_{\alpha, E} y \Leftrightarrow x \succsim_\alpha y$.

A bet on an event is an act that yields a 'good' outcome x if this event occurs and a 'bad' outcome y otherwise. Savage's fourth axiom requires preferences over bets to be independent of the choice of x and y . His axiom can again be rendered as an *intra-* or *inter-*context condition:

Axiom 4* (comparative probability, local version): For all contexts $\alpha \in \Gamma$, events $E, D \subseteq S_\alpha$, and outcomes $x \succ_\alpha y$ and $x' \succ_\alpha y'$ in X_α , $x_E y_{S_\alpha \setminus E} \succsim_\alpha x_D y_{S_\alpha \setminus D} \Leftrightarrow x'_E y'_{S_\alpha \setminus E} \succsim_\alpha x'_D y'_{S_\alpha \setminus D}$.

Axiom 4 (comparative probability, global version): For all contexts $\alpha, \alpha' \in \Gamma$ with same state space $S := S_\alpha = S_{\alpha'}$, events $E, D \subseteq S$, and outcomes $x \succ_\alpha y$ in X_α and $x' \succ_{\alpha'} y'$ in $X_{\alpha'}$, $x_E y_{S \setminus E} \succsim_\alpha x_D y_{S \setminus D} \Leftrightarrow x'_E y'_{S \setminus E} \succsim_{\alpha'} x'_D y'_{S \setminus D}$.

Remark 5 *Axiom 4* is the restriction of Axiom 4 to the case that $\alpha = \alpha'$.*

¹⁴Replacing sure-thing reasoning by ambiguity aversion in our setting is an interesting avenue.

I will use Axiom 4 rather than 4*. Axiom 4 applies the reasoning underlying Savage’s fourth axiom all the way through, regardless of barriers of context. Another familiar notion can now be defined in our setting:

Definition 9 (*comparative beliefs*) In a context $\alpha \in \Gamma$, an event $E \subseteq S_\alpha$ is **at least as probable as** another $D \subseteq S_\alpha$ – written $E \succsim_\alpha D$ – if $x_{E \setminus S_\alpha \setminus E} \succsim_\alpha x_{D \setminus S_\alpha \setminus D}$ for some (hence under Axiom 4 any) outcomes $x \succ_\alpha y$ in X_α .

Savage’s fifth and sixth axioms have the following counterparts:

Axiom 5 (non-triviality): For all context $\alpha \in \Gamma$, there are acts $f \succ_\alpha g$ in F_α .

Axiom 6* (Archimedean, local version): For all contexts $\alpha \in \Gamma$, acts $f \succ_\alpha g$ in F_α , and outcomes $x \in X_\alpha$, one can partition S_α into events E_1, \dots, E_n such that $f_{S_\alpha \setminus E_i} x_{E_i} \succ_\alpha g$ and $f \succ_\alpha g_{S_\alpha \setminus E_i} x_{E_i}$ for all E_i .

However, just as Savage’s sixth postulate, Axiom 6* is very demanding. It forces the agent to conceive plenty of small events, ultimately forcing all state spaces S_α to be infinite (assuming Axiom 5 for non-triviality). Our framework allows for a cognitively less demanding Archimedean axiom, which permits all state spaces S_α to be finite. To avoid ‘state-space explosion’, it allows the events E_1, \dots, E_n to be *not yet* conceived: they are conceived in *some* possibly different context β . So the agent can presently have limited state awareness, as long as he can refine states by moving to a new context. The slogan is: ‘*refinable* rather than (already) *refined* states’. Indeed, many real people rarely consider events of probability less than 1%, but are (if needed) perfectly able to conceive them by refining their state concept.¹⁵ The next axiom renders this idea.

Definition 10 Acts $f \in F_\alpha$ and $g \in F_\beta$ ($\alpha, \beta \in \Gamma$) are **(objectively) equivalent** if $f^* = g^*$.

Definition 11 A partition S **refines** or is **at least as fine as** a partition T if, for some equivalence relation on S , $T = \{\cup_{A \in E} A : E \text{ is an equivalence class}\}$.¹⁶

Axiom 6 (Archimedean, global version 1):** For all contexts $\alpha \in \Gamma$, acts $f \succ_\alpha g$ in F_α , and outcomes $x \in X_\alpha$, there is a context $\beta \in \Gamma$ with state space S_β at least as fine as S_α and outcome space $X_\beta \supseteq X_\alpha$ (ensuring that F_β contains acts f' and g' equivalent to f resp. g) such that one can partition S_β into events E_1, \dots, E_n for which $f'_{S_\beta \setminus E_i} x_{E_i} \succ_\beta g'$ and $f' \succ_\beta g'_{S_\beta \setminus E_i} x_{E_i}$ for all E_i .

¹⁵It suffices to incorporate, say, the results of three independent tosses of a fair dice. Here the refined state describes the ‘old’ state *and* the triple of dicing results. The refined state space can thus be partitioned into the $6^3 = 216$ small-probability events of the sort ‘the triple of dicing results is (i, j, k) ’, where $i, j, k \in \{1, 2, \dots, 6\}$.

¹⁶In other words, T **coarsens** or is **at least as coarse as** S .

Remark 6 *Axiom 6* is the restriction of Axiom 6** to the case that $\alpha = \beta$.*

Axiom 6** is not yet fully suitable. It fails to ensure any connection between \succsim_β and \succsim_α , allowing even that $g \succ_\beta f$ although $f \succ_\alpha g$. I thus use a variant of Axiom 6**, which indirectly guarantees a connection. It requires that the objective events represented by E_1, \dots, E_n – say $A_1, \dots, A_n \subseteq \mathbf{S}$ – are of a special ‘innocuous’ kind. Informally, A_1, \dots, A_n must belong to an algebra of *risky* objective events, e.g., roulette events or coin flipping events. Formally, they must belong to a so-called ‘robust’ algebra of ‘incorporable’ objective events. Before defining these terms, I anticipate the axiom’s definitive statement (simpler axioms could also be used, as seen later in Sections 7 and 8):

Axiom 6 (Archimedean, global version 2): There is a robust algebra \mathcal{R} of incorporable objective events such that, for all contexts $\alpha \in \Gamma$, acts $f \succ_\alpha g$ in F_α , and outcomes $x \in X_\alpha$, one can partition \mathbf{S} into some $A_1, \dots, A_n \in \mathcal{R}$ such that, in some context $\beta \in \Gamma$ with state space $S_\beta = S_\alpha \vee \{A_1, \dots, A_n\}$ (ensuring that each A_i is representable by an $E_i \subseteq S_\beta$) and outcome space $X_\beta \supseteq X_\alpha$ (ensuring that F_β contains acts f' and g' equivalent to f resp. g), we have $f'_{S_\beta \setminus E_i} x_{E_i} \succ_\beta g'$ and $f' \succ_\beta g'_{S_\beta \setminus E_i} x_{E_i}$ for all E_i .

Axiom 6 of course allows that $\alpha = \beta$; then A_1, \dots, A_n are already representable in context α . The label ‘ \mathcal{R} ’ is meant to suggest ‘risky’ or ‘robust’. I now gradually build up the axiom’s terminology. I start with the familiar join operator:

Definition 12 *The join of partitions S and T is $S \vee T := \{s \cap t : s \in S, t \in T\} \setminus \{\emptyset\}$.¹⁷*

An objective event may or may not be representable in a context. Formally:

Definition 13_{exh} *In a context $\alpha \in \Gamma$, an objective event $A \subseteq \mathbf{S}$ is (**subjectively**) **representable** if it corresponds to some subjective event, which is then called its (**subjective**) **representation** and denoted $A_\alpha (= \{s \in S_\alpha : s \subseteq A\})$.*

An objective event $\{r, s, t\} \subseteq \mathbf{S}$ might be represented by $\{\{r, s\}, \{t\}\} \subseteq S_\alpha$ in a context α , and by $\{\{r, s, t\}\} \subseteq S_\beta$ in a context β , while being non-representable in a context γ in which the agent lacks appropriate state awareness.

An algebra¹⁸ \mathcal{R} on \mathbf{S} is *robust* if the ranking of \mathcal{R} -determined acts is stable:

Definition 14 *For an algebra \mathcal{R} on \mathbf{S} , an act f is **\mathcal{R} -determined** if the inverse image $f^{-1}(x)$ of any of its outcomes x represents an objective event in \mathcal{R} .*

¹⁷To be precise, S and T are partitions in the generalized sense of possibly containing \emptyset . Note that S and T could partition *different* sets, a case relevant later under non-exhaustive states.

¹⁸ \mathcal{R} is an algebra on \mathbf{S} if (a) $\mathbf{S} \in \mathcal{R}$, (b) $A \in \mathcal{R} \Rightarrow \overline{A} \in \mathcal{R}$, and (c) $A, B \in \mathcal{R} \Rightarrow A \cup B \in \mathcal{R}$.

Remark 7_{exh} An act f is \mathcal{R} -determined (given an algebra \mathcal{R} on \mathbf{S}) if and only if f^* is \mathcal{R} -measurable.¹⁹

Definition 15_{exh} An algebra \mathcal{R} on \mathbf{S} is **robust** if, for all contexts $\alpha, \beta \in \Gamma$, we have $f \succsim_\alpha g \Leftrightarrow f' \succsim_\beta g'$ whenever $f \in F_\alpha$ and $f' \in F_\beta$ are equivalent \mathcal{R} -determined acts, and $g \in F_\alpha$ and $g' \in F_\beta$ are also equivalent \mathcal{R} -determined acts.

Robustness is plausible if \mathcal{R} contains *risky* objective events, so that \mathcal{R} -determined acts are *risky acts*, because the agent presumably has fixed ‘preferences under risk’. The idea is that a *risky* objective event tends to get the same subjective probability regardless of the state space S_α in which it is represented: the event that a fair coin lands heads *always* has 1/2 probability, objectively and thus (where conceived) subjectively. This translates into a stable evaluation of risky acts, hence into robustness. I now introduce another natural notion:

Definition 16 A preference relation \succsim_β is **faithful to** another \succsim_α ($\alpha, \beta \in \Gamma$) if it preserves all comparisons made by \succsim_α : given any acts $f, g \in F_\alpha$, we have $f \succsim_\alpha g \Leftrightarrow f' \succsim_\beta g'$ for some (unique) acts $f', g' \in F_\beta$ equivalent to f resp. g .

If \succsim_β is faithful to \succsim_α , then any act in F_α is equivalent to one in F_β . So in context β the agent must conceive the same outcomes and at least as fine states:

Remark 8 If \succsim_β is faithful to \succsim_α , then (a) S_β is at least as fine as S_α (assuming $|X_\alpha| > 1$), and (b) $X_\beta \supseteq X_\alpha$ (hence $X_\beta = X_\alpha$ under exhaustive outcomes).

An objective event is *incorporable* if, whenever it is not representable, the agent can refine states to make it representable, without ‘preference perturbation’.

Definition 17 An objective event $A \subseteq \mathbf{S}$ is **incorporable** if it is always representable after (if needed) a preference-neutral state refinement: for every context $\alpha \in \Gamma$ there is a context $\beta \in \Gamma$ (possibly equal to α) such that S_β refines S_α to make A representable, i.e., $S_\beta = S_\alpha \vee \{A, \bar{A}\}$, and \succsim_β is faithful to \succsim_α .

The paradigmatic example of incorporability is, once again, *risky* objective events, as these are trivial in many respects. Refining states such that a coin toss becomes representable is an easy mechanical task (at least in principle), and the new preferences should be faithful to the old ones since the ranking of previously conceived (hence, coin-toss-independent) acts will hardly change.

Our axioms generalize Savage’s well-known axioms (stated in Appendix C.2):

Remark 9 In the single-context case $\Gamma = \{\gamma\}$, the variable Savage framework $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ is equivalent to an ordinary Savage framework $(X, S, \succsim) = (X_\gamma, S_\gamma, \succsim_\gamma)$, and our axioms reduce to Savage’s axioms, i.e.,

¹⁹ \mathcal{R} -measurability of f^* means that $(f^*)^{-1}(x) \in \mathcal{R}$ for all outcomes x of f^* , i.e., of f .

- (a) *Axiom 1 is equivalent to Savage's Axiom P1,*
- (b) *Axioms 2 and 2* are equivalent to Savage's Axiom P2,*
- (c) *Axiom 3 is equivalent to Savage's Axiom P3,*
- (d) *Axioms 4 and 4* are equivalent to Savage's Axiom P4,*
- (e) *Axiom 5 is equivalent to Savage's Axiom P5,*
- (f) *Axioms 6, 6* and 6** are equivalent to Savage's Axiom P6.²⁰*

4 Objective instability, subjective stability

Interestingly, whether an agent who obeys our axioms is stable or context-dependent in his preferences and beliefs depends on the chosen level of description.

4.1 Instability at the objective level

When modelling options as *objective* acts, choice reversals happen easily. Just imagine that in two contexts $\alpha, \beta \in \Gamma$ the agent chooses between the same objective acts $f, g \in \mathbf{F}$, which he subjectively represents as $f_\alpha, g_\alpha \in F_\alpha$ in context α , and as $f_\beta, g_\beta \in F_\beta$ in context β (see Definition 4). Then he will choose f in context α if $f_\alpha \succ_\alpha g_\alpha$, but g in context β if $g_\beta \succ_\beta f_\beta$. Such reversals are driven by changes in representation, i.e., description. All this is consistent with Axioms 1–6. One may view such reversals as *preference* reversals, by ‘lifting’ preferences to the objective level. I shall talk then of ‘effective’ preferences:

Definition 18 (*preference over objective acts*) *In a context $\alpha \in \Gamma$, an objective act $f \in \mathbf{F}$ is (effectively) weakly preferred to another one $g \in \mathbf{F}$ – written $f \succsim_\alpha g$ – if f and g are representable and the representations satisfy $f_\alpha \succsim_\alpha g_\alpha$.*

The (effective) preference between $f, g \in \mathbf{F}$ is reversible, as possibly $f \succsim_\alpha g$ but $g \succ_\beta f$. In a similar vein, (effective) *beliefs* are reversible. The agent may attach high probability to the event $\{\{s\}, \{t\}\}$ (where conceived), but low probability to the event $\{\{s, t\}\}$ (where conceived), although both events represent the same objective event $\{s, t\}$. Formally, we may lift the agent’s comparative beliefs to the objective level, talking then of ‘effective’ beliefs:

Definition 19 (*comparative belief about objective events*) *In a context $\alpha \in \Gamma$, an objective event $A \subseteq \mathbf{S}$ is (effectively) at least as probable as another one $B \subseteq \mathbf{S}$ – written $A \succsim_\alpha B$ – if A and B are representable and the representations satisfy $A_\alpha \succsim_\alpha B_\alpha$.*

Nothing prevents a belief $A \succsim_\alpha B$ ($A, B \subseteq \mathbf{S}$) to reverse into $B \succ_\beta A$. However:

²⁰ Axioms 6* and 6** imply Axioms 6 by letting \mathcal{R} contain *all* representable objective events.

Proposition 1 (*stability of comparative belief on robust algebras*) Under Axioms 2, 4 and 5, objective events from a robust algebra \mathcal{R} on \mathbf{S} are ranked the same way wherever representable: $A \succsim_{\alpha} B \Leftrightarrow A \succsim_{\beta} B$ for all objective events $A, B \in \mathcal{R}$ representable in both contexts α and β (where $\alpha, \beta \in \Gamma$).

4.2 Stability at the subjective level

Despite ‘objective instability’, our axioms imply stable preferences over *subjective* acts (and outcomes) and stable comparative beliefs about *subjective* events.

Proposition 2 (*preference stability*) Under Axiom 2, acts are ranked the same way wherever conceived: $f \succsim_{\alpha} g \Leftrightarrow f \succsim_{\beta} g$ for all acts conceived in both contexts $f, g \in F_{\alpha} \cap F_{\beta}$ (where $\alpha, \beta \in \Gamma$).

So, under Axiom 2 the context affects only which acts are conceived, not how acts are ranked *when conceived*. Saying ‘only’ is perhaps an understatement, as Proposition 2 has a bite only for those pairs of contexts $\alpha, \beta \in \Gamma$ for which $F_{\alpha} \cap F_{\beta} \neq \emptyset$, i.e., for which $S_{\alpha} = S_{\beta}$ and $X_{\alpha} \cap X_{\beta} \neq \emptyset$. If awareness varies so drastically that no distinct contexts share any acts, then Proposition 2 is vacuous.

Proposition 3 (*outcome-preference stability*) Under Axiom 6, outcomes are ranked the same way wherever conceived: $x \succsim_{\alpha} y \Leftrightarrow x \succsim_{\beta} y$ for all outcomes conceived in both contexts $x, y \in X_{\alpha} \cap X_{\beta}$ (where $\alpha, \beta \in \Gamma$).

One might at first take stability over outcomes to be a special case of stability over acts, by identifying outcomes with constant acts. In fact, both stability properties are independent, as the same outcome $x \in X_{\alpha} \cap X_{\beta}$ is identified with *distinct* constant acts $x_{S_{\alpha}} \in F_{\alpha}$ and $x_{S_{\beta}} \in F_{\beta}$ if $S_{\alpha} \neq S_{\beta}$.

Proposition 4 (*comparative-belief stability*) Under Axioms 2, 4, 5 and 6, events are ranked the same way wherever conceived: $A \succsim_{\alpha} B \Leftrightarrow A \succsim_{\beta} B$ for all events conceived in both contexts $A, B \subseteq S_{\alpha} \cap S_{\beta}$ (where $\alpha, \beta \in \Gamma$).

5 The representation theorem

I now state the theorem; it will be restated in Section 8 using a simpler sixth axiom and an exogenous notion of risk. I start with terminology:

Definition 20 For the variable Savage framework $(X_{\alpha}, S_{\alpha}, \succsim_{\alpha})_{\alpha \in \Gamma}$, a (**variable**) **expected-utility representation** is a system $(U_{\alpha}, P_{\alpha})_{\alpha \in \Gamma}$ of non-constant ‘utility’ functions $U_{\alpha} : X_{\alpha} \rightarrow \mathbb{R}$ and probability measures²¹ $P_{\alpha} : 2^{S_{\alpha}} \rightarrow [0, 1]$ such that

$$f \succsim_{\alpha} g \Leftrightarrow \mathbb{E}_{P_{\alpha}}(U_{\alpha} \circ f) \geq \mathbb{E}_{P_{\alpha}}(U_{\alpha} \circ g) \text{ for all contexts } \alpha \in \Gamma \text{ and acts } f, g \in F_{\alpha}.$$

²¹The term ‘probability measure’ is used in its *finitely* additive sense.

Definition 21 A probability measure on an algebra \mathcal{R} is **fine** if for all $\epsilon > 0$ there are mutually exclusive and exhaustive $A_1, \dots, A_n \in \mathcal{R}$ of probabilities at most ϵ .²²

I call a function ρ on a set \mathcal{R} of objective events *uncontroversial* among probability measures P_α on 2^{S_α} ($\alpha \in \Gamma$) if, roughly speaking, each P_α assigns probability $\rho(A)$ to the event representing $A \in \mathcal{R}$. The precise definition is more general: it allows an $A \in \mathcal{R}$ to be *not (yet) representable* in a context α , in which case the probability $\rho(A)$ is derived not from P_α itself, but from a version of P_α defined on a refined state space that makes A representable. Formally:

Definition 22 Given a context $\alpha \in \Gamma$, a function P on 2^{S_α} induces a function P^* on the set of representable objective events $A \in 2^{\mathbf{S}}$ via $P^*(A) := P(A_\alpha)$.

Definition 23 A function ρ on a set \mathcal{R} of objective events is **uncontroversial** among functions P_α on 2^{S_α} ($\alpha \in \Gamma$) if each induced function P_α^* matches ρ ‘modulo extension’: for any $A \in \mathcal{R}$, each P_α^* has an extension P_β^* (for some $\beta \in \Gamma$) such that $S_\beta = S_\alpha \vee \{A, \bar{A}\}$ (so that $P_\beta^*(A)$ is defined) and $P_\beta^*(A) = \rho(A)$.

Remark 10 If ρ (defined on \mathcal{R}) is uncontroversial among functions P_α ($\alpha \in \Gamma$), then $P_\alpha^*(A) = \rho(A)$ whenever $A \in \mathcal{R}$ is representable in context $\alpha \in \Gamma$.

Our axioms characterize expected-utility preferences with certain revision rules:

Theorem 1 The variable Savage framework $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ satisfies Axioms 1–6 if and only if it has an expected-utility representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ satisfying three revision rules: (R1) any U_α is an increasing affine transformation of any U_β on $X_\alpha \cap X_\beta$, (R2) any P_α is proportional to any P_β on $2^{S_\alpha \cap S_\beta}$, and (R3) some fine (‘objective’) probability measure on some algebra on \mathbf{S} is uncontroversial among the measures P_α . Each P_α is unique and each U_α is unique up to increasing affine transformation.²³

Rules R1–R3 describe how utilities and probabilities are revised as the agent’s outcome/state concepts change. By R1 and R2 utilities and probabilities are affinely resp. proportionally rescaled where concepts are stable. So if the agent, say, splits an outcome $x \in X_\alpha$ into y and z , resulting in a context β with $X_\beta = (X_\alpha \setminus \{x\}) \cup \{y, z\}$ and $S_\beta = S_\alpha$, then $P_\beta = P_\alpha$ by R2, and utilities are essentially unchanged on $X_\alpha \setminus \{x\}$ by R1. By R3 certain (‘objective’) probabilities are robust.

Remark 11 In Theorem 1’s representation, probabilities are independent of outcome awareness, and utilities are independent of state awareness: if from context α to context β only the outcome space changes then $P_\alpha = P_\beta$, and if only the state space changes then $U_\alpha = U_\beta$ for suitably normalised utility functions. So:

²²Fineness implies Savage’s *atomlessness*, and is equivalent to atomlessness if \mathcal{R} is a σ -algebra.

²³Formally, if $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in the theorem’s sense, then $(U'_\alpha, P'_\alpha)_{\alpha \in \Gamma}$ is also one if and only if, for all $\alpha \in \Gamma$, $P'_\alpha = P_\alpha$ and $U'_\alpha = a_\alpha U_\alpha + b_\alpha$ for some $a_\alpha > 0$ and $b_\alpha \in \mathbb{R}$.

- if state awareness (i.e., S_α) is the same in all contexts α , $P_\alpha = P$ is fixed,
- if outcome awareness (i.e., X_α) is the same in all contexts α , $U_\alpha = U$ is fixed given suitably normalised utility functions.

Remark 12 In the single-context case $\Gamma = \{\gamma\}$, Theorem 1 reduces to Savage’s Theorem (for the ordinary Savage framework $(X_\gamma, S_\gamma, \succsim_\gamma)$), as Axioms 1–6 reduce to Savage’s Axioms P1–P6 (by Remark 9), rules R1 and R2 hold trivially, and R3 reduces to atomlessness of P_γ .²⁴

Remark 13 In contrast to Savage’s Theorem, Theorem 1’s representation allows that all state spaces S_α are finite; but it forces the objective state space \mathbf{S} to be infinite (as by R3 there is an infinite algebra on \mathbf{S}). \mathbf{X} can be finite or infinite.

Remark 14 R3 has an equivalent formulation: ‘the objective probability function induced by the P_α is fine’. This draws on a well-defined, purely endogenous and preference-based notion of ‘objective probability’.²⁵

Rule R2 implies stable probability *ratios*, an interesting analogy to Ahn-Ergin’s (2010) ‘partition-dependent probabilities’ and Karni-Viero’s (2013) ‘reverse Bayesianism’. But the functions P_α need not admit an Ahn-Ergin-type representation.²⁶

It is tempting to further increase uniqueness of the representation by requiring any U_α and U_β to *coincide* on the domain overlap; one could then replace the family of U_α functions by a *single* function U on $\cup_{\alpha \in \Gamma} X_\alpha$. This may not work:

Remark 15 In Theorem 1 it may be impossible to ‘simultaneously scale’ the utility functions such that any U_α and U_β coincide on the domain overlap $X_\alpha \cap X_\beta$.

Indeed, after scaling U_β to match U_α on $X_\alpha \cap X_\beta$, and scaling U_γ to match U_β on $X_\beta \cap X_\gamma$, U_γ might fail to match U_α on $X_\alpha \cap X_\gamma$. This shows the genuine need for utility revision. I now give two examples of Theorem 1’s representation.

A trivial example (with only *risky* contingencies and full outcome awareness): A fair coin is tossed infinitely often. Let $\mathbf{X} = \{x, y\}$ and $\mathbf{S} = \{0, 1\}^\mathbb{N}$. In an objective

²⁴Indeed, R3 reduces to fineness of P_γ , and so to atomlessness of P_γ by footnote 22.

²⁵The (*endogenous* or *revealed*) *objective probability function* is the uncontroversial function ρ with *largest* domain. It is fully determined by preferences as all P_α are unique. It is a probability measure, albeit in the generalized sense that its domain \mathcal{R} need not be an algebra (it need not be closed under union, but is closed under complement and contains \mathbf{S}). That is, $\rho(\mathbf{S}) = 1$ and ρ is additive, i.e., $\rho(A \cup B) = \rho(A) + \rho(B)$ if $A, B, A \cup B \in \mathcal{R}$ and $A \cap B = \emptyset$. ρ is *fine* if it defines a fine (ordinary) probability measure on some (sub)algebra $\mathcal{R}' \subseteq \mathcal{R}$.

²⁶That is, even if all S_α are finite, there may not exist any (possibly non-additive) function μ on $\cup_{\alpha \in \Gamma} S_\alpha$ which simultaneously induces each P_α in the sense that P_α and μ are proportional as functions on S_α . Indeed, there may exist contexts $\alpha, \beta, \gamma \in \Gamma$ and states $r \in S_\alpha \cap S_\beta, s \in S_\beta \cap S_\gamma, t \in S_\gamma \cap S_\alpha$ such that $P_\alpha(r) = P_\alpha(t), P_\beta(r) = P_\beta(s), P_\gamma(s) \neq P_\gamma(t)$; here, an inducing μ would have to satisfy $\mu(r) = \mu(t), \mu(r) = \mu(s), \mu(s) \neq \mu(t)$, a contradiction.

state $(s_i)_{i \in \mathbb{N}} \in \mathbf{S}$, an s_i is 1 resp. 0 depending on whether the i -th toss resulted in heads resp. tails. An objective event $A \subseteq \mathbf{S}$ is *finitely complex* if it concerns only finitely many tosses, i.e., $A = \{(s_i)_{i \in \mathbb{N}} \in \mathbf{S} : (s_i)_{i \in I} \in B\}$ for some finite subset $I \subseteq \mathbb{N}$ and some $B \subseteq \{0, 1\}^I$. An example is the objective event ‘first toss heads, fourth tails’ (here $I = \{1, 4\}$). Identifying contexts with state spaces, let Γ be the set of (finite non-singleton) partitions α of \mathbf{S} into finitely complex objective events, and let $S_\alpha := \alpha$ and $X_\alpha := \{\{x\}, \{y\}\}$. An example is $\alpha = S_\alpha = \{\text{‘first toss heads’}, \text{‘first toss tails, second heads’}, \text{‘first toss tails, second tails’}\}$. In each context $\alpha \in \Gamma$, let the agent hold expected-utility preferences given by a context-invariant (non-constant) utility function $U = U_\alpha$ on X_α and the probability measure P_α on 2^{S_α} which assigns probability $\frac{|B|}{2^{|I|}}$ to state $\{(s_i)_{i \in \mathbb{N}} \in \mathbf{S} : (s_i)_{i \in I} \in B\} \in S_\alpha$.²⁷ So P_α mirrors that the agent knows that the coin is fair and the tosses are independent. Rules R1 and R2 hold. Also R3 holds, because the objective tossing probabilities are fine and uncontroversial.²⁸

A refined example: I now enrich the previous example by including non-risky contingencies, which may get different probabilities depending on the subjective representation. Let \mathbf{S}' be a non-empty set of ‘non-risky objective states’, representing contingencies without objective probability such as the weather or the music at tonight’s concert. I redefine the objective state space as $\{0, 1\}^{\mathbb{N}} \times \mathbf{S}'$, whose members (s, s') have two parts: a ‘risky objective state’ $s \in \{0, 1\}^{\mathbb{N}}$ with the same coin-toss interpretation as before, and a ‘non-risky objective state’ $s' \in \mathbf{S}'$. Let Γ' be some non-empty set of finite partitions of \mathbf{S}' ; they represent the agent’s possible awareness levels relative to non-risky contingencies. If one also wishes to model *non-exhaustive* awareness (to which we return in Section 6), one should more generally let Γ' contain partitions of (non-empty) *subsets* of \mathbf{S}' . I redefine the set of contexts as $\Gamma \times \Gamma'$, where Γ is the old set of contexts. In context $\gamma = (\alpha, \alpha') \in \Gamma \times \Gamma'$, the outcome space is still $X_\gamma = \{\{x\}, \{y\}\}$, and the state space is $S_\gamma := \{A \times A' : A \in \alpha, A' \in \alpha'\}$. Each state $A \times A' \in S_\gamma$ is thus composed of ‘risky state’ A and a ‘non-risky state’ A' . In each context $\gamma = (\alpha, \alpha')$, let the agent hold expected-utility preferences given by a context-invariant (non-constant) utility function $U = U_\gamma$ on X_γ , and the probability function P_γ which to any state $A \times A' \in S_\gamma$ assigns the probability $P_\gamma(A \times A') := P_\alpha(A)P_{\alpha'}(A')$, where P_α is the earlier-defined probability measure for the ‘risky state space’ α , and $P_{\alpha'}$ is a probability measure for the ‘non-risky state space’ α' . This reflects the plausible idea that coin tosses are independent of non-risky contingencies. I assume the functions

²⁷The value $\frac{|B|}{2^{|I|}}$ does not depend on the pair (I, B) used to represent the state in S_α (the most natural representation takes the minimal I).

²⁸Formally, the \mathbb{N} -fold product $\bigotimes_{i=1}^{\infty} \text{Bernoulli}(\frac{1}{2})$ of the uniform Bernoulli measure is fine and uncontroversial. I define it on an algebra (not ‘ σ ’-algebra) on $\{0, 1\}^{\mathbb{N}}$: the \mathbb{N} -fold product of the power-set algebra on $\{0, 1\}$, which consists precisely of the finitely complex objective events.

$P_{\alpha'} (\alpha' \in \Gamma')$ are related to each other: let there be an arbitrary function μ assigning to any ever conceived non-risky state $s' \in \cup_{\alpha' \in \Gamma'} S_{\alpha'}$ a ‘plausibility’ $\mu(s') > 0$, and let μ induce each $P_{\alpha'} (\alpha' \in \Gamma')$ in the sense that $P_{\alpha'}$ and μ are proportional as functions on $S_{\alpha'}$ (so $P_{\alpha'}$ arises from normalising μ within $S_{\alpha'}$). Rules R1 and R2 hold, as one may verify. Rule R3 holds since, as before, the objective tossing probabilities are fine and uncontroversial.²⁹

6 The general case

I now lift the temporary restriction to exhaustive states. Recall that the above ‘theorem’ and ‘propositions’ and most ‘definitions’ and ‘remarks’ continue to apply *as stated*. The three exceptions, namely Definitions 13_{exh} and 15_{exh} and Remark 7_{exh}, will now be re-stated in their general form, using the same numbering but without index ‘exh’. The general statements are equivalent to their earlier counterparts in case of exhaustive states. In light of the generalized statements, readers can afterwards reconsider Sections 3–5 without restriction to exhaustive states. This will pose no problems, but two details should be kept in mind. For one, two partitions (e.g., state spaces) of which one refines the other must be partitions of the *same* set (see Definition 11). Further, the (unchanged) Definitions 17 and 23 and Axiom 6, when applied with non-exhaustive states, require forming the join of partitions of *possibly distinct* sets (namely \mathbf{S}_α and \mathbf{S}). This join then partitions the intersection of the two sets (here \mathbf{S}_α), by Definition 12. I now state the three generalizations.

First, I generalize the definition of representations of objective events:

Definition 13 *In a context $\alpha \in \Gamma$, an objective event $A \subseteq \mathbf{S}$ is (**subjectively**) **representable** if its encompassed part $A \cap \mathbf{S}_\alpha$ corresponds to a subjective event, called then A ’s (**subjective**) **representation**, denoted $A_\alpha (= \{s \in S_\alpha : s \subseteq A\})$.*

Second, the notion of an act f being *determined* by an algebra \mathcal{R} on \mathbf{S} , while defined as before, has a generalized ‘measurability characterization’:

Remark 7 *A subjective act $f \in F_\alpha (\alpha \in \Gamma)$ is \mathcal{R} -determined (given an algebra \mathcal{R} on \mathbf{S}) if and only if f^* is \mathcal{R}' -measurable where $\mathcal{R}' = \{A \cap \mathbf{S}_\alpha : A \in \mathcal{R}\}$ is the trace of \mathcal{R} in \mathbf{S}_α .*

Third, I generalize the definition of *robustness* of an algebra \mathcal{R} , through replacing ‘equivalent \mathcal{R} -determined acts’ by ‘corresponding \mathcal{R} -determined acts’:

Definition 24 *Two acts $f \in F_\alpha$ and $f' \in F_\beta$ (where $\alpha, \beta \in \Gamma$) are **corresponding***

²⁹Formally, letting ρ be the fine and uncontroversial measure of the ‘trivial’ example and \mathcal{R} its underlying algebra on $\{0, 1\}^\mathbb{N}$ (see footnote 28), we obtain a fine uncontroversial measure ρ' for the refined example by defining ρ' on the algebra $\{A \times \mathbf{S}' : A \in \mathcal{R}\}$ by $\rho'(A \times \mathbf{S}') := \rho(A)$.

\mathcal{R} -determined acts (for an algebra \mathcal{R} on \mathbf{S}) if both are given by an identical \mathcal{R} -measurable function, i.e., there is an \mathcal{R} -measurable function \mathbf{f} on \mathbf{S} such that $\mathbf{f}(\mathbf{s}) = f(s)$ whenever $\mathbf{s} \in s \in S_\alpha$ and $\mathbf{f}(\mathbf{s}) = f'(s)$ whenever $\mathbf{s} \in s \in S_\beta$ (i.e., such that, $\mathbf{f}_{S_\alpha} = f^*$ and $\mathbf{f}_{S_\beta} = f'^*$).

Definition 15 An algebra \mathcal{R} on \mathbf{S} is **robust** if, for all contexts $\alpha, \beta \in \Gamma$, we have $f \succsim_\alpha g \Leftrightarrow f' \succsim_\beta g'$ whenever $f \in F_\alpha$ and $f' \in F_\beta$ are corresponding \mathcal{R} -determined acts, and $g \in F_\alpha$ and $g' \in F_\beta$ are also corresponding \mathcal{R} -determined acts.

Definition 15 indeed generalizes Definition 15_{exh}, for a simple reason:

Remark 16 In case of exhaustive states, then acts $f \in F_\alpha$ and $f' \in F_\beta$ ($\alpha, \beta \in \Gamma$) are corresponding \mathcal{R} -determined acts if and only if they are equivalent (i.e., $f^* = f'^*$) and \mathcal{R} -determined.

The label ‘corresponding \mathcal{R} -determined acts’ is explained by a simple fact:

Remark 17 Each of two corresponding \mathcal{R} -determined acts is \mathcal{R} -determined.

7 The special case of fine states

I now apply our theorem to the case of fine states. Here all $s \in S_\alpha$ are singleton, and just one kind of state awareness changes: the level of state exhaustiveness.

Remark 18 In case of fine states, all objective events are representable in each context (hence, are trivially incorporable).

As a result, the fine-state case allows us to work with a simpler sixth axiom:

Axiom $\tilde{6}$ (Archimedean, fine-state version): There is a robust algebra \mathcal{R} on \mathbf{S} such that, for all contexts $\alpha \in \Gamma$, acts $f \succ_\alpha g$ in F_α , and outcomes $x \in X_\alpha$, one can partition S_α into events $E_1, \dots, E_n \subseteq S_\alpha$ representing objective events from \mathcal{R} such that $f_{S_\alpha \setminus E_i} x_{E_i} \succ_\alpha g$ and $f \succ_\alpha g_{S_\alpha \setminus E_i} x_{E_i}$ for all E_i .

We may also work with a more basic notion than ‘uncontroversial measures’:

Definition 25 The **commonality** of functions P_α on 2^{S_α} ($\alpha \in \Gamma$) is the meet (greatest common subfunction) of all P_α^* ($\alpha \in \Gamma$).

Remark 19 The commonality of probability measures P_α ($\alpha \in \Gamma$) is itself a probability measure, namely the restriction of each P_α^* to the algebra $\{A \subseteq \mathbf{S} : \text{in all contexts } \alpha \in \Gamma \text{ the probability } P_\alpha^*(A) \text{ is defined}^{30} \text{ and identical}\}$.

³⁰Definedness is equivalent to representability of A , and comes for free under fine states.

Interpretively, the commonality of probability measures is their ‘objective overlap’ and captures objective probabilities.³¹ In general, its domain can be as small as $\{\emptyset, \mathbf{S}\}$ or as large as $2^{\mathbf{S}}$. Theorem 1’s fine-state corollary follows via two lemmas:

Lemma 1 *Under fine states, Axioms 6 and $\tilde{6}$ are equivalent given Axiom 2.*

Lemma 2 *Under fine states, R3 holds if and only if the P_α have fine commonality.*

Corollary 1 *Under fine states, Axioms 1–5 and $\tilde{6}$ hold if and only if there is an expected-utility representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ satisfying R1, R2, and a third revision rule: the functions P_α have fine commonality.³²*

8 Exogenizing risk

I now restate Theorem 1 using an exogenous notion of ‘risky objective events’ (but leaving the ‘objective’ probabilities of these events endogenous). I introduce an exogenous algebra \mathcal{R} (on \mathbf{S}) of ‘risky’ objective events, and replace Axiom 6 by three axioms with \mathcal{R} as parameter:

Axiom $6_{\mathcal{R}}$ (Archimedean, global version 3): This axiom states like Axiom 6, but without the initial quantification ‘There is a robust algebra \mathcal{R} of incorporable objective events such that’.

Axiom $7_{\mathcal{R}}$ (robust risk preference): The algebra \mathcal{R} is robust.

Axiom $8_{\mathcal{R}}$ (risk incorporability): All objective events in \mathcal{R} are incorporable.

Theorem 2 *Given an exogenous (risky) algebra \mathcal{R} on \mathbf{S} , the variable Savage framework $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ satisfies Axioms 1–5 and $6_{\mathcal{R}}-8_{\mathcal{R}}$ if and only if it has an expected-utility representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ satisfying R1, R2, and a third revision rule: some fine (‘objective’) probability measure on \mathcal{R} is uncontroversial among the measures P_α . Each P_α is unique and each U_α is unique up to increasing affine transformation.*

Remarks 11, 13 and 15 apply analogously to Theorem 2. To obtain Theorem 2’s fine-state corollary, I simplify Axiom $6_{\mathcal{R}}$, drop Axiom $8_{\mathcal{R}}$ (which comes for free) and simplify rule R3, drawing on two lemmas:

Axiom $\tilde{6}_{\mathcal{R}}$: This axiom states like Axiom $\tilde{6}$, but without the initial quantification ‘There is a robust algebra \mathcal{R} such that’.

³¹Under fine states it is just the *endogenous objective probability function* of footnote 25.

³²Fine states are essentially objective states. So, had this paper focused exclusively on fine states, we could have introduced each S_α as a primitive set (not a partition), redefined the ‘objective state space’ as $\cup_{\alpha \in \Gamma} S_\alpha$, and redefined accordingly all concepts that refer to objective states (such as ‘robust algebras’ and the ‘commonality’ of functions).

Lemma 3 *Under fine states and an exogenous (risky) algebra \mathcal{R} on \mathbf{S} , Axioms $6_{\mathcal{R}}$ and $\tilde{6}_{\mathcal{R}}$ are equivalent given Axiom 2.*

Lemma 4 *Under fine states and an exogenous (risky) algebra \mathcal{R} on \mathbf{S} , Theorem 2's third revision rule holds if and only if the commonality of the measures P_{α} extends a fine measure on \mathcal{R} .*

Corollary 2 *Under fine states and an exogenous (risky) algebra \mathcal{R} on \mathbf{S} , Axioms 1–5, $\tilde{6}_{\mathcal{R}}$ and $7_{\mathcal{R}}$ hold if and only if there is an expected-utility representation $(U_{\alpha}, P_{\alpha})_{\alpha \in \Gamma}$ satisfying R1, R2, and a third revision rule: the commonality of the functions P_{α} includes (i.e., extends) a fine measure on \mathcal{R} .*

9 Concluding remarks

I have presented a unified theorem for preferences under uncertainty and changing awareness. Preferences are governed by expected utility with three rules for revising utilities and probabilities. The theorem has many special cases, including (i) fixed awareness, where we recover the classic Savage theorem, (ii) fixed outcome awareness, where utilities are stable, (iii) fixed state awareness, where probabilities are stable, (iv) exhaustive state awareness, where some definitions simplify, and (v) fine state awareness, where Axiom 6 simplifies and the third revision rule is expressible in terms of the *commonality* ('objective overlap') of the probability functions. Just as Savage's axioms have been weakened over time, giving rise to 'non-expected-utility' theories, it would be interesting to relax the current axioms and explore alternative representations with other revision rules.

In our analysis, the agent 'looks' stable or unstable (in his preferences and beliefs) depending on whether the subjective or objective level of description is chosen. This suggests that instability and context-dependence are phenomena driven by a changing subjective perception of the objective world.

A When are probabilities objectively stable?

A stronger belief stability condition than R2 requires any P_{α} to *coincide* with any P_{β} on the domain overlap. An even stronger condition (natural from an orthodox full-rationality perspective) requires stable probabilities of *objective* events, irrespective of their subjective representation. Formally:

(R2⁺) **Objective belief stability:** *For all $\alpha, \beta \in \Gamma$, $A \subseteq S_{\alpha}$ and $B \subseteq S_{\beta}$, if $A^* = B^*$, then $P_{\alpha}(A) = P_{\beta}(B)$. (So the functions P_{α} are given by fixed function of objective events.)*

This for instance forces the event $\{\{s, s'\}, \{s''\}\}$ (where conceived) to have the same probability as $\{\{s\}, \{s', s''\}\}$ (where conceived), as both events represent

$\{s, s', s''\}$. Rule $R2^+$ holds in Section 4's 'trivial example', but not in the 'refined example'. How do we need to strengthen our axioms to enforce $R2^+$? Under exhaustive states, the following additional axiom fills the gap:

Axiom ROB (robustness): For all contexts $\alpha, \beta \in \Gamma$, we have $f \succsim_\alpha g \Leftrightarrow f' \succsim_\beta g'$ whenever acts $f, g \in F_\alpha$ are objectively equivalent to acts $f', g' \in F_\beta$, respectively.

Corollary 3 *Under exhaustive states, preferences satisfy Axioms 1–6 and ROB if and only if they have an expected-utility representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ satisfying R1, $R2^+$ and R3.*

Proof. Assume exhaustive states. First, let preferences admit Corollary 3's representation. Ax. 1–6 hold by Theorem 1. Regarding Ax. ROB, consider $\alpha, \beta \in \Gamma$; w.l.o.g. let U_α and U_β coincide on $X_\alpha \cap X_\beta$. It suffices to consider $f \in F_\alpha$ and $f' \in F_\beta$ with $f^* = f'^*$ and show that $\mathbb{E}_{P_\alpha}(U_\alpha \circ f) = \mathbb{E}_{P_\beta}(U_\beta \circ f')$, or equivalently, that $\mathbb{E}_{P_\alpha^*}(U_\alpha \circ f^*) = \mathbb{E}_{P_\beta^*}(U_\beta \circ f'^*)$. This holds because $U_\alpha \circ f^* = U_\beta \circ f'^*$ and because $P_\alpha^*(A) = P_\beta^*(A)$ for all A in the domains of both P_α^* and P_β^* .

Conversely, let Ax. 1–6 and ROB hold. By Theorem 1, there is a representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ with the properties specified there. Pick a fine and uncontroversial probability measure ρ . To show $R2^+$, consider $\alpha, \beta \in \Gamma$ and an $A \subseteq \mathbf{S}$ from the domains of P_α^* and P_β^* . I show that $P_\alpha^*(A) = P_\beta^*(A)$. For a contradiction, assume $P_\alpha^*(A) \neq P_\beta^*(A)$, say $P_\alpha^*(A) < P_\beta^*(A)$. As ρ is fine, its range is dense in $[0, 1]$. So we can pick a B in ρ 's domain such that $P_\alpha^*(A) < \rho(B) < P_\beta^*(A)$. As ρ is uncontroversial, P_α^* extends to some $P_{\alpha'}^*$ with state space $S_{\alpha'} = S_\alpha \vee \{B, \overline{B}\}$, where $P_{\alpha'}^*(B) = \rho(B)$; and similarly, P_β^* extends to some $P_{\beta'}^*$ with state space $S_{\beta'} = S_\beta \vee \{B, \overline{B}\}$, where $P_{\beta'}^*(B) = \rho(B)$. The inequalities $P_\alpha^*(A) < \rho(B)$ and $\rho(B) < P_\beta^*(A)$ now reduce to $P_{\alpha'}^*(A) < P_{\alpha'}^*(B)$ and $P_{\beta'}^*(B) < P_{\beta'}^*(A)$. This contradicts belief-stability on robust algebras (Prop. 1), since A and B belong to $2^{\mathbf{S}}$, a robust algebra by Ax. ROB. ■

Theorem 2 has an analogous corollary. Axiom ROB essentially requires preferences to be independent of the state concept. For instance, whether the agent prefers getting 100 Dollars in the objective event $A = \{s, s', s''\}$ (and nothing otherwise) to getting 50 Dollars for sure should not depend on whether he represents A as $\{\{s, s'\}, \{s''\}\}$, $\{\{s\}, \{s', s''\}\}$, $\{\{s, s', s''\}\}$ or $\{\{s\}, \{s'\}, \{s''\}\}$. However, if we take the idea of limited awareness seriously, there is little reason to believe in Axiom ROB or objective belief stability. An agent who fails to conceive objective states will not know whether two acts from different contexts are objectively equivalent. This undermines Axiom ROB's plausibility.

B Proof of the stability propositions

This and the following appendices contain proofs, starting with the stability propositions (App. B), followed by Thm. 1 under exhaustive states (App. C), Thm. 1 in general (App. D), Thm. 2 (App. E), and finally Lem. 1–4 and some technical lemmas stated in due course whose proofs are relegated to the end to avoid distraction (App. F). Proofs use the following notation:

- Recall the notation ‘ x_α ’, ‘ s_α ’, ‘ E^* ’ and ‘ f^* ’ (see Sect. 2.3), as well as ‘ A_α ’ (Def. 13_{exh} resp. 13) and P^* (Def. 22).
- For any set of events T , define the set of objective events $T^* := \{A^* : A \in T\}$.
- For any set F of acts, define the set of functions $F^* := \{f^* : f \in F\}$.
- For any $\alpha \in \Gamma$ and $f \in F_\alpha^*$, let $f_\alpha \in F_\alpha$ be the act given by $(f_\alpha)^* = f$. (‘ f_α ’ was also used for the representation of an objective act $f \in \mathbf{F}$; see Def. 4.)

Proof of Prop. 2. Just take $E = S_\alpha = S_{\alpha'}$, $f = f'$ and $g = g'$ in Ax. 2. ■

Proof of Prop. 3. Ax. 6 implies existence of a robust algebra \mathcal{R} on \mathbf{S} . The claim holds as \mathcal{R} is robust and as any two constant acts with same outcome (on possibly distinct state spaces) are corresponding \mathcal{R} -determined acts. ■

Proof of Prop. 4. Assume Ax. 2, 4, 5 and 6, and let $\alpha, \beta \in \Gamma$ and $A, B \subseteq S_\alpha \cap S_\beta$. I suppose $A \succsim_\alpha B$ and show that $A \succsim_\beta B$ (the converse is analogous). Using Ax. 5, pick outcomes $x \succ_\alpha y$ in X_α and $x' \succ_\beta y'$ in X_β . Using independence between outcome and state awareness, pick a $\gamma \in \Gamma$ with $S_\gamma = S_\beta$ and $X_\gamma = X_\alpha$. As $x \succ_\alpha y$ we have $x \succ_\gamma y$ by Prop. 3. As $A \succsim_\alpha B$ and $x \succ_\alpha y$, by Ax. 4 $x_A y_{S_\alpha \setminus A} \succsim_\alpha x_B y_{S_\alpha \setminus B}$. So $x_A y_{S_\gamma \setminus A} \succsim_\gamma x_B y_{S_\gamma \setminus B}$ by Ax. 2 applied to the event $A \cup B \subseteq S_\alpha \cap S_\gamma$. Hence $x'_A y'_{S_\beta \setminus A} \succsim_\beta x'_B y'_{S_\beta \setminus B}$ by Ax. 4. So $A \succsim_\beta B$. ■

Proof of Prop. 1. Assume Ax. 2, 4 and 5. Consider a robust algebra \mathcal{R} on \mathbf{S} , $\alpha, \beta \in \Gamma$, $A, B \subseteq S_\alpha$, and $\tilde{A}, \tilde{B} \subseteq S_\beta$, such that A and \tilde{A} represent an identical objective event $A' \in \mathcal{R}$, and B and \tilde{B} also represent an identical $B' \in \mathcal{R}$. I show that $A \succsim_\alpha B \Rightarrow \tilde{A} \succsim_\beta \tilde{B}$ (the converse direction ‘ \Leftarrow ’ holds analogously). Let $A \succsim_\alpha B$. Using Ax. 5, pick outcomes $x \succ_\alpha y$ in X_α and $x' \succ_\beta y'$ in X_β . Using independence between outcome and state awareness, pick a $\gamma \in \Gamma$ with $S_\gamma = S_\beta$ and $X_\gamma = X_\alpha$. As $x \succ_\alpha y$ we have $x \succ_\gamma y$, by Prop. 3 (more exactly, a version of Prop. 3 based not on Ax. 6, but only on the existence of a robust algebra). Also, as $A \succsim_\alpha B$ and $x \succ_\alpha y$, by Ax. 4 $x_A y_{S_\alpha \setminus A} \succsim_\alpha x_B y_{S_\alpha \setminus B}$. So $x_{\tilde{A}} y_{S_\gamma \setminus \tilde{A}} \succsim_\gamma x_{\tilde{B}} y_{S_\gamma \setminus \tilde{B}}$, because \mathcal{R} is robust, $x_A y_{S_\alpha \setminus A}$ and $x_{\tilde{A}} y_{S_\gamma \setminus \tilde{A}}$ are corresponding \mathcal{R} -determined acts (as both stem from the \mathcal{R} -measurable function $x_{A'} y_{\mathbf{S} \setminus A'}$), and $x_B y_{S_\alpha \setminus B}$ and $x_{\tilde{B}} y_{S_\gamma \setminus \tilde{B}}$ are also corresponding \mathcal{R} -determined acts (as both stem from the \mathcal{R} -measurable function $x_{B'} y_{\mathbf{S} \setminus B'}$). As $x_{\tilde{A}} y_{S_\gamma \setminus \tilde{A}} \succsim_\gamma x_{\tilde{B}} y_{S_\gamma \setminus \tilde{B}}$ and $x \succ_\gamma y$, by Ax. 4 $\tilde{A} \succsim_\gamma \tilde{B}$. So $\tilde{A} \succsim_\beta \tilde{B}$, by Prop. 4 (more precisely, a version of Prop. 4, like before). ■

C Proof of Theorem 1 under exhaustive states

Proof strategy: This appendix assumes exhaustive states (the general proof follows in App. D). While a relation \succsim_α ($\alpha \in \Gamma$) may violate Savage’s Archimedean axiom, we will ‘extrapolate’ it to a relation to which ‘Savage applies’. To get an idea, note that for incorporable objective events $I_1, I_2, \dots \subseteq \mathbf{S}$, we can successively refine the state space S_α to $S_{\alpha_1} = S_\alpha \vee \{I_1, \bar{I}_1\}$ (for a context α_1), then to $S_{\alpha_2} = S_{\alpha_1} \vee \{I_2, \bar{I}_2\}$ (for a context α_2), and so on. In each step another I_i becomes representable, and the new relation \succsim_{α_i} remains faithful to the earlier ones if it has the same outcome space. These refinements do not lead far enough: if S_α was finite, then all S_{α_i} are finite, hence still too small ‘for Savage’. We will thus go further: we will faithfully extrapolate each \succsim_α to a relation whose state space incorporates infinitely many and indeed *all* incorporable $I \subseteq \mathbf{S}$. This high state sophistication is purely hypothetical: it might never be reached by the agent in any context in Γ . The proof proceeds as follows, leaving out various difficulties:

- **Sufficiency** of the axioms is established by (i) showing that under Ax. 1–6 each extrapolated relation, denoted \succsim_α^+ , satisfies Savage’s axioms, (ii) deducing an expected-utility representation of each \succsim_α^+ via Savage’s Theorem in Kopylov’s (2007) version, and (iii) deducing suitable representations (U_α, P_α) of the original relations \succsim_α satisfying rules R1–R3.
- **Necessity** of the axioms is trivial in the case of the ‘local’ Ax. 1, 3 and 5, while the ‘non-local’ Ax. 2, 4 and 6 are proved using rules R1–R3.
- **The uniqueness property** of the representation is established by reducing it to the uniqueness property when representing the extrapolated relations, which is in turn obtained via Savage’s Theorem in Kopylov’s (2007) version.

C.1 Definition of extrapolated preferences

As mentioned, we extrapolate each relation \succsim_α by incorporating into the state space all incorporable objective events. In fact, we even incorporate all *weakly incorporable* objective events (in a shortly defined sense), because weakly incorporable objective events are more canonical. They form an algebra, and are probably the largest class suitable for incorporation along with preference extrapolation.

Definition 26 *An objective event $A \subseteq \mathbf{S}$ is **weakly incorporable** if there is a finite partition \mathcal{P} of \mathbf{S} at least as fine as $\{A, \bar{A}\}$ which the agent can always represent after (if necessary) refining states in a preference-neutral way: for all contexts $\alpha \in \Gamma$ there is a $\beta \in \Gamma$ (possibly equal to α) with $S_\beta = S_\alpha \vee \mathcal{P}$ and with \succsim_β faithful to \succsim_α . Let $\mathcal{I} := \{A \subseteq \mathbf{S} : A \text{ is weakly incorporable}\}$.*

Remark 20 *Incorporability implies weak incorporability: here $\mathcal{P} = \{A, \bar{A}\}$.*

Remark 21 *The set \mathcal{I} of weakly incorporable objective events is an algebra on \mathbf{S} : (i) $\mathbf{S} \in \mathcal{I}$; (ii) if $I \in \mathcal{I}$ (in virtue of partition \mathcal{P}) then $\bar{I} \in \mathcal{I}$ (in virtue of \mathcal{P}); (iii) if $I, I' \in \mathcal{I}$ (in virtue of \mathcal{P} resp. \mathcal{P}') then $I \cap I' \in \mathcal{I}$ (in virtue of $\mathcal{P} \vee \mathcal{P}'$).*

Given what was announced, one might expect that I refine each state space S_α to a partition S' of \mathbf{S} (a hypothetical subjective state space) in which all $I \in \mathcal{I}$ are representable, and to extrapolate the relation \succsim_α to one on $X_\alpha^{S'}$. It will in fact be easier to work not with a (hypothetical) subjective state space S' , but with the *objective* state space \mathbf{S} . So I will extrapolate \succsim_α to a relation on the set $X_\alpha^{\mathbf{S}}$ of ‘semi-objective acts’, which map objective states to subjective outcomes.

Definition 27 *A partition of \mathbf{S} **harmlessly refines** another one S if it is the join of S and some finite partition of \mathbf{S} into weakly incorporable objective events.*

Definition 28 *For a contexts $\alpha \in \Gamma$, the **extrapolated relation** \succsim_α^+ on $X_\alpha^{\mathbf{S}}$ is given as follows: $f \succsim_\alpha^+ g$ if and only if $f_\beta \succsim_\beta g_\beta$ for some context $\beta \in \Gamma$ such that (i) $f, g \in F_\beta^*$ (so f_β and g_β are defined) and (ii) S_β harmlessly refines S_α .³³*

C.2 Sufficiency of the axioms

Using extrapolated preferences, I now gradually prove sufficiency.

Definition 29 *Events $A \subseteq S_\alpha$ and $B \subseteq S_\beta$ ($\alpha, \beta \in \Gamma$) are **(objectively) equivalent** if $A^* = B^*$.*

Definition 30 *The join $\mathcal{R} \vee \mathcal{R}'$ of algebras \mathcal{R} and \mathcal{R}' on \mathbf{S} is the smallest algebra $\mathcal{A} \supseteq \mathcal{R} \cup \mathcal{R}'$ on \mathbf{S} , i.e., the closure of $\mathcal{R} \cup \mathcal{R}'$ under complement and finite union.*

An extrapolated relation \succsim_α^+ may still violate one of Savage’s axioms, by failing completeness: many functions in $X_\alpha^{\mathbf{S}}$ may be non-ranked. But \succsim_α^+ will be shown to be complete among functions measurable w.r.t. the ‘extrapolated algebra’:

Definition 31 *The **extrapolated algebra** for context $\alpha \in \Gamma$ is the set \mathcal{E}_α of objective events that are representable after a harmless state refinement: $\mathcal{E}_\alpha := \{A^* : A \subseteq S \text{ for some harmless refinement } S \text{ of } S_\alpha\}$.*

Lemma 5 *For all contexts $\alpha \in \Gamma$, \mathcal{E}_α is an algebra on \mathbf{S} , characterizable as*

- (1) *the join $(2^{S_\alpha})^* \vee \mathcal{I}$ of the algebra of representable objective events $(2^{S_\alpha})^*$ ($= \{A^* : A \subseteq S_\alpha\}$) and the algebra \mathcal{I} ,*
- (2) *the union $\cup_{\beta \in \Gamma: S_\beta \text{ harmlessly refines } S_\alpha} (2^{S_\beta})^*$ of each algebra $(2^{S_\beta})^*$ of representable objective events after some harmless refinement.*

³³Clause (ii) ensures that \succsim_α^+ is intimately linked to (i.e., ‘extrapolates’) \succsim_α .

I now recall Savage's theorem in the generalized version in which acts are measurable w.r.t. an arbitrary event algebra, not necessarily a σ -algebra, let alone the power set of the state space. It operates in a generalized framework:

Definition 32 A *generalized Savage framework* is a tuple $(X, (S, \mathcal{E}), \succsim)$ of a non-empty finite set X of '*outcomes*', a non-empty set S of '*states*' endowed with an algebra \mathcal{E} on S (the '*event algebra*'), and a '*preference*' relation \succsim on the set of \mathcal{E} -measurable functions from S to X ('*acts*').

An ordinary Savage framework (X, S, \succsim) is identified with the generalized one $(X, (S, 2^S), \succsim)$. In a generalized Savage framework $(X, (S, \mathcal{E}), \succsim)$ with sets of acts denoted F , Savage's well-known postulates can be stated as follows.

- P1:** \succsim is a transitive and complete relation on F .
- P2:** For all $f, g, f', g' \in F$ and $E \in \mathcal{E}$, if $f_E = f'_E$, $g_E = g'_E$, $f_{S \setminus E} = g_{S \setminus E}$ and $f'_{S \setminus E} = g'_{S \setminus E}$, then $f \succsim g \Leftrightarrow f' \succsim g'$.
- P3:** For all $x, y \in X$ and non-null $E \in \mathcal{E}$, $x \succsim_E y \Leftrightarrow x \succsim y$.³⁴
- P4:** For all $A, B \in \mathcal{E}$ and all $x \succ y$ and $x' \succ y'$ in X , $x_A y_{S \setminus A} \succsim x_B y_{S \setminus B} \Leftrightarrow x'_A y'_{S \setminus A} \succsim_\alpha x'_B y'_{S \setminus B}$.
- P5:** There exist $f, g \in F$ such that $f \succ g$.
- P6:** For all $f \succ g$ in F and $x \in X$, one can partition S into $A_1, \dots, A_n \in \mathcal{E}$ such that $f_{S \setminus A_i} x_{A_i} \succ g$ and $f \succ g_{S \setminus A_i} x_{A_i}$ for $i = 1, \dots, n$.

Lemma 6 (Savage's Theorem for arbitrary event algebras; see Kopylov 2007) A generalized Savage framework $(X, (S, \mathcal{E}), \succsim)$ satisfies Ax. P1–P6 if and only if there exist a non-constant utility function $U : X \rightarrow \mathbb{R}$ and a fine probability measure $P : \mathcal{E} \rightarrow [0, 1]$ such that $f \succsim g \Leftrightarrow \mathbb{E}_P(U \circ f) \geq \mathbb{E}_P(U \circ g)$ for all $f, g \in F$. Further, P is unique and U is unique up to increasing affine transformation.³⁵

Lemma 7 If Ax. 1–6 hold, then for each context $\alpha \in \Gamma$ Ax. P1–P6 hold for the generalized Savage framework $(X_\alpha, (S, \mathcal{E}), \succsim)$ in which (i) \mathcal{E} is \mathcal{E}_α or more generally any algebra such that $\mathcal{R} \subseteq \mathcal{E} \subseteq \mathcal{E}_\alpha$ for some algebra \mathcal{R} as in Ax. 6, and (ii) \succsim is \succsim_α^+ restricted to the set of acts $F = \{f \in X_\alpha^S : f \text{ is } \mathcal{E}\text{-measurable}\}$.

Lem. 7's proof rests on some technical lemmas (shown in App. F):

Lemma 8 Under Ax. 2, a relation \succsim_β is faithful to another \succsim_α if $X_\beta \supseteq X_\alpha$ and S_β harmlessly refines S_α .

³⁴Elements of C are identified with constant acts. An event is *null* if all acts that agree outside it are indifferent. An act (or outcome) f is *weakly preferred* to another g given $E \in \mathcal{E}$ – written $f \succsim_E g$ – if $f' \succsim g'$ for some acts f' and g' such that $f_E = f'_E$, $g_E = g'_E$ and $f'_{S \setminus E} = g'_{S \setminus E}$.

³⁵Kopylov proves this theorem for the case that \mathcal{E} is a *mosaic*, a more general structure than an algebra. My statement uses the condition that P is *fine*, which is equivalent in the algebra case to his condition that P is *finely ranged*. In Savage's special case $\mathcal{E} = 2^S$, a probability measure P on \mathcal{E} is fine if and only if it is atomless, and if and only if for all $A \in \mathcal{E}$ and $0 < \lambda < 1$ there is a $B \subseteq A$ in \mathcal{E} such that $P(B) = \lambda P(A)$. In general, fineness is a weaker condition.

Lemma 9 Under Ax. 2, whenever $f \succsim_{\alpha}^{+} g$ (where $\alpha \in \Gamma$ and $f, g \in X_{\alpha}^{\mathbf{S}}$), then

- (a) $f_{\beta} \succsim_{\beta} g_{\beta}$ for all (not just some) $\beta \in \Gamma$ satisfying (i)–(ii) in Def. 28,
- (b) $f_{\beta} \succsim_{\beta} g_{\beta}$ for some $\beta \in \Gamma$ such that (i)–(ii) in Def. 28 hold and \succsim_{β} is faithful to \succsim_{α} (in particular, $X_{\beta} \supseteq X_{\alpha}$).

Lemma 10 For any context $\alpha \in \Gamma$ and finite set $\mathcal{B} \subseteq \mathcal{E}_{\alpha}$, there is a context $\beta \in \Gamma$ such that (i) all $B \in \mathcal{B}$ are representable (i.e., $\mathcal{B} \subseteq (2^{S_{\beta}})^{*}$), (ii) S_{β} harmlessly refines S_{α} , and (iii) \succsim_{β} is faithful to \succsim_{α} .

Lemma 11 For all contexts $\alpha \in \Gamma$ and finite sets \mathcal{G} of \mathcal{E}_{α} -measurable functions from \mathbf{S} to X_{α} , there is a context $\beta \in \Gamma$ such that (i) $\mathcal{G} \subseteq F_{\beta}^{*}$, (ii) S_{β} harmlessly refines S_{α} , and (iii) \succsim_{β} is faithful to \succsim_{α} .

Lemma 12 Assume Ax. 2 and 5 and let $\alpha \in \Gamma$. For all acts f, g and events A of Lem. 7's generalized Savage framework, the conditional preference $f \succsim_A g$, i.e., $f \succsim_{\alpha, A}^{+} g$, holds if and only if $f_{\beta} \succsim_{\beta, A_{\beta}} g_{\beta}$ holds for some $\beta \in \Gamma$ such that $f, g \in F_{\beta}^{*}$, A is representable in context β , and S_{β} harmlessly refines S_{α} . The equivalence remains true when also requiring that \succsim_{β} is faithful to \succsim_{α} .

Lemma 13 Assume Ax. 2 and let $\alpha \in \Gamma$. An event A in Lem. 7's generalized Savage framework is non-null if and only if A_{β} is a non-null event in some context $\beta \in \Gamma$ such that A is representable (i.e., A_{β} is defined) and S_{β} harmlessly refines S_{α} . The equivalence remains true when also requiring that \succsim_{β} is faithful to \succsim_{α} .

Proof of Lem. 7. Assume Ax. 1–6. Let $\alpha, \mathcal{R}, \mathcal{E}$ be as specified. I show P1–P6 for the extrapolated relation \succsim_{α}^{+} restricted to $F := \{f \in X_{\alpha}^{\mathbf{S}} : f \text{ is } \mathcal{E}\text{-measurable}\}$.

Claim 1: P1 holds. To show completeness, let $f, g \in F$. Using Lem. 11, pick a $\beta \in \Gamma$ such that $f, g \in F_{\beta}^{*}$ and S_{β} harmlessly refines S_{α} . By Ax. 1, $f_{\beta} \succsim_{\beta} g_{\beta}$ or $g_{\beta} \succsim_{\beta} f_{\beta}$. In the first case $f \succsim_{\alpha}^{+} g$, in the second $g \succsim_{\alpha}^{+} f$. To show transitivity, let $f, g, h \in F$ such that $f \succsim_{\alpha}^{+} g$ and $g \succsim_{\alpha}^{+} h$. Using Lem. 11, pick a $\beta \in \Gamma$ such that $f, g, h \in F_{\beta}^{*}$ and S_{β} harmlessly refines S_{α} . So, as $f \succsim_{\alpha}^{+} g$ and $g \succsim_{\alpha}^{+} h$, we have $f_{\beta} \succsim_{\beta} g_{\beta}$ and $g_{\beta} \succsim_{\beta} h_{\beta}$ by Lem. 9. Hence, $f_{\beta} \succsim_{\beta} h_{\beta}$ by Ax. 1, and so $f \succsim_{\alpha}^{+} g$.

Claim 2: P2 holds. Consider $f, g, f', g' \in F$ and $E \in \mathcal{E}$ such that $f_E = f'_E$, $g_E = g'_E$, $f_{S \setminus E} = g_{S \setminus E}$ and $f'_{S \setminus E} = g'_{S \setminus E}$. Pick an $h \in F$ taking one value on E and another on \overline{E} (h exists as $|X_{\alpha}| \geq 2$ by Ax. 5). Using Lem. 11, pick a $\beta \in \Gamma$ such that $f, g, f', g', h \in F_{\beta}^{*}$ and S_{β} harmlessly refines S_{α} . As $f, g, g', g' \in F_{\beta}^{*}$, the acts $f_{\beta}, g_{\beta}, f'_{\beta}, g'_{\beta} \in F_{\beta}$ are defined; and as $h \in F_{\beta}^{*}$, the event E is representable in context β , so that E_{β} is defined (the sole purpose of introducing h was indeed to ensure representability of E). Note that $(f_{\beta})_{E_{\beta}} = (f'_{\beta})_{E_{\beta}}$, $(g_{\beta})_{E_{\beta}} = (g'_{\beta})_{E_{\beta}}$, $(f_{\beta})_{S \setminus E_{\beta}} = (g_{\beta})_{S \setminus E_{\beta}}$ and $(f'_{\beta})_{S \setminus E_{\beta}} = (g'_{\beta})_{S \setminus E_{\beta}}$. So, by Ax. 2 (or just 2*), $f_{\beta} \succsim_{\beta} g_{\beta} \Leftrightarrow f'_{\beta} \succsim_{\beta} g'_{\beta}$. This equivalence reduces to $f \succsim_{\alpha}^{+} g \Leftrightarrow f' \succsim_{\alpha}^{+} g'$ by Lem. 9.

Claim 3: P3 holds. Let $x, y \in X_{\alpha}$. Let $A \in \mathcal{E}$ be non-null. I show $x \succsim_{\alpha, A}^{+} y \Leftrightarrow x \succsim_{\alpha}^{+} y$. By Lem. 13, A_{β} is non-null for a $\beta \in \Gamma$ such that A is representable,

$S_\beta = S_\alpha \vee \mathcal{P}$ with a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbf{S} , and \succsim_β is faithful to \succsim_α . First, if $x \succsim_\alpha^+ y$, then $x \succsim_\beta y$ by Lem. 9, so $x \succsim_{\beta, A_\beta} y$ by Ax. 3 and A_β 's non-nullness, hence $x \succsim_\alpha^+ y$ by Lem. 12. Now let $x \succsim_{\alpha, A}^+ y$. By Lem. 12, $x \succsim_{\gamma, A_\gamma} y$ for a $\gamma \in \Gamma$ such that A is representable, $S_\gamma = S_\alpha \vee \mathcal{Q}$ with a finite partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbf{S} , and \succsim_γ is faithful to \succsim_α . Using Lem. 10, pick $\delta \in \Gamma$ such that $\mathcal{P} \cup \mathcal{Q} \subseteq (2^{S_\delta})^*$, $S_\delta = S_\alpha \vee \mathcal{P}'$ with a finite partition $\mathcal{P}' \subseteq \mathcal{I}$ of \mathbf{S} , and \succsim_δ is faithful to \succsim_α . W.l.o.g. X_β and X_γ equal X_α (by independence between outcome and state awareness) and \mathcal{P}' refines \mathcal{P} and \mathcal{Q} (otherwise replace \mathcal{P}' by $\mathcal{P}' \vee \mathcal{P} \vee \mathcal{Q}$). Now \succsim_δ is faithful to \succsim_β and \succsim_γ , each time by Lem. 8, using that $X_\delta \supseteq X_\beta = X_\gamma (= X_\alpha)$ and that $S_\delta = S_\beta \vee \mathcal{P}' = S_\gamma \vee \mathcal{P}'$ (since each set equals $S_\alpha \vee \mathcal{P}'$ as \mathcal{P}' refines \mathcal{P} and \mathcal{Q}). As $A_\beta (\subseteq S_\beta)$ is non-null and \succsim_δ is faithful to \succsim_β , $A_\delta (\subseteq S_\delta)$ is non-null. As $x \succsim_{\gamma, A_\gamma} y$ and \succsim_δ is faithful to \succsim_γ , $x \succsim_{\delta, A_\delta} y$. So $x \succsim_\delta y$ by Ax. 3. Thus $x \succsim_\alpha^+ y$.

Claim 4: P4 holds. Let $A, B \in \mathcal{E}$ and $x, y, x', y' \in X_\alpha$ such that $x \succ_\alpha^+ y$ and $x' \succ_\alpha^+ y'$. I show $x_A y_{\mathbf{S} \setminus A} \succsim_\alpha^+ x_B y_{\mathbf{S} \setminus B} \Leftrightarrow x'_A y'_{\mathbf{S} \setminus A} \succsim_\alpha^+ x'_B y'_{\mathbf{S} \setminus B}$. Via Lem. 11, pick a $\beta \in \Gamma$ such that $x_A y_{\mathbf{S} \setminus A}, x_B y_{\mathbf{S} \setminus B}, x'_A y'_{\mathbf{S} \setminus A}, x'_B y'_{\mathbf{S} \setminus B} \in F_\beta^*$ and S_β harmlessly refines S_α . By Lem. 9, $x \succ_\beta y$ and $x' \succ_\beta y'$. So the claimed equivalence reduces to $(x_A y_{\mathbf{S} \setminus A})_\beta \succsim_\beta (x_B y_{\mathbf{S} \setminus B})_\beta \Leftrightarrow (x'_A y'_{\mathbf{S} \setminus A})_\beta \succsim_\beta (x'_B y'_{\mathbf{S} \setminus B})_\beta$, i.e., $x_{A_\beta} y_{S_\beta \setminus A_\beta} \succsim_\beta x_{B_\beta} y_{S_\beta \setminus B_\beta} \Leftrightarrow x'_{A_\beta} y'_{S_\beta \setminus A_\beta} \succsim_\beta x'_{B_\beta} y'_{S_\beta \setminus B_\beta}$. This holds by Ax. 4.

Claim 5: P5 holds. Using Ax. 5, pick $f \succ_\alpha g$ in F_α . Clearly, $f^* \succ_\alpha^+ g^*$.

Claim 6: P6 holds. Let $f \succ_\alpha^+ g$ in F and $x \in X_\alpha$. As $f \succ_\alpha^+ g$, we have $f_\beta \succsim_\beta g_\beta$ for a $\beta \in \Gamma$ such that $f, g \in F_\beta^*$, $S_\beta = S_\alpha \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbf{S} , and \succsim_β is faithful to \succsim_α . Note $x \in X_\beta$; and $g_\beta \not\succsim_\beta f_\beta$ as $g \not\succsim_\alpha^+ f$. So $f_\beta \succ_\beta g_\beta$. As $\mathcal{R} (\subseteq \mathcal{E})$ is as in Ax. 6, one can partition \mathbf{S} into A_1, \dots, A_n from \mathcal{R} (hence from \mathcal{E}) such that, for some $\gamma \in \Gamma$ with $S_\gamma = S_\beta \vee \{A_1, \dots, A_n\}$ and $X_\gamma \supseteq X_\beta$, $(f_\gamma)_{S_\gamma \setminus (A_i)_\gamma} x_{(A_i)_\gamma} \succ_\gamma g_\gamma$ and $f_\gamma \succ_\gamma (g_\gamma)_{S_\gamma \setminus (A_i)_\gamma} x_{(A_i)_\gamma}$ for all i , i.e., $(f_{\mathbf{S} \setminus A_i} x_{A_i})_\gamma \succ_\gamma g_\gamma$ and $f_\gamma \succ_\gamma (g_{\mathbf{S} \setminus A_i} x_{A_i})_\gamma$ for all i . So (as S_γ harmlessly refines S_α , being the join of S_α and $\mathcal{P} \vee \{A_1, \dots, A_n\} \subseteq \mathcal{I}$), $f_{\mathbf{S} \setminus A_i} x_{A_i} \succ_\alpha^+ g$ and $f \succ_\alpha^+ g_{\mathbf{S} \setminus A_i} x_{A_i}$ for all i .³⁶ ■

Given Ax. 1–6, for each $\alpha \in \Gamma$ we now use Lem. 6 and 7 to pick a utility function U_α on X_α and a fine probability measure P_α^+ on \mathcal{E}_α which represent the extrapolated relation \succsim_α^+ on $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{E}_\alpha\text{-measurable}\}$:

$$f \succsim_\alpha^+ g \Leftrightarrow \mathbb{E}_{P_\alpha^+}(U_\alpha \circ f) \geq \mathbb{E}_{P_\alpha^+}(U_\alpha \circ g) \text{ for all } \mathcal{E}\text{-measurable } f, g \in X_\alpha^{\mathbf{S}}.$$

Each P_α^+ induces a probability measure P_α on the subjective event space 2^{S_α} via

$$P_\alpha(E) := P_\alpha^+(E^*) \text{ for all } E \subseteq S_\alpha.$$

The next four lemmas complete the sufficiency proof by establishing that the functions P_α and U_α ($\alpha \in \Gamma$) have all properties required in Thm. 1.

³⁶To make the last step, one needs to first decompose each strict preference (\succ_γ) into a weak preference (\succsim_γ) without weak dispreference ($\not\succsim_\gamma$), then infer corresponding extended weak preferences (\succsim_α^+) without weak dispreference ($\not\succsim_\alpha^+$) using Lem. 9, which implies extended strict preferences (\succ_α^+).

Lemma 14 *Under Ax. 1–6, the above-defined system $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is an expected-utility representation.*

Lemma 15 *Under Ax. 1–6, the above-defined functions U_α satisfy R1.*

Lemma 16 *Under Ax. 1–6, the above-defined functions P_α satisfy R2.*

Lemma 17 *Under Ax. 1–6, for each algebra \mathcal{R} as in Ax. 6,*

- (a) *all above-defined measures P_α^+ have identical restriction $\rho := P_\alpha^+|_{\mathcal{R}}$,*
- (b) *the above-defined measures P_α satisfy R3 in virtue of ρ .*

I begin by proving the first of these four ‘sufficiency lemmas’.

Proof of Lem. 14. Assume Ax. 1–6. Let $\alpha \in \Gamma$ and $f, g \in F_\alpha$. Let U_α, P_α and P_α^+ be as above. I show $f \succsim_\alpha g \Leftrightarrow \mathbb{E}_{P_\alpha}(U_\alpha \circ f) \geq \mathbb{E}_{P_\alpha}(U_\alpha \circ g)$. The left side reduces to $f^* \succsim_\alpha^+ g^*$ by Lem. 9, and the right side to $\mathbb{E}_{P_\alpha^+}(U_\alpha \circ f^*) \geq \mathbb{E}_{P_\alpha^+}(U_\alpha \circ g^*)$ because, letting $\tau : \mathbf{S} \rightarrow S_\alpha$ map any $s \in \mathbf{S}$ to its subjectivization $\tau(s) = s_\alpha$, we have $f^* = f \circ \tau$, $g^* = g \circ \tau$, and P_α is P_α^+ ’s image under τ . To complete the proof, note $f^* \succsim_\alpha^+ g^* \Leftrightarrow \mathbb{E}_{P_\alpha^+}(U_\alpha \circ f^*) \geq \mathbb{E}_{P_\alpha^+}(U_\alpha \circ g^*)$ by definition of U_α and P_α^+ . ■

Proving the other three ‘sufficiency lemmas’ requires further results. I begin with two cornerstone results from the literature:

Lemma 18 (Niiniluoto 1972, Wakker 1981) *Every fine and tight qualitative probability relation on an algebra \mathcal{E} on \mathbf{S} (not necessarily a σ -algebra) is uniquely representable by a probability measure on \mathcal{E} .*

Lemma 19 (Wakker 1981, Kopylov 2007³⁷) *A probability measure on an algebra \mathcal{E} on \mathbf{S} (not necessarily a σ -algebra) is fine if and only if the represented qualitative probability relation is fine and tight.*

I also need five technical lemmas (proved in App. F), the last two about extrapolated preferences, and the first three about the *extrapolated belief* relation over objective events induced by \succsim_α^+ and denoted again by ‘ \succsim_α^+ ’.

Lemma 20 (extrapolated comparative beliefs) *Under Ax. 2, 4 and 5, for all $\alpha \in \Gamma$ and $A, B \subseteq \mathbf{S}$, $A \succsim_\alpha^+ B$ if and only if $A_\beta \succsim_\beta B_\beta$ for some $\beta \in \Gamma$ such that A and B are representable (i.e., A_β and B_β are defined) and S_β harmlessly refines S_α . The equivalence remains true when also requiring \succsim_β to be faithful to \succsim_α .*

Lemma 21 *Under Ax. 2, 4, 5 and 6, whenever $A \succsim_\alpha^+ B$ (where $\alpha \in \Gamma$ and $A, B \subseteq \mathbf{S}$), then $A_\beta \succsim_\beta B_\beta$ for each (not just some) context $\beta \in \Gamma$ in which A and B are representable (so that A_β and B_β are defined) and S_β harmlessly refines S_α .*

³⁷Lem. 19 is implicit in Wakker (1981) and a special case of Kopylov’s (2007) Thm. A.1.

Lemma 22 Under Ax. 2, 4 and 5, the extrapolated relations \succsim_α^+ ($\alpha \in \Gamma$) agree (as belief relations on $2^{\mathbf{S}}$) on each robust algebra \mathcal{R} of incorporable objective events.

Lemma 23 Under Ax. 1–6, the restriction of the above-defined measure P_α^+ to an algebra \mathcal{R} of type in Ax. 6 is (a) fine, and (b) the same for all $\alpha \in \Gamma$.

Lemma 24 Given Ax. 1 and 2, for any contexts $\alpha, \beta \in \Gamma$, if S_β harmlessly refines S_α then $\mathcal{E}_\alpha = \mathcal{E}_\beta$, and if moreover \succsim_β is faithful to \succsim_α then $\succsim_\beta^+ = \succsim_\alpha^+$.

Lemma 25 (stability of nullness) Under Ax. 2, any null event A of some context is null in all contexts $\alpha \in \Gamma$ where it is conceived, i.e., where $A \subseteq S_\alpha$.

Proof of Lem. 17. Assume Ax. 1–6. Let \mathcal{R} be as in Ax. 6, and P_α and P_α^+ ($\alpha \in \Gamma$) as above. By Lem. 23, $\rho := P_\alpha^+|_{\mathcal{R}}$ is fine and independent of $\alpha \in \Gamma$. I show ρ is uncontroversial. Let $A \in \mathcal{R}$ and $\alpha \in \Gamma$. I must show existence of a $\beta \in \Gamma$ such that $S_\beta = S_\alpha \vee \{A, \bar{A}\}$, P_β^* extends P_α^* , and $P_\beta^*(A) = \rho(A)$. As $A \in \mathcal{R}$, A is incorporable; so pick a $\beta \in \Gamma$ such that $S_\beta = S_\alpha \vee \{A, \bar{A}\}$ and \succsim_β is faithful to \succsim_α . By Lem. 24, $\mathcal{E}_\alpha = \mathcal{E}_\beta$ and $\succsim_\alpha^+ = \succsim_\beta^+$. So $P_\alpha^+ = P_\beta^+$. Thus $P_\beta^* (= P_\beta^+|_{(2^{S_\beta})^*})$ extends $P_\alpha^* (= P_\alpha^+|_{(2^{S_\alpha})^*})$. Finally, $P_\beta^*(A) = P_\beta^+(A) = \rho(A)$. ■

Proof of Lem. 16. Assume Ax. 1–6. Let P_α, P_α^+ ($\alpha \in \Gamma$) be as above, $\alpha, \beta \in \Gamma$, and $S := S_\alpha \cap S_\beta$. If S is null in both contexts, P_α and P_β are zero, so proportional, on 2^S . Now let S be non-null in one, hence by Lem. 25 both, contexts. Let \mathcal{R} be as in Ax. 6. Put $\mathcal{E} := \{A^* : A \subseteq S \vee \mathcal{P} \text{ for a finite partition } \mathcal{P} \subseteq \mathcal{R} \text{ of } \mathbf{S}\}$. Here $S \vee \mathcal{P}$ joins partitions of *distinct* sets S^* and \mathbf{S} ; Def. 12 still applies.

Claim 1: The measures P_α^+ and P_β^+ are ordinally equivalent on \mathcal{E} . Note \mathcal{E} is an algebra on S^* , not \mathbf{S} .³⁸ Let $A, B \in \mathcal{E}$. I show $P_\alpha^+(A) \geq P_\alpha^+(B) \Leftrightarrow P_\beta^+(A) \geq P_\beta^+(B)$, or equivalently (as P_α^+ and P_β^+ represent \succsim_α^+ resp. \succsim_β^+) $A \succsim_\alpha^+ B \Leftrightarrow A \succsim_\beta^+ B$. As $A, B \in \mathcal{E}$, we may pick finite partitions $\mathcal{P}_A, \mathcal{P}_B \subseteq \mathcal{R}$ of \mathbf{S} such that $A \in (2^{S \vee \mathcal{P}_A})^*$ and $B \in (2^{S \vee \mathcal{P}_B})^*$. Clearly, $\mathcal{P} := \mathcal{P}_A \vee \mathcal{P}_B$ is again a finite partition of \mathbf{S} . Using that all $X \in \mathcal{P}$ are incorporable (as $\mathcal{P} \subseteq \mathcal{R}$), pick $\alpha', \beta' \in \Gamma$ such that $S_{\alpha'} = S_\alpha \vee \mathcal{P}$ and $S_{\beta'} = S_\beta \vee \mathcal{P}$. Now A and B are representable in context α' (as $S_{\alpha'}$ refines $S_\alpha \vee \mathcal{P}_A$); so $A \succsim_\alpha^+ B \Leftrightarrow A_{\alpha'} \succsim_{\alpha'} B_{\alpha'}$ by Lem. 20 and 21. Similarly, A and B are representable in context β' ; so $A \succsim_\beta^+ B \Leftrightarrow A_{\beta'} \succsim_{\beta'} B_{\beta'}$. It remains to show $A_{\alpha'} \succsim_{\alpha'} B_{\alpha'} \Leftrightarrow A_{\beta'} \succsim_{\beta'} B_{\beta'}$. This holds by comparative-belief stability (Prop. 4), since $A_{\alpha'} = A_{\beta'}$ and $B_{\alpha'} = B_{\beta'}$ as $S_{\alpha'}$ and $S_{\beta'}$ agree within S^* ($\supseteq A, B$).

Claim 2: P_α and P_β are proportional on 2^S . By Claim 1, the conditional probability measures $P_\alpha^+(\cdot|S^*)$ and $P_\beta^+(\cdot|S^*)$ are ordinally equivalent on \mathcal{E} . Their restrictions $P_\alpha^+(\cdot|S^*)|_{\mathcal{E}}$ and $P_\beta^+(\cdot|S^*)|_{\mathcal{E}}$ are probability measures on \mathcal{E} (an algebra on S^*), which are fine as P_α^+ and P_β^+ are fine. So $P_\alpha^+(\cdot|S^*)|_{\mathcal{E}} = P_\beta^+(\cdot|S^*)|_{\mathcal{E}}$ by Lem. 18 and 19. Hence, P_α^+ is proportional to P_β^+ on \mathcal{E} , and thus on $(2^S)^*$ ($\subseteq \mathcal{E}$). So, P_α is proportional to P_β on 2^S . ■

³⁸ \mathcal{E} is also the join of algebras on S^* : $\{A^* : A \subseteq S\}$ and $\{A \cap S^* : A \in \mathcal{R}\}$ (\mathcal{R} 's trace in S^*).

Proof of Lem. 15. Assume Ax. 1–6. Let U_α, P_α^+ ($\alpha \in \Gamma$) be as above. Fix $\alpha, \beta \in \Gamma$. Put $X := X_\alpha \cap X_\beta$. For all $x, y \in X$, $x \succsim_\alpha y \Leftrightarrow x \succsim_\beta y$ by outcome-preference stability (Prop. 3); so $U_\alpha(x) \geq U_\alpha(y) \Leftrightarrow U_\beta(x) \geq U_\beta(y)$ by Lem. 14. If U_α (and so U_β) is constant on X , then U_α is an increasing affine transformation of U_β on X . Now let U_α (and so U_β) be non-constant on X . Let \mathcal{R} be as in Ax. 6. As (U_α, P_α^+) represents \succsim_α^+ restricted to the \mathcal{E}_α -measurable acts, $(U_\alpha|_X, P_\alpha^+|_{\mathcal{R}})$ represents \succsim_α^+ restricted further to \mathcal{R} -measurable acts mapping into X , i.e., to $F := \{f \in X^{\mathbf{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$. For analogous reasons, $(U_\beta|_X, P_\beta^+|_{\mathcal{R}})$ represents \succsim_β^+ restricted to F . (In fact $P_\alpha^+|_{\mathcal{R}} = P_\beta^+|_{\mathcal{R}}$ by Lem. 17.) Next I show that \succsim_α^+ and \succsim_β^+ coincide on F . Let $f, g \in F$. As $f, g \in \mathcal{E}_\alpha$ we may by Lem. 11 pick an $\alpha' \in \Gamma$ such that $f, g \in F_{\alpha'}^*$ and $S_{\alpha'}$ harmlessly refines S_α . Analogously, as $f, g \in \mathcal{E}_\beta$ we may pick a $\beta' \in \Gamma$ such that $f, g \in F_{\beta'}^*$ and $S_{\beta'}$ harmlessly refines S_β . Now $f \succsim_\alpha^+ g \Leftrightarrow f \succsim_\beta^+ g$, as by Lem. 9 this reduces to $f_{\alpha'} \succsim_{\alpha'} g_{\alpha'} \Leftrightarrow f_{\beta'} \succsim_{\beta'} g_{\beta'}$, which holds since $f (= (f_{\alpha'})^* = (f_{\beta'})^*)$ and $g (= (g_{\alpha'})^* = (g_{\beta'})^*)$ are measurable w.r.t. a robust algebra. As just shown, $(U_\alpha|_X, P_\alpha^+|_{\mathcal{R}})$ and $(U_\beta|_X, P_\beta^+|_{\mathcal{R}})$ represent the *same* relation on F ; note also that $U_\alpha|_X$ and $U_\beta|_X$ are non-constant and $P_\alpha^+|_{\mathcal{R}}$ and $P_\beta^+|_{\mathcal{R}}$ are fine. So $U_\alpha|_X$ is an increasing affine transformation of $U_\beta|_X$ by Lem. 7. ■

C.3 Necessity of the axioms

I now show that our representation implies all axioms. I start by the ‘local’ Axioms 1, 3 and 5, and then turn to the ‘global’ (cross-context) Axioms 2, 4 and 6.

Lemma 26 *Given an expected-utility representation, Ax. 1, 3 and 5 hold.*

Proof. If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is such a representation, then Ax. 1 holds trivially, Ax. 5 holds by non-constancy of all U_α , and Ax. 3 holds by definition of conditional preference (using that non-null events $E \subseteq S_\alpha$ have probabilities $P_\alpha(E) \neq 0$). ■

Lemma 27 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1’s sense, Ax. 2 holds.*

Proof. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ be a representation. Let $\alpha, \alpha' \in \Gamma$, $f, g \in F_\alpha$, $f', g' \in F_{\alpha'}$, and $E \subseteq S_\alpha \cap S_{\alpha'}$, such that $f_E = f'_E$, $g_E = g'_E$, $f_{S_\alpha \setminus E} = g_{S_\alpha \setminus E}$ and $f'_{S_{\alpha'} \setminus E} = g'_{S_{\alpha'} \setminus E}$. I must show $f \succsim_\alpha g \Leftrightarrow f' \succsim_{\alpha'} g'$, i.e., $\mathbb{E}_{P_\alpha}(U_\alpha \circ f) \geq \mathbb{E}_{P_\alpha}(U_\alpha \circ g) \Leftrightarrow \mathbb{E}_{P_{\alpha'}}(U_{\alpha'} \circ f') \geq \mathbb{E}_{P_{\alpha'}}(U_{\alpha'} \circ g')$, or $\int_E U_\alpha \circ f \, dP_\alpha \geq \int_E U_\alpha \circ g \, dP_\alpha \Leftrightarrow \int_E U_{\alpha'} \circ f' \, dP_{\alpha'} \geq \int_E U_{\alpha'} \circ g' \, dP_{\alpha'}$ as $f_{S_\alpha \setminus E} = g_{S_\alpha \setminus E}$ and $f'_{S_{\alpha'} \setminus E} = g'_{S_{\alpha'} \setminus E}$. The latter holds as (i) P_α is proportional to $P_{\alpha'}$ within E , (ii) $f_E = f'_E$ and $g_E = g'_E$, and (iii) U_α is an increasing affine transformation of $U_{\alpha'}$ on $X_\alpha \cap X_{\alpha'}$ (where by (ii)–(iii) $U_\alpha \circ f$ is an increasing affine transformation of $U_{\alpha'} \circ f'$ on E , and $U_\alpha \circ g$ is one of $U_{\alpha'} \circ g'$ on E). ■

Lemma 28 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1’s sense, Ax. 4 holds.*

Proof. Assume $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation. Let $\alpha, \alpha' \in \Gamma$ such that $S := S_\alpha = S_{\alpha'}$, let $E, D \subseteq S$, and consider $x \succ_\alpha y$ in X_α and $x' \succ_{\alpha'} y'$ in $X_{\alpha'}$. I claim that $x_E y_{S \setminus E} \succ_\alpha x_D y_{S \setminus D} \Leftrightarrow x'_E y'_{S \setminus E} \succ_{\alpha'} x'_D y'_{S \setminus D}$. Noting that $U_\alpha(x) > U_\alpha(y)$ (as $x \succ_\alpha y$) and $U_{\alpha'}(x') > U_{\alpha'}(y')$ (as $x' \succ_{\alpha'} y'$), the claimed equivalence reduces to the equivalence $P_\alpha(E) \geq P_\alpha(D) \Leftrightarrow P_{\alpha'}(E) \geq P_{\alpha'}(D)$, which holds as P_α is proportional (in fact, identical) to $P_{\alpha'}$ on the full domain $2^S (= 2^{S_\alpha} = 2^{S_{\alpha'}})$. ■

Lemma 29 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1's sense, with a fine uncontroversial $\rho : \mathcal{R} \rightarrow [0, 1]$ in R3, then Ax. 6 holds in virtue of algebra \mathcal{R} .*

Proof. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$, ρ and \mathcal{R} be as assumed.

Claim 1: All $A \in \mathcal{R}$ are incorporable. Let $A \in \mathcal{R}$ and $\alpha \in \Gamma$. As ρ is uncontroversial, there is a $\beta \in \Gamma$ where $S_\beta = S_\alpha \vee \{A, \bar{A}\}$ and P_β^* extends P_α^* . W.l.o.g. $X_\beta = X_\alpha$ (by independence between outcome and state awareness); so \succ_β is faithful to \succ_α , using R1 and the fact that P_β^* extends P_α^* .

Claim 2: \mathcal{R} is robust. Let $\alpha_1, \alpha_2 \in \Gamma$. Consider \mathcal{R} -determined acts $f_1, g_1 \in F_{\alpha_1}$ and $f_2, g_2 \in F_{\alpha_2}$ such that $f_1^* = f_2^* =: f$ and $g_1^* = g_2^* =: g$. I show that $f_1 \succ_{\alpha_1} g_1 \Leftrightarrow f_2 \succ_{\alpha_2} g_2$, i.e., $\mathbb{E}_{P_{\alpha_1}}(U_{\alpha_1} \circ f_1) \geq \mathbb{E}_{P_{\alpha_1}}(U_{\alpha_1} \circ g_1) \Leftrightarrow \mathbb{E}_{P_{\alpha_2}}(U_{\alpha_2} \circ f_2) \geq \mathbb{E}_{P_{\alpha_2}}(U_{\alpha_2} \circ g_2)$. As f and g are \mathcal{R} -measurable and $P_{\alpha_i}^*(A) = \rho(A)$ for all $i \in \{1, 2\}$ and all $A \in \mathcal{R} \cap (2^{S_{\alpha_i}})^*$, the desired equivalence reduces to $\mathbb{E}_\rho(U_{\alpha_1} \circ f) \geq \mathbb{E}_\rho(U_{\alpha_1} \circ g) \Leftrightarrow \mathbb{E}_\rho(U_{\alpha_2} \circ f) \geq \mathbb{E}_\rho(U_{\alpha_2} \circ g)$, which holds by R1 and the fact that $X_{\alpha_1} \cap X_{\alpha_2} \supseteq f(S_{\alpha_1}), g(S_{\alpha_2})$.

Claim 3: \mathcal{R} has the additional property required in Ax. 6. Let $\alpha \in \Gamma$, $f \succ_\alpha g$ in F_α , and $x \in X_\alpha$. For any $\epsilon > 0$, pick (i) a finite partition $\mathcal{P}_\epsilon \subseteq \mathcal{R}$ of \mathbf{S} such that $\rho(A) \leq \epsilon$ for all $A \in \mathcal{P}_\epsilon$ (using ρ 's fineness) and (ii) an $\alpha_\epsilon \in \Gamma$ such that $S_{\alpha_\epsilon} = S_\alpha \vee \mathcal{P}_\epsilon$, $X_{\alpha_\epsilon} = X_\alpha$, and $P_{\alpha_\epsilon}^*$ extends P_α^* (using ρ 's uncontroversialness and the independence between outcome and state awareness); let $f^\epsilon, g^\epsilon \in F_{\alpha_\epsilon}$ be the acts equivalent to f resp. g . It suffices to show that (*) for small enough $\epsilon > 0$, $\mathbb{E}_{P_{\alpha_\epsilon}}(U \circ ((f^\epsilon)_{S_{\alpha_\epsilon} \setminus A_{\alpha_\epsilon}} x_{A_{\alpha_\epsilon}})) > \mathbb{E}_{P_{\alpha_\epsilon}}(U \circ g^\epsilon)$ for all $A \in \mathcal{P}_\epsilon$, and (**) for small enough $\epsilon > 0$, $\mathbb{E}_{P_{\alpha_\epsilon}}(U \circ f^\epsilon) > \mathbb{E}_{P_{\alpha_\epsilon}}(U \circ ((g^\epsilon)_{S_{\alpha_\epsilon} \setminus A_{\alpha_\epsilon}} x_{A_{\alpha_\epsilon}}))$ for all $A \in \mathcal{P}_\epsilon$. As all U_{α_ϵ} have same domain as U_α , they are increasing affine transformations of U_α . W.l.o.g. let $U_{\alpha_\epsilon} = U_\alpha =: U$ for all $\epsilon > 0$. Given $\epsilon > 0$, each $(f^\epsilon)_{S_{\alpha_\epsilon} \setminus A_{\alpha_\epsilon}} x_{A_{\alpha_\epsilon}}$ ($A \in \mathcal{P}_\epsilon$) differs from f^ϵ at most on A_{α_ϵ} , hence at most with $(P_{\alpha_\epsilon}$ -)probability ϵ . Put $\Delta := \max_{x, y \in X_\alpha} |U(x) - U(y)|$. Now $|\mathbb{E}_{P_{\alpha_\epsilon}}(U \circ ((f^\epsilon)_{S_{\alpha_\epsilon} \setminus A_{\alpha_\epsilon}} x_{A_{\alpha_\epsilon}})) - \mathbb{E}_{P_{\alpha_\epsilon}}(U \circ f^\epsilon)| \leq \epsilon \Delta$ for all $A \in \mathcal{P}_\epsilon$. This implies (*) since $\mathbb{E}_{P_{\alpha_\epsilon}}(U \circ f^\epsilon) = \mathbb{E}_{P_\alpha}(U \circ f) > \mathbb{E}_{P_\alpha}(U \circ g) = \mathbb{E}_{P_{\alpha_\epsilon}}(U \circ g^\epsilon)$, where the ' $>$ ' holds as $f \succ_\alpha g$, and both '=' hold as f^ϵ (g^ϵ) is equivalent to f (g) and $P_{\alpha_\epsilon}^*$ extends P_α^* . An analogous argument shows (**). ■

C.4 Uniqueness of the representation

I now prove uniqueness, based on two technical lemma (shown in App. F):

Lemma 30 *If a probability measure ρ on an \mathcal{R} is uncontroversial among probability measures P_α on 2^{S_α} ($\alpha \in \Gamma$) which satisfy R2, then for each $\alpha \in \Gamma$ there is a*

probability measure ρ_α on $\mathcal{R}_\alpha := \mathcal{R} \vee (2^{S_\alpha})^*$ which extends all P_β^* for which S_β is the join on S_α and a finite partition $\mathcal{P} \subseteq \mathcal{R}$ (so ρ_α extends ρ as ρ is uncontroversial).

Lemma 31 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1's sense with a fine uncontroversial measure ρ on \mathcal{R} , and ρ_α and \mathcal{R}_α ($\alpha \in \Gamma$) are as in Lem. 30, then, for all $\alpha \in \Gamma$, (U_α, ρ_α) represents the restriction of \succsim_α^+ to $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{R}_\alpha\text{-measurable}\}$ in Lem. 6's sense.*

Lemma 32 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ and $(U'_\alpha, P'_\alpha)_{\alpha \in \Gamma}$ are representations in Thm. 1's sense, then each P_α equals P'_α and each U_α is an increasing affine transformation of U'_α .*

Proof. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ be a representation, with a fine uncontroversial measure $\rho : \mathcal{R} \rightarrow [0, 1]$; so Ax. 1–6 hold. Let ρ_α and \mathcal{R}_α ($\alpha \in \Gamma$) be as in Lem. 30 and 31. Let $(U'_\alpha, P'_\alpha)_{\alpha \in \Gamma}$ be the representation defined in in App. C.2 under Ax. 1–6 using the objects \succsim_α^+ and P_α^+ (it was formerly denoted ' $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ '). Fix $\alpha \in \Gamma$. I show $P'_\alpha = P_\alpha$ and $U'_\alpha = a_\alpha U_\alpha + b_\alpha$ with $a_\alpha > 0$ and $b_\alpha \in \mathbb{R}$. As (U'_α, P'_α) represents \succsim_α^+ on $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{E}_\alpha\text{-measurable}\}$ (in Lem. 6's sense), $(U'_\alpha, P'_\alpha|_{\mathcal{R}_\alpha})$ represents the same relation as (U_α, ρ_α) on $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{R}_\alpha\text{-measurable}\}$ by Lem. 31. So by Lem. 6 $\rho_\alpha = P'_\alpha|_{\mathcal{R}_\alpha}$ and U_α is an increasing affine transformation of U'_α . Finally, $P_\alpha = P'_\alpha$: for all $E \subseteq S_\alpha$, $P_\alpha(E) = P_\alpha^*(E^*) = \rho_\alpha(E^*) = P_\alpha^+(E^*) = P'_\alpha(E)$. ■

D Proof of Theorem 1 for the general case

From now on states can be non-exhaustive. I prove Thm. 1 by reduction to the case of exhaustive states where it has been established. Let Π be the partition of Γ into non-empty sets of contexts such that $\alpha, \beta \in \Gamma$ belong to the same set in Π if and only if $\mathbf{S}_\alpha = \mathbf{S}_\beta$. This yields for each $\Delta \in \Pi$ a (sub)framework $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Delta}$ with exhaustive states, called the Δ -**subframework**, to which we may apply Thm. 1; let \mathbf{S}_Δ be its set of objective states. For all $\gamma \in \Gamma$, let Δ_γ be the member of Π containing γ , and, generalizing earlier objects, let \succsim_γ^+ , \mathcal{E}_γ and P_γ^+ be defined as in App. C, but *w.r.t. the Δ_γ -subframework* (which has exhaustive states, ensuring well-definedness); so \succsim_γ^+ is a relation on $X_\gamma^{\mathbf{S}_{\Delta_\gamma}}$, and \mathcal{E}_γ is an algebra on $\mathbf{S}_{\Delta_\gamma}$. The *trace in \mathbf{S}' ($\subseteq \mathbf{S}$)* of an algebra \mathcal{R} on \mathbf{S} is the algebra on \mathbf{S}' given by $\mathcal{R}|_{\mathbf{S}'} := \{A \cap \mathbf{S}' : A \in \mathcal{R}\}$.³⁹

D.1 Sufficiency of the axioms

Our reductive proof draws on a technical lemma shown in App. F:

³⁹ A direct, non-reductive proof of Theorem 1 would also work, by generalizing App. C's proof strategy and defining the objects \mathcal{I} , \succsim_γ^+ , \mathcal{E}_γ and P_γ^+ ($\gamma \in \Gamma$) directly relative to the *general* framework (here \succsim_γ^+ is a relation on $X_\gamma^{\mathbf{S}}$, and \mathcal{E}_γ an algebra on \mathbf{S}).

Lemma 33 *If Ax. 1–6 hold, then they hold for each Δ -subframework ($\Delta \in \Pi$).*

Now assume Ax. 1–6. By Lem. 33, each Δ -subframework ($\Delta \in \Pi$) satisfies Ax. 1–6. So by Thm. 1 each Δ -subframework ($\Delta \in \Pi$) has a representation $(U_\alpha, P_\alpha)_{\alpha \in \Delta}$ in Thm. 1's sense. Joining these representations together, we obtain a grand system $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$, which is now shown to represent the general framework.

Lemma 34 *Under Ax. 1–6, the above-defined system $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is an expected-utility representation.*

Proof. This property is inherited from the subsystems $(U_\alpha, P_\alpha)_{\alpha \in \Delta}$ ($\Delta \in \Pi$). ■

I now reduce R3 to subframeworks, using another lemma shown in App. F:

Lemma 35 *Given Ax. 1–6 and the above-defined functions P_α , if \mathcal{R} is an algebra as in Ax. 6 and for each Δ -subframework ($\Delta \in \Pi$) ρ_Δ is a fine probability measure on $\mathcal{R}|_{\mathbf{S}_\Delta}$ uncontroversial among $(P_\alpha)_{\alpha \in \Delta}$, then the assignment $A \mapsto \rho_\Delta(A \cap \mathbf{S}_\Delta)$ defines a fine probability measure on \mathcal{R} which does not depend on $\Delta \in \Pi$ and is uncontroversial among $(P_\alpha)_{\alpha \in \Gamma}$.*

Lemma 36 *Under Ax. 1–6, the above-defined measures P_α satisfy R3.*

Proof. Assume Ax. 1–6, with \mathcal{R} as in Ax. 6. By Lem. 35 it suffices to show that for each $\Delta \in \Pi$ there is a fine probability measure ρ_Δ on $\mathcal{R}|_{\mathbf{S}_\Delta}$ uncontroversial among the above-defined (Δ) -family $(P_\alpha)_{\alpha \in \Delta}$. Let $\Delta \in \Pi$. As $(P_\alpha)_{\alpha \in \Delta}$ satisfies R3, some fine measure ρ_Δ is uncontroversial among $(P_\alpha)_{\alpha \in \Delta}$. By Lem. 33's proof, the Δ -subframework satisfies Ax. 6 in virtue of the trace algebra $\mathcal{R}|_{\mathbf{S}_\Delta}$. So by Thm. 1's proof we may w.l.o.g. let ρ_Δ have domain $\mathcal{R}|_{\mathbf{S}_\Delta}$. ■

Lemma 37 *Under Ax. 1–6, the above-defined functions P_α satisfy R2.*

Proof. The proof states literally like that of Lem. 16, the corresponding lemma under exhaustive states. As a tiny addition, α' and β' in Claim 1 must be chosen from Δ_α resp. Δ_β ,⁴⁰ so that Lem. 20 and 21 can be applied to the Δ_α - resp. Δ_β -subframework; both lemmas hadn't been stated for a general framework.⁴¹ ■

I finally prove R1, again using a technical lemma shown in App. F:

Lemma 38 *Under Ax. 1–6, for any contexts $\alpha, \beta \in \Gamma$, algebra \mathcal{R} as in Ax. 6 and \mathcal{R} -measurable functions $f, g : \mathbf{S} \rightarrow X_\alpha \cap X_\beta$, $f_{\mathbf{S}_\alpha} \succsim_\alpha^+ g_{\mathbf{S}_\alpha} \Leftrightarrow f_{\mathbf{S}_\beta} \succsim_\beta^+ g_{\mathbf{S}_\beta}$.*

Lemma 39 *Under Ax. 1–6, the above-defined functions U_α satisfy R1.*

⁴⁰This is possible, because $\mathbf{S}_{\alpha'} = \mathbf{S}_\alpha$ and (by independence between outcome and state awareness) w.l.o.g. $X_{\alpha'} = X_\alpha$, and because $\mathbf{S}_{\beta'} = \mathbf{S}_\beta$ and w.l.o.g. $X_{\beta'} = X_\beta$.

⁴¹Lem. 25 is applied to the general framework for which it had not been stated but still holds.

Proof. Suppose Ax. 1–6. Let $\alpha, \beta \in \Gamma$. Put $X := X_\alpha \cap X_\beta$. Let U_α and U_β be the above-defined functions. W.l.o.g. they are both non-constant on X .⁴² Let \mathcal{R} be as in Ax. 6, and ρ a (by Lem. 36 and its proof existing) fine uncontroversial measure on \mathcal{R} . Let \geq be the relation on $F := \{f \in X^{\mathbf{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$ given by $f \geq g \Leftrightarrow f_{\mathbf{S}_\gamma} \succsim_\gamma^+ g_{\mathbf{S}_\gamma}$ for some (hence by Lem. 38 *any*) $\gamma \in \{\alpha, \beta\}$. To show that $U_\alpha|_X$ is an increasing affine transformations of $U_\beta|_X$, I prove that $(U_\alpha|_X, \rho)$ and $(U_\beta|_X, \rho)$ both represent $(X, (S, R), \geq)$ in Lem. 6's sense. Let $\gamma \in \{\alpha, \beta\}$. I show $\mathbb{E}_\rho(U_\gamma \circ f) \geq \mathbb{E}_\rho(U_\gamma \circ g) \Leftrightarrow f \geq g$ for all $f, g \in F$. As $f \geq_\gamma g$ reduces to $f_{\mathbf{S}_\gamma} \succsim_\gamma^+ g_{\mathbf{S}_\gamma}$, hence to $\mathbb{E}_{P_\gamma^+}(U_\gamma \circ f_{\mathbf{S}_\gamma}) \geq \mathbb{E}_{P_\gamma^+}(U_\gamma \circ g_{\mathbf{S}_\gamma})$, it suffices to show that $\mathbb{E}_\rho(U_\gamma \circ f) = \mathbb{E}_{P_\gamma^+}(U_\gamma \circ f_{\mathbf{S}_\gamma})$ for all $f \in F$. Let $f \in F$; I prove $\rho(f^{-1}(x)) = P_\gamma^+(f_{\mathbf{S}_\gamma}^{-1}(x))$ for all $x \in X$. By Lem. 35 (and Lem. 36's proof), we may write $\rho = \rho_{\Delta_\gamma}(\cdot \cap \mathbf{S}_\gamma)$ for a fine measure ρ_{Δ_γ} on $\mathcal{R}|_{\mathbf{S}_\gamma}$ uncontroversial among $(P_\delta)_{\delta \in \Delta_\gamma}$. Not only ρ_{Δ_γ} , but also $P_\gamma^+|_{\mathcal{R}|_{\mathbf{S}_\gamma}}$ is uncontroversial among $(P_\delta)_{\delta \in \Delta_\gamma}$, by Lem. 4 (applied to the Δ_γ -subframework). So $\rho_{\Delta_\gamma} = P_\gamma^+|_{\mathcal{R}|_{\mathbf{S}_\gamma}}$. For any $x \in X$, $\rho(f^{-1}(x)) = \rho_{\Delta_\gamma}(f^{-1}(x) \cap \mathbf{S}_\gamma) = \rho_{\Delta_\gamma}(f_{\mathbf{S}_\gamma}^{-1}(x)) = P_\gamma^+(f_{\mathbf{S}_\gamma}^{-1}(x))$, where the first equality holds as $\rho = \rho_{\Delta_\gamma}(\cdot \cap \mathbf{S}_\gamma)$, and the last one as $\rho_{\Delta_\gamma} = P_\gamma^+|_{\mathcal{R}|_{\mathbf{S}_\gamma}}$. ■

D.2 Necessity of the axioms and uniqueness

Necessity of Ax. 1–5 holds by the same arguments as under exhaustive states (App. C). I now prove necessity of Ax. 6 and uniqueness of the representation, both by reduction to subframeworks via the following technical lemma (shown in App. F):

Lemma 40 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1's sense (with a fine uncontroversial measure ρ on an algebra \mathcal{R}), then each subsystem $(U_\alpha, P_\alpha)_{\alpha \in \Delta}$ ($\Delta \in \Pi$) represents the Δ -subframework in Thm. 1's sense (with a fine uncontroversial measure ρ_Δ on $\mathcal{R}|_{\mathbf{S}_\Delta}$ given by $\rho(\cdot) = \rho_\Delta(\cdot \cap \mathbf{S}_\Delta)$).*

Lemma 41 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ and $(U'_\alpha, P'_\alpha)_{\alpha \in \Gamma}$ are representations in Thm. 1's sense, then any P_α equals P'_α and any U_α is an increasing affine transformation of U'_α .*

Proof. This property follows via Lem. 40 from the uniqueness property for subframeworks, which is guaranteed by Thm. 1 applied to subframeworks. ■

Lemma 42 *If $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ is a representation in Thm. 1's sense, with a fine uncontroversial measure on an algebra \mathcal{R} , then Ax. 6 holds in virtue of algebra \mathcal{R} .*

Proof. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ and \mathcal{R} be as specified. Let $\alpha \in \Gamma$, $f \succ_\alpha g$ in F_α , and $x \in X_\alpha$. Put $\Delta := \Delta_\alpha$. By Lem. 40, $(U_\alpha, P_\alpha)_{\alpha \in \Delta}$ represents the Δ -subframework, with a fine uncontroversial measure on $\mathcal{R}|_{\mathbf{S}_\Delta}$. So by Lem. 29, Ax. 6 holds for

⁴²The argument is like in the proof of Lemma 15, but using Lem. 34 rather than 14.

this subframework in virtue of algebra $\mathcal{R}|_{\mathbf{S}_\Delta}$. Hence one can partition \mathbf{S}_Δ into $A_1, \dots, A_n \in \mathcal{R}|_{\mathbf{S}_\Delta}$ and pick a $\beta \in \Gamma$ where $S_\beta = S_\alpha \vee \{A_1, \dots, A_n\}$ (so any A_i is representable by an $E_i \subseteq S_\beta$), $X_\beta \supseteq X_\alpha$ (so F_β contains acts f' and g' equivalent to f resp. g), and $f'_{S_\beta \setminus E_i} x_{E_i} \succ_\beta g'$ and $f' \succ_\beta g'_{S_\beta \setminus E_i} x_{E_i}$ for all E_i . Each A_i is in $\mathcal{R}|_{\mathbf{S}_\Delta}$; so $A_i = B_i \cap \mathbf{S}_\Delta$ for a $B_i \in \mathcal{R}$. W.l.o.g. B_1, \dots, B_n partition \mathbf{S} .⁴³ Ax. 6 for the general framework follows since $S_\beta = S_\alpha \vee \{B_1, \dots, B_n\}$ (as $S_\beta = S_\alpha \vee \{A_1, \dots, A_n\}$ and each A_i matches B_i within \mathbf{S}_Δ) and any B_i is, like A_i , represented by E_i . ■

E Proof of Theorem 2

I now reduce Thm. 2 to Thm. 1. The proof is stated so as to be useful also for readers focusing on exhaustive states.

Let a ('risky') algebra \mathcal{R} on \mathbf{S} be given. First assume Ax. 1–5 and $6_{\mathcal{R}}\text{--}8_{\mathcal{R}}$. As Ax. $6_{\mathcal{R}}\text{--}8_{\mathcal{R}}$ imply Ax. 6, Thm. 1's representation $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$ exists. This representation satisfies even Thm. 2's modified third rule, as the uncontroversial measure can be defined on any algebra as in Ax. 6, e.g., on the risky algebra \mathcal{R} , using Lem. 17 (under exhaustive states) or more generally Lem. 36.

Conversely, if preferences admit Thm. 2's representation, then Ax. 1–6 hold by Thm. 1. In fact Ax. 6 holds in virtue of the risky algebra \mathcal{R} , by Lem. 29 (under exhaustive states) or more generally Lem. 42. This implies Ax. $6_{\mathcal{R}}\text{--}8_{\mathcal{R}}$. ■

F Proof of the technical lemmas

Proof of Lem. 1. Assume fine states. Ax. $\tilde{6}$ implies Ax. 6 in virtue of the same \mathcal{R} and the special case $\alpha = \beta$, because incorporability of all $A \in \mathcal{R}$ comes for free (see Rem. 18) and whenever $E_1, \dots, E_n \subseteq S_\alpha$ partition S_α and represent some $A_1, \dots, A_n \in \mathcal{R}$, then we may choose A_1, \dots, A_n such as to partition \mathbf{S} . Conversely, assume Ax. 6 and 2. Pick an algebra \mathcal{R} as in Ax. 6. To show that Ax. $\tilde{6}$ holds in virtue of \mathcal{R} , consider $\alpha \in \Gamma$, acts $f \succ_\alpha g$ in F_α , and an outcome $x \in X_\alpha$. Pick $\beta \in \Gamma$, $E_1, \dots, E_n \subseteq S_\beta$ and $f', g' \in F_\beta$ as given by Ax. 6; so $f'_{S_\beta \setminus E_i} x_{E_i} \succ_\beta g'$ and $f' \succ_\beta g'_{S_\beta \setminus E_i} x_{E_i}$ for $i = 1, \dots, n$. By state fineness, $S_\beta = S_\alpha$; so $E_1, \dots, E_n \subseteq S_\alpha$ and (as also $X_\beta \supseteq X_\alpha$) $f' = f$ and $g' = g$. So, by preference stability (see Prop. 2, which uses Ax. 2), $f_{S_\alpha \setminus E_i} x_{E_i} \succ_\alpha g$ and $f \succ_\alpha g_{S_\alpha \setminus E_i} x_{E_i}$ for $i = 1, \dots, n$. ■

Proof of Lem. 2. Assume fine states. First, fineness of the commonality implies R3 since the commonality is uncontroversial. Conversely, assume there is a fine uncontroversial ρ . As states are fine, all objective events are representable in all contexts. So the commonality extends ρ , hence it itself fine. ■

⁴³Otherwise replace each B_i by $B_i \setminus \bigcup_{j=1}^{i-1} B_j$ if $i < n$ and by $(B_i \setminus \bigcup_{j=1}^{i-1} B_j) \cup (\overline{\bigcup_{j=1}^n B_j})$ if $i = n$, which yields sets in \mathcal{R} that are exclusive (by the ' $\setminus \bigcup_{j=1}^{i-1} B_j$ ') and exhaustive (by the ' $\cup (\overline{\bigcup_{j=1}^n B_j})$ ').

Lem. 3 and 4 are provable analogously to Lem. 1 resp. 2.

We now turn to App. C's technical lemmas. Let states be exhaustive until otherwise stated.

Proof of Lem. 5. Let $\alpha \in \Gamma$. Let \mathcal{R}_1 and \mathcal{R}_2 be the sets in (1) resp. (2). Since \mathcal{R}_1 is obviously an algebra, it suffices to show that $\mathcal{E}_\alpha = \mathcal{R}_1 = \mathcal{R}_2$.

Claim 1: $\mathcal{R}_1 \subseteq \mathcal{R}_2$. Note that \mathcal{R}_2 includes $(2^{S_\alpha})^*$ as S_α harmlessly refines S_α ; and \mathcal{R}_2 also includes \mathcal{I} as each $I \in \mathcal{I}$ is by definition representable in some harmless refinement S_β of S_α , meaning that $I \in (2^{S_\alpha})^*$. Hence \mathcal{R}_2 includes the join $\mathcal{R}_1 = (2^{S_\alpha})^* \vee \mathcal{I}$.

Claim 2: $\mathcal{R}_2 \subseteq \mathcal{E}_\alpha$. Let $E \in \mathcal{R}_2$. Then we may pick a context $\beta \in \Gamma$ such that S_β harmlessly refines S_α and $E \in (2^{S_\beta})^*$. So $E = A^*$ for some $A \in 2^{S_\beta}$, i.e., some $A \subseteq S_\beta$. Hence, $E \in \mathcal{E}_\alpha$.

Claim 3: $\mathcal{E}_\alpha \subseteq \mathcal{R}_1$. Let $E \in \mathcal{E}_\alpha$. Then we may pick a finite partition $S \subseteq \mathcal{I}$ of \mathbf{S} such that $E = A^*$ where $A \subseteq S_\alpha \vee \mathcal{P}$. Note that E can be represented as

$$E = \cup_{I \in \mathcal{P}} \cup_{s \in S_\alpha: s \cap I \in A} (s \cap I) = \cup_{I \in \mathcal{P}} (I \cap (\cup_{s \in S_\alpha: s \cap I \in A} s)).$$

So E is a Boolean combination of members of \mathcal{I} and $(2^{S_\alpha})^*$, showing that $E \in \mathcal{R}_1$. \blacksquare

Lemma 43 *Given any finite set $\mathcal{J} \subseteq \mathcal{I}$, there is a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbf{S} refining each $\{J, \bar{J}\}$ ($J \in \mathcal{J}$) such that for all contexts $\alpha \in \Gamma$ there is a $\beta \in \Gamma$ for which $S_\beta = S_\alpha \vee \mathcal{P}$ and \lesssim_β is faithful to \lesssim_α .*

Proof of Lem. 43. This can be shown by induction on the size of \mathcal{J} . The claim holds trivially if $\mathcal{J} = \emptyset$, namely in virtue of the partition $\mathcal{P} = \{\mathbf{S}\}$. Now assume the claim holds for some sets $\mathcal{J}_1, \mathcal{J}_2 \subseteq \mathcal{I}$, say in virtue of partitions \mathcal{P}_1 resp. \mathcal{P}_2 . Then the claim also holds for $\mathcal{J}_1 \cup \mathcal{J}_2$, namely in virtue of the partition $\mathcal{J}_1 \vee \mathcal{J}_2$, because for any $\alpha \in \Gamma$ we may first pick a context $\beta' \in \Gamma$ such that $S_{\beta'} = S_\alpha \vee \mathcal{P}_1$ and $\lesssim_{\beta'}$ is faithful to \lesssim_α , and then pick a context $\beta \in \Gamma$ such that $S_\beta = S_{\beta'} \vee \mathcal{P}_2 = S_\alpha \vee \mathcal{P}$ and \lesssim_β is faithful to $\lesssim_{\beta'}$, hence to \lesssim_α . \blacksquare

Proof of Lem. 8. Assume Ax. 2, and let $\alpha, \beta \in \Gamma$ such that $X_\alpha \subseteq X_\beta$ and $S_\beta = S_\alpha \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbf{S} . Using Lem. 43, pick a finite partition $\mathcal{P}' \subseteq \mathcal{I}$ of \mathbf{S} refining \mathcal{P} such that there are $\alpha', \beta' \in \Gamma$ where $S_{\alpha'} = S_\alpha \vee \mathcal{P}'$, $S_{\beta'} = S_\beta \vee \mathcal{P}'$, $\lesssim_{\alpha'}$ is faithful to \lesssim_α , and $\lesssim_{\beta'}$ is faithful to \lesssim_β . Now $S_{\alpha'} = S_{\beta'}$ (as $S_\beta = S_\alpha \vee \mathcal{P}$) and $X_{\alpha'}, X_{\beta'} \supseteq X_\alpha$ (as by faithfulness $X_{\alpha'} \supseteq X_\alpha$ and $X_{\beta'} \supseteq X_\beta$, and as $X_\beta \supseteq X_\alpha$). So $F_{\alpha'} \cap F_{\beta'} \supseteq X_{\alpha'}^{S_{\alpha'}}$. Hence, by Ax. 2, $\lesssim_{\beta'}$ matches $\lesssim_{\alpha'}$ on $X_{\alpha'}^{S_{\alpha'}}$, hence is (like $\lesssim_{\alpha'}$) faithful to \lesssim_α . As $\lesssim_{\beta'}$ is faithful to \lesssim_α and \lesssim_β (and as each $f \in F_\alpha$ is objectively equivalent to some $g \in F_\beta$), \lesssim_β is faithful to \lesssim_α . \blacksquare

Proof of Lem. 9. Assume Ax. 2 and $f \lesssim_\alpha^+ g$, where $\alpha \in \Gamma$ and $f, g \in X_\alpha^{\mathbf{S}}$.

(a) Let $\beta \in \Gamma$ satisfy the conditions (i)–(ii) in Def. 28. I show that $f_\beta \lesssim_\beta g_\beta$. As $f \lesssim_\alpha^+ g$, we have $f_{\beta'} \lesssim_{\beta'} g_{\beta'}$ for some $\beta' \in \Gamma$ satisfying these conditions. As S_β

and $S_{\beta'}$ harmlessly refine S_α , we may pick finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{I}$ of \mathbf{S} such that $S_\beta = S_\alpha \vee \mathcal{P}$ and $S_{\beta'} = S_\alpha \vee \mathcal{P}'$. Using Lem. 43, there is a partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbf{S} which refines \mathcal{P} and \mathcal{P}' and contexts $\gamma, \gamma' \in \Gamma$ such that $S_\gamma = S_\beta \vee \mathcal{Q}$, $S_{\gamma'} = S_{\beta'} \vee \mathcal{Q}$, \succsim_γ is faithful to \succsim_β , and $\succsim_{\gamma'}$ to $\succsim_{\beta'}$. Note that $S_\gamma = S_{\gamma'} = S_\alpha \vee \mathcal{Q}$, so that $f_\gamma = f_{\gamma'}$ and $g_\gamma = g_{\gamma'}$. Hence, by preference stability (Prop. 2), $f_\gamma \succsim_\gamma g_\gamma \Leftrightarrow f_{\gamma'} \succsim_{\gamma'} g_{\gamma'}$. This equivalence reduces to $f_\beta \succsim_\beta g_\beta \Leftrightarrow f_{\beta'} \succsim_{\beta'} g_{\beta'}$ by faithfulness of \succsim_γ to \succsim_β and of $\succsim_{\gamma'}$ to $\succsim_{\beta'}$. As $f_{\beta'} \succsim_{\beta'} g_{\beta'}$, it follows that $f_\beta \succsim_\beta g_\beta$.

(b) Pick any $\beta \in \Gamma$ satisfying (i)–(ii) in Def. 28 hold. Pick a context $\beta' \in \Gamma$ such that $S_{\beta'} = S_\beta$ and $X_{\beta'} = X_\alpha$. Clearly, also $\beta' \in \Gamma$ satisfies (i)–(ii) in Def. 28. Moreover, $\succsim_{\beta'}$ is faithful to \succsim_α by Lem. 8. ■

Proof of Lem. 10. Consider an $\alpha \in \Gamma$ and a finite $\mathcal{B} \subseteq \mathcal{E}_\alpha$. For each $B \in \mathcal{B}$, pick a partition \mathcal{P}_B of \mathbf{S} refining $\{B, \overline{B}\}$ and having the property stated in the definition of weak incorporability (note that $\mathcal{P}_B \subseteq \mathcal{I}$). Let B_1, \dots, B_n be all $n = |\mathcal{B}|$ members of \mathcal{B} in any given order. We may pick, first, a context $\beta_1 \in \Gamma$ such that $S_{\beta_1} = S_\alpha \vee \mathcal{P}_{B_1}$ and \succsim_{β_1} is faithful to \succsim_α ; second, a context $\beta_2 \in \Gamma$ such that $S_{\beta_2} = S_{\beta_1} \vee \mathcal{P}_{B_2}$ and \succsim_{β_2} is faithful to \succsim_{β_1} ; and so on for contexts β_3, \dots, β_n . Let $\beta := \beta_n$. Property (i) holds because each B_i is representable in context β_i , hence in context β . Property (ii) holds as $S_\beta = S_\alpha \vee \mathcal{P}$ with $\mathcal{P} := \mathcal{P}_{B_1} \vee \dots \vee \mathcal{P}_{B_n}$. Property (iii) holds by transitivity of faithfulness. ■

Proof of Lem. 11. This claim follows from Lem. 10 applied to the (finite) set $\mathcal{B} = \{f^{-1}(x) : f \in \mathcal{G}, x \in X_\alpha\}$, by noting that for any $\beta \in \Gamma$ F_β^* is characterizable as the set of $(2^{S_\beta})^*$ -measurable function from \mathbf{S} to X_α . ■

Proof of Lem. 12. Assume Ax. 2 and 5, let $\alpha \in \Gamma$ and consider Lem. 7's generalized Savage framework, with set of acts denoted F . Let f, g, A be as specified. First assume $f \succsim_{\alpha, A}^+ g$. Then, by definition, $f' \succsim_\alpha^+ g'$ for some $f', g' \in F$ agreeing with f resp. g on A and with each other outside A . Choose any $h \in F$ taking one value on A and another on \overline{A} (it exists as $|X_\alpha| \geq 2$ by Ax. 5). Using Lem. 11, we pick a $\beta \in \Gamma$ such that $f, g, f', g', h \in F_\beta^*$ and S_β harmlessly refines S_α (and \succsim_β is faithful to \succsim_α , which is only needed if the modified equivalence is to be proved). As $h \in F_\beta^*$, A is representable. As $f' \succsim_\alpha^+ g'$, we have $f'_\beta \succsim_\beta g'_\beta$ by Lem. 9. Noting that f'_β and g'_β agree with f_β resp. g_β on A_β and with each other outside A_β (because of inheriting these properties from analogous properties of f' and g'), it follows that $f_\beta \succsim_{\beta, A} g_\beta$.

Conversely, assume that $f_\beta \succsim_{\beta, A_\beta} g_\beta$ for some $\beta \in \Gamma$ satisfying the specified properties. Then there are two functions in F_β – we may write them as f'_β and g'_β for certain $f, g \in F_\beta^*$ – such that $f'_\beta \succsim_\beta g'_\beta$ and such that f'_β and g'_β agree with f_β resp. g_β on A_β and with each other outside A_β . From $f'_\beta \succsim_\beta g'_\beta$ (and the properties of β) it follows that $f' \succsim_\alpha^+ g'$, which in turn implies that $f \succsim_{\alpha, A}^+ g$ since f' and g' agree with f resp. g on A and with each other outside A (they inherit this

behaviour from f_β and g_β because $f = (f_\beta)^*$, $g = (g_\beta)^*$ and $A = (A_\beta)^*$. ■

Proof of Lem. 13. Assume Ax. 2 and let $\alpha \in \Gamma$. Consider Lem. 7's generalized Savage framework and an event $A \in \mathcal{E}$.

First assume A is non-null. Then there are $f, g \in F$ such that $f_{\bar{A}} = g_{\bar{A}}$ and $f \not\sim_\alpha^+ g$. Pick any $h \in F$ taking one value on A and another on \bar{A} (h exists as $|X_\alpha| \geq 2$ by the fact that F contains distinct functions f, g). By Lem. 11, we may choose a context $\beta \in \Gamma$ such that $f, g, h \in F_\beta^*$ and S_β faithfully refines S_α (and such that \succsim_β is faithful to \succsim_α , something we need to add when proving the equivalence in its modified version). As $h \in F_\beta^*$, A is representable in context β , i.e., A_β is defined. As $f \not\sim_\alpha^+ g$, we have $f_\beta \not\sim_\beta g_\beta$, which (since f_β and g_β agree outside A_β) shows that A_β is non-null.

Conversely, assume A_β is non-null (under \succsim_β) for some $\beta \in \Gamma$ with the specified properties. Then we may pick two non-indifferent acts in F_β which agree outside A ; we may write them as f_β and g_β for some $f, g \in F_\beta^*$. Since $f_\beta \not\sim_\beta g_\beta$, we have $f \not\sim_\alpha^+ g$ by Lem. 9. So, as f and g agree outside A , A is non-null. ■

Proof of Lem. 20. Assume Ax. 2, 4 and 5. Let $\alpha \in \Gamma$ and $A, B \subseteq \mathbf{S}$. By Ax. 5 there are $x \succ_\alpha y$ in X_α .

First assume $A \succsim_\alpha^+ B$. Then there exist $x, y \in X_\alpha$ such that $x \succ_\alpha^+ y$ and $x_A y_{\bar{A}} \succsim_\alpha^+ x_B y_{\bar{B}}$. So by Lem. 11 there is a context $\beta \in \Gamma$ such that $x_A y_{\bar{A}}, x_B y_{\bar{B}} \in F_\beta^*$ (hence, A and B are representable), S_β harmlessly refines S_α , and \succsim_β is faithful to \succsim_α (the latter is needed when proving the modified equivalence). By Lem. 9, it follows that $x_{S_\beta} \succ_\beta y_{S_\beta}$ and $(x_A y_{\bar{A}})_\beta \succsim_\beta (x_B y_{\bar{B}})_\beta$. In other words $x \succ_\beta y$ and $x_{A_\beta} y_{S_\beta \setminus A_\beta} \succsim_\beta x_{B_\beta} y_{S_\beta \setminus B_\beta}$. So, $A_\beta \succsim_\beta B_\beta$.

Conversely, assume $A_\beta \succsim_\beta B_\beta$ for a $\beta \in \Gamma$ such that A and B are representable and S_β harmlessly refines S_α . It follows that $A, B \in \mathcal{E}_\alpha$. So by Lem. 10 we may pick a context $\beta' \in \Gamma$ such that A and B are representable, $S_{\beta'}$ harmlessly refines S_α , and $\succsim_{\beta'}$ is faithful to \succsim_α . In particular, $X_{\beta'} = X_\alpha$. As $A_\beta \succsim_\beta B_\beta$ we have $A_{\beta'} \succsim_{\beta'} B_{\beta'}$ by belief stability (see Prop. 4, which uses Ax. 2, 4 and 5). Hence there are $x', y' \in X_{\beta'} (= X_\alpha)$ such that $x' \succ_{\beta'} y'$ and $x'_{A_{\beta'}} y'_{S_{\beta'} \setminus A_{\beta'}} \succsim_{\beta'} x'_{B_{\beta'}} y'_{S_{\beta'} \setminus B_{\beta'}}$. In other words, $(x'_S)_{\beta'} \succ_{\beta'} (y'_S)_{\beta'}$ and $(x'_A y'_{\bar{A}})_{\beta'} \succsim_{\beta'} (x'_B y'_{\bar{B}})_{\beta'}$. By Lem. 9 it follows that $x'_S \succ_\alpha^+ y'_S$ (i.e., $x' \succ_\alpha^+ y'$) and $x'_A y'_{\bar{A}} \succsim_\alpha^+ x'_B y'_{\bar{B}}$. So $A \succsim_\alpha^+ B$. ■

Proof of Lem. 21. Assume Ax. 2, 4, 5 and 6. Let $A \succsim_\alpha^+ B$, where $\alpha \in \Gamma$ and $A, B \subseteq \mathbf{S}$. Let $\beta \in \Gamma$ satisfy the conditions stated. I show that $A_\beta \succsim_\beta B_\beta$. As $A \succsim_\alpha^+ B$, we have $A_{\beta'} \succsim_{\beta'} B_{\beta'}$ for some $\beta' \in \Gamma$ satisfying the analogous conditions, by Lem. 20 (which uses Ax. 2, 4 and 5). As S_β and $S_{\beta'}$ harmlessly refine S_α , we may pick finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{I}$ of \mathbf{S} such that $S_\beta = S_\alpha \vee \mathcal{P}$ and $S_{\beta'} = S_\alpha \vee \mathcal{P}'$. Using Lem. 43, there is a partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbf{S} which refines \mathcal{P} and \mathcal{P}' and contexts $\gamma, \gamma' \in \Gamma$ such that $S_\gamma = S_\beta \vee \mathcal{Q}$, $S_{\gamma'} = S_{\beta'} \vee \mathcal{Q}$, \succsim_γ is faithful to \succsim_β , and $\succsim_{\gamma'}$ to $\succsim_{\beta'}$. Note that $S_\gamma = S_{\gamma'} = S_\alpha \vee \mathcal{Q}$, so that $A_\gamma = A_{\gamma'}$ and $B_\gamma = B_{\gamma'}$. So, by comparative-

belief stability (Prop. 4, which uses Ax. 2, 4, 5 and 6), $A_\gamma \succsim_\gamma B_\gamma \Leftrightarrow A_{\gamma'} \succsim_{\gamma'} B_{\gamma'}$. This equivalence reduces to $A_\beta \succsim_\beta B_\beta \Leftrightarrow A_{\beta'} \succsim_{\beta'} B_{\beta'}$ by faithfulness of \succsim_γ to \succsim_β and of $\succsim_{\gamma'}$ to \succsim_γ . As $A_{\beta'} \succsim_{\beta'} B_{\beta'}$, it follows that $A_\beta \succsim_\beta B_\beta$. ■

Proof of Lem. 22. Assume Ax. 2, 4 and 5. Let \mathcal{R} be a robust algebra of incorporable objective events, and let $A, B \in \mathcal{R}$ and $\alpha, \beta \in \Gamma$. I assume $A \succsim_\alpha^+ B$ and have to prove $A \succsim_\beta^+ B$. By Lem. 22, as $A \succsim_\alpha^+ B$ we have $A_\gamma \succsim_\gamma B_\gamma$ for a $\gamma \in \Gamma$ such that A and B are representable and S_γ harmlessly refines S_α . Meanwhile, as $A, B \in \mathcal{R} \subseteq \mathcal{I} \subseteq \mathcal{E}_\beta$, by Lem. 10 there exists a $\delta \in \Gamma$ such that A and B are representable in context δ and S_δ harmlessly refines S_β . As $A_\gamma \succsim_\gamma B_\gamma$ and as A and B belong to a robust algebra (i.e., \mathcal{R}), we have $A_\delta \succsim_\delta B_\delta$ by belief stability on robust algebras (Prop. 4). So $A \succsim_\beta^+ B$ by Lem. 20. ■

Proof of Lem. 23. Assume Ax. 1–6. Let \mathcal{R} be as in Ax. 6. Let U_α and P_α^+ ($\alpha \in \Gamma$) be as defined above.

Claim 1: $P_\alpha^+|_{\mathcal{R}}$ is fine for all $\alpha \in \Gamma$. Let $\alpha \in \Gamma$. The pair (U_α, P_α^+) represents \succsim_α^+ on $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{E}_\alpha\text{-measurable}\}$ (in Lem. 6's sense). So $(U_\alpha, P_\alpha^+|_{\mathcal{R}})$ represents \succsim_α^+ on $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$. By Lem. 7 (applied with $\mathcal{E} = \mathcal{R}$), there is a fine probability measure on \mathcal{R} representing the (belief) relation induced by \succsim_α^+ on \mathcal{R} . This measure represents the same (belief) relation on \mathcal{R} as $P_\alpha^+|_{\mathcal{R}}$, and thus coincides with $P_\alpha^+|_{\mathcal{R}}$ by Lem. 18 and 19. So $P_\alpha^+|_{\mathcal{R}}$ is fine.

Claim 2: $\rho := P_\alpha^+|_{\mathcal{R}}$ is the same for all $\alpha \in \Gamma$. Let $\alpha, \beta \in \Gamma$. By Lem. 22, the functions $P_\alpha^+|_{\mathcal{R}}$ and $P_\beta^+|_{\mathcal{R}}$ are ordinally equivalent. Since these are fine probability measures by Claim 1, they must coincide by Lem. 18 and 19. ■

Lemma 44 Under Ax. 1, for any context $\alpha \in \Gamma$, two functions $f, g \in X_\alpha^{\mathbf{S}}$ are \succsim_α^+ -comparable (i.e., $f \succsim_\alpha^+ g$ or $g \succsim_\alpha^+ f$) if and only if both are \mathcal{E}_α -measurable.

Proof. Assume Ax. 1. Let $\alpha \in \Gamma$ and $f, g \in X_\alpha^{\mathbf{S}}$. First assume f and g are comparable under \succsim_α^+ . Then f_β and g_β are comparable for some context $\beta \in \Gamma$ such that $f, g \in F_\beta^*$ and $S_\beta = S_\alpha \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbf{S} . Since $f, g \in F_\beta^*$, f and g are $(2^{S_\beta})^*$ -measurable, which implies \mathcal{E}_α -measurability as $(2^{S_\beta})^* = (2^{S_\alpha \vee \mathcal{P}})^* \subseteq \mathcal{E}_\alpha$. Conversely, if f and g are \mathcal{E}_α -measurable, then by Lem. 11 there is a context $\beta \in \Gamma$ such that $f, g \in F_\beta^*$ and S_β harmlessly refines S_α . By Ax. 1, $f_\beta \succsim_\beta g_\beta$ or $g_\beta \succsim_\beta f_\beta$, which implies that $f \succsim_\alpha^+ g$ or $g \succsim_\alpha^+ f$. ■

Proof of Lem. 24. Assume Ax. 1 and 2. Let $\alpha, \beta \in \Gamma$. Assume S_β harmlessly refines S_α . Then $\mathcal{E}_\alpha = \mathcal{E}_\beta$ by definition of extrapolated algebras. Now suppose that in addition \succsim_β is faithful to \succsim_α . In view of Lem. 44 it suffices to show that \succsim_α^+ and \succsim_β^+ coincide on the set of \mathcal{E}_β - (resp. \mathcal{E}_α -)measurable functions in $X_\alpha^{\mathbf{S}}$. Let $f, g \in X_\alpha^{\mathbf{S}}$ be \mathcal{E}_β - (hence, \mathcal{E}_α -)measurable. Then by Lem. 11 there is a context $\gamma \in \Gamma$ such that $f, g \in F_\gamma^*$ and S_γ harmlessly refines S_β , hence, also S_α . We have $f \succsim_\alpha^+ g \Leftrightarrow f \succsim_\beta^+ g$ because each side is equivalent to $f_\gamma \succsim_\gamma g_\gamma$ by Lem. 9. ■

Proof of Lem. 25. Assume Ax. 2. Let $\alpha, \beta \in \Gamma$ and $A \subseteq S_\alpha \cap S_\beta$. We assume A is non-null in α and prove non-nullness in β . By assumption, there exist $f, g \in F_\alpha$ such that $f_{S_\alpha \setminus A} = g_{S_\alpha \setminus A}$ and $f \not\sim_\alpha g$. Pick any $f', g' \in F_\beta$ such that $f_A = f'_A$, $g_A = g'_A$, and $f'_{S_\beta \setminus A} = g'_{S_\beta \setminus A}$. As $f \not\sim_\alpha g$ we have $f' \sim_\beta g'$ by Ax. 2. So A is non-null in β . ■

Proof of Lem. 30. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$, ρ , \mathcal{R} and \mathcal{R}_α be as specified. Fix $\alpha \in \Gamma$.

Claim 1: $\mathcal{R}_\alpha = \cup_{\beta \in \Gamma: S_\beta = S_\alpha \vee \mathcal{P}} \text{ for some finite partition } \mathcal{P} \subseteq \mathcal{R} \text{ of } \mathbf{S} (2^{S_\beta})^*$. This claim is provable analogously to the proof of Lem. 5.

Claim 2: For all $\beta \in \Gamma$ and finite partitions $\mathcal{P} \subseteq \mathcal{R}$ of \mathbf{S} , there is a $\gamma \in \Gamma$ such that $S_\gamma = S_\beta \vee \mathcal{P}$ and P_γ^* extends P_β^* . Consider such β and \mathcal{P} . Write $\mathcal{P} = \{I_1, \dots, I_n\}$. As each I_i is incorporable and ρ is uncontroversial, we can let $\beta_0 := \beta$ and successively pick $\beta_1, \dots, \beta_n \in \Gamma$ such that, for each β_i , $P_{\beta_i}^*$ extends $P_{\beta_{i-1}}^*$ and $S_{\beta_i} = S_{\beta_{i-1}} \vee \{I_i, \bar{I}_i\}$. Clearly, $P_{\beta_n}^*$ extends P_β^* and $S_{\beta_n} = S_\beta \vee \{I_1, \bar{I}_1\} \vee \dots \vee \{I_n, \bar{I}_n\} = S_\beta \vee \mathcal{P}$.

Claim 3: The measures P_β^* with $S_\beta = S_\alpha \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbf{S} agree pairwise on the domain overlap. Let $\beta, \beta' \in \Gamma$ such that $S_\beta = S_\alpha \vee \mathcal{P}$ and $S_{\beta'} = S_\alpha \vee \mathcal{P}'$ for finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{R}$ of \mathbf{S} . I show that P_β^* and $P_{\beta'}^*$ agree on the domain overlap. By Claim 2, there are $\gamma, \gamma' \in \Gamma$ such that $S_\gamma = S_\beta \vee \mathcal{P}'$, $S_{\gamma'} = S_{\beta'} \vee \mathcal{P}$, P_γ^* extends P_β^* , and $P_{\gamma'}^*$ extends $P_{\beta'}^*$. It suffices to show that $P_\gamma^* = P_{\gamma'}^*$. As P_γ and $P_{\gamma'}$ have the same domain $2^{S_\gamma} = 2^{S_{\gamma'}} (= 2^{S_\alpha \vee \mathcal{P} \vee \mathcal{P}'})$, $P_\gamma = P_{\gamma'}$ by R2, whence $P_\gamma^* = P_{\gamma'}^*$.

Claim 4: All desired properties are met by the function ρ_α which to each $A \in \mathcal{R}_\alpha$ assigns $P_\beta^*(A)$ for a (by Claim 1 existing and by Claim 3 arbitrary) $\beta \in \Gamma$ such that $S_\beta = S_\alpha \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbf{S} . By definition, ρ_α extends all P_β^* such that $S_\beta = S_\alpha \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbf{S} . It remains to show that ρ_α is a probability measure. Clearly, $\rho_\alpha(\mathbf{S}) = P_\alpha^*(\mathbf{S}) = 1$. Now consider disjoint $A, B \in \mathcal{R}_\alpha$. By Claim 1 we may pick $\beta, \gamma \in \Gamma$ such that $A \in (2^{S_\beta})^*$, $B \in (2^{S_\gamma})^*$, $S_\beta = S_\alpha \vee \mathcal{P}$ and $S_\gamma = S_\alpha \vee \mathcal{Q}$, for finite partitions $\mathcal{P}, \mathcal{Q} \subseteq \mathcal{R}$ of \mathbf{S} . By Claim 2 we may pick a $\delta \in \Gamma$ such that $S_\delta = S_\alpha \vee \mathcal{P} \vee \mathcal{Q}$. Now $A, B \in (2^{S_\delta})^*$ and $\rho_\alpha(A) + \rho_\alpha(B) = P_\delta^*(A) + P_\delta^*(B) = P_\delta^*(A \cup B) = \rho_\alpha(A \cup B)$. ■

Proof of Lem. 31. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$, ρ , \mathcal{R} , ρ_α and \mathcal{R}_α be as specified. Fix $\alpha \in \Gamma$. The proof proceeds in two steps.

Claim 1: $\mathbb{E}_{\rho_\alpha}(U_\alpha \circ f) \geq \mathbb{E}_{\rho_\alpha}(U_\alpha \circ g) \Leftrightarrow f \succeq_\alpha^+ g$ for all \mathcal{R}_α -measurable $f, g \in X_\alpha^\mathbf{S}$. Let $f, g \in X_\alpha^\mathbf{S}$ be \mathcal{R}_α -measurable. We may pick a finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbf{S} such that f and g are $(2^{S_\alpha \vee \mathcal{P}})^*$ -measurable, and then pick a $\gamma \in \Gamma$ such that $S_\gamma = S_\alpha \vee \mathcal{P}$ (for details see Claims 1 and 2 in Lem. 30's proof). W.l.o.g. $X_\gamma = X_\alpha$ by independence of outcome and state awareness. The desired equivalence holds as $\mathbb{E}_{\rho_\alpha}(U_\alpha \circ f) \geq \mathbb{E}_{\rho_\alpha}(U_\alpha \circ g) \Leftrightarrow \mathbb{E}_{P_\gamma}(U_\gamma \circ f_\gamma) \geq \mathbb{E}_{P_\gamma}(U_\gamma \circ g_\gamma) \Leftrightarrow f_\gamma \succeq_\gamma g_\gamma \Leftrightarrow f \succeq_\alpha^+ g$, where the last ' \Leftrightarrow ' holds by Lem. 9 and the first ' \Leftrightarrow ' holds as ρ_α extends P_γ^* and U_γ is an increasing affine transformation of U_α (by R1 and the fact that $X_\gamma = X_\alpha$).

Claim 2: ρ_α is fine and U_α is non-constant. Non-constancy of U_α holds as U_α is part of representation in Thm. 1's sense. Further, as $\mathcal{R} \subseteq \mathcal{R}_\alpha \subseteq \mathcal{E}_\alpha$ where by Lem. 29 \mathcal{R} is an algebra as in Ax. 6 (and \mathcal{E}_α is the extrapolated algebra), we know by Lem. 6 that the restriction of \succsim_α^+ to $\{f \in X_\alpha^{\mathbf{S}} : f \text{ is } \mathcal{R}_\alpha\text{-measurable}\}$ has a representation (U'_α, P'_α) in Lem. 6's sense; in particular, P'_α is a fine probability measure on \mathcal{R}_α . By Claim 1, ρ represents the same probability order on \mathcal{R}_α as P'_α . Hence $\rho = P'_\alpha$ by Lem. 18 and 19. So ρ is itself fine. ■

From now on the restriction to exhaustive states is lifted.

Lemma 45 *If an algebra \mathcal{R} on \mathbf{S} is robust, then w.r.t. any Δ -subframework ($\Delta \in \Pi$) the (trace) algebra $\mathcal{R}|_{\mathbf{S}_\Delta}$ on \mathbf{S}_Δ is robust.*

Proof. Consider a robust algebra \mathcal{R} on \mathbf{S} , a $\Delta \in \Pi$, contexts $\alpha, \beta \in \Delta$, and $\mathcal{R}|_{\mathbf{S}_\Delta}$ -determined acts $f, g \in F_\alpha$ and $f', g' \in F_\beta$ such that f is equivalent to f' , and g to g' . We must show that $f \succsim_\alpha g \Leftrightarrow f' \succsim_\beta g'$. This holds because (i) \mathcal{R} is robust, and (ii) the $\mathcal{R}|_{\mathbf{S}_\Delta}$ -determinedness of the four acts implies (in fact, is equivalent to) their \mathcal{R} -determinedness. ■

Lemma 46 *Assume Ax. 2. If an objective event $A \subseteq \mathbf{S}$ is incorporable, then w.r.t. any Δ -subframework ($\Delta \in \Pi$) $A \cap \mathbf{S}_\Delta$ is incorporable.*

Proof. Let $A \subseteq \mathbf{S}$ be incorporable w.r.t. $(X_\alpha, S_\alpha, \succsim_\alpha)_{\alpha \in \Gamma}$ and let $\Delta \in \Pi$. Let $\alpha \in \Delta$. By A 's incorporability, there is a context $\beta \in \Gamma$ (perhaps not in Δ) such that $S_\beta = S_\alpha \vee \{A, \bar{A}\}$ and \succsim_β is faithful to \succsim_α . By independence between outcome and state awareness, we can pick a context $\gamma \in \Gamma$ such that $X_\gamma = X_\alpha$ and $S_\gamma = S_\beta$. As $X_\gamma = X_\alpha$ and as $\mathbf{S}_\gamma = \mathbf{S}_\beta = \mathbf{S}_\alpha$ (the last identity holds because S_β refines S_α), we have $\gamma \in \Delta$. So it remains to show two things:

- $S_\gamma = S_\alpha \vee \{A \cap \mathbf{S}_\Delta, \mathbf{S}_\Delta \setminus (A \cap \mathbf{S}_\Delta)\}$: this holds because

$$S_\gamma = S_\beta = S_\alpha \vee \{A, \bar{A}\} = S_\alpha \vee \{A \cap \mathbf{S}_\Delta, \mathbf{S}_\Delta \setminus (A \cap \mathbf{S}_\Delta)\}.$$

- \succsim_γ is faithful to \succsim_α : As \succsim_β is faithful to \succsim_α , $X_\beta \supseteq X_\alpha$, i.e., $X_\beta \supseteq X_\gamma$. So, as also $S_\beta = S_\gamma$, the relation \succsim_γ is the restriction of \succsim_β to $F_\gamma (\subseteq F_\beta)$ by preference stability (see Prop. 2, which uses Ax. 2). Hence, not only \succsim_β , but also \succsim_γ is faithful to \succsim_α . ■

Proof of Lem. 33. Let $\Delta \in \Pi$. The Δ -subframework trivially inherits the first five axioms. We now show that also Ax. 6 is inherited, given Ax. 2. Assume Ax. 2 and 6. Pick an algebra \mathcal{R} on \mathbf{S} as in Ax. 6 (for the general framework). I show that the subframework satisfies Ax. 6 in virtue of the trace algebra $\mathcal{R}|_{\mathbf{S}_\Delta}$. By Lem. 45 and 46, $\mathcal{R}|_{\mathbf{S}_\Delta}$ is, w.r.t. the subframework, a robust algebra (on \mathbf{S}_Δ) composed of incorporable objective events. Now consider an $\alpha \in \Delta$, acts $f \succ_\alpha g$ in F_α , and an $x \in X_\alpha$. By Ax. 6 for the general framework, we may partition \mathbf{S} into some

$A_1, \dots, A_n \in \mathcal{R}$ such that, in some context $\beta \in \Gamma$ where $S_\beta = S_\alpha \vee \{A_1, \dots, A_n\}$ (so each A_i is representable by an $E_i \subseteq S_\beta$) and $X_\beta \supseteq X_\alpha$ (so F_β contains acts f' and g' equivalent to f resp. g), we have $f'_{S_\beta \setminus E_i} x_{E_i} \succ_\beta g'$ and $f' \succ_\beta g'_{S_\beta \setminus E_i} x_{E_i}$ for $i = 1, \dots, n$. To complete the proof of Ax. 6 for the subframework, it suffices to note that (i) $\beta \in \Delta$ because $\mathbf{S}_\beta = \mathbf{S}_\alpha$ (as $S_\beta = S_\alpha \vee \{A_1, \dots, A_n\}$), and (ii) \mathbf{S}_Δ is partitioned into (the non-empty sets among) $A_1 \cap \mathbf{S}_\Delta, \dots, A_n \cap \mathbf{S}_\Delta \in \mathcal{R}|_{\mathbf{S}_\Delta}$, where each such $A_i \cap \mathbf{S}_\Delta$ is represented by E_i . ■

Proof of Lem. 35. Assume Ax. 1–6. Let \mathcal{R} , $(P_\alpha)_{\alpha \in \Gamma}$, and $(\rho_\Delta)_{\Delta \in \Pi}$ be as specified. Each ρ_Δ induces a function π_Δ on \mathcal{R} via $\pi_\Delta(A) := \rho_\Delta(A \cap \mathbf{S}_\Delta)$ ($A \in \mathcal{R}$).

Claim 1: Each π_Δ ($\Delta \in \Pi$) is a fine probability measure. Let $\Delta \in \Pi$. First, π_Δ is a probability measure as ρ_Δ is one, or more precisely, as $\pi_\Delta(\mathbf{S}) = \rho_\Delta(\mathbf{S}_\Delta) = 1$ and as for disjoint $A, B \in \mathcal{R}$ we have $\pi_\Delta(A \cup B) = \rho_\Delta((A \cup B) \cap \mathbf{S}_\Delta) = \rho_\Delta((A \cap \mathbf{S}_\Delta) \cup (B \cap \mathbf{S}_\Delta)) = \rho_\Delta(A \cap \mathbf{S}_\Delta) + \rho_\Delta(B \cap \mathbf{S}_\Delta) = \pi_\Delta(A) + \pi_\Delta(B)$. Second, I show fineness. Let $\epsilon > 0$. As ρ_Δ is fine, we may partition \mathbf{S}_Δ into $A_1, \dots, A_n \in \mathcal{R}|_{\mathbf{S}_\Delta}$ such that $\rho_\Delta(A_i) < \epsilon$ for all A_i . As each A_i belongs to $\mathcal{R}|_{\mathbf{S}_\Delta}$, we may write it as $A_i = B_i \cap \mathbf{S}_\Delta$ for some $B_i \in \mathcal{R}$. We may take B_1, \dots, B_n to partition \mathbf{S} , by the argument in fn. 43. Now π_Δ is fine as $\pi_\Delta(B_i) = \rho_\Delta(B_i \cap \mathbf{S}_\Delta) = \rho_\Delta(A_i) < \epsilon$ for all i .

Claim 2: π_Δ is the same for all $\Delta \in \Pi$. Let $\Delta, \Delta' \in \Pi$; we show that $\pi_\Delta = \pi_{\Delta'}$. By Claim 1 and Lem. 18 and 19, it suffices to show that π_Δ and $\pi_{\Delta'}$ are ordinally equivalent. Let $A, B \in \mathcal{R}$. As A and B are incorporable, we may pick a context $\alpha \in \Delta$ in which A and B are representable. The events $A_\alpha, B_\alpha (\subseteq S_\alpha)$ representing A resp. B also represent $A \cap \mathbf{S}_\Delta$ resp. $B \cap \mathbf{S}_\Delta$. Now (*) $\pi_\Delta(A) \geq \pi_\Delta(B) \Leftrightarrow A_\alpha \succeq_\alpha B_\alpha$, since $\pi_\Delta(A) \geq \pi_\Delta(B) \Leftrightarrow \rho_\Delta(A \cap \mathbf{S}_\Delta) \geq \rho_\Delta(B \cap \mathbf{S}_\Delta) \Leftrightarrow P_\alpha(A_\alpha) \geq P_\alpha(B_\alpha) \Leftrightarrow A_\alpha \succeq_\alpha B_\alpha$, where the second equivalence holds as ρ_Δ is uncontroversial among $(P_\delta)_{\delta \in \Delta}$ and A_α and B_α represent $A \cap \mathbf{S}_\Delta$ resp. $B \cap \mathbf{S}_\Delta$. Analogously, as A and B are incorporable we may pick an $\alpha' \in \Delta'$ where A and B are representable; as before, (**) $\pi_{\Delta'}(A) \geq \pi_{\Delta'}(B) \Leftrightarrow A_{\alpha'} \succeq_{\alpha'} B_{\alpha'}$. As A and B belong to the robust algebra \mathcal{R} , $A_\alpha \succeq_\alpha B_\alpha \Leftrightarrow A_{\alpha'} \succeq_{\alpha'} B_{\alpha'}$ by Prop. 1, and so $\pi_\Delta(A) \geq \pi_\Delta(B) \Leftrightarrow \pi_{\Delta'}(A) \geq \pi_{\Delta'}(B)$ by (*) and (**), as required.

Claim 3: The (by Claim 2 Δ -independent) probability measure $\rho := \pi_\Delta$ is uncontroversial among the P_α ($\alpha \in \Gamma$). For any $\alpha \in \Gamma$, recall that P_α^* is the function of (representable) objective events $A \subseteq \mathbf{S}$ induced by P_α ; let P_α^{**} be the analogous function induced by P_α w.r.t. the Δ_α -subframework. So P_α^{**} is a function of (representable) $A \subseteq \mathbf{S}_{\Delta_\alpha}$. Now let $A \in \mathcal{R}$, $\alpha \in \Gamma$, and $\Delta := \Delta_\alpha$. As ρ_Δ is uncontroversial among $(P_\gamma)_{\gamma \in \Delta}$, there is a $\beta \in \Delta$ such that P_β^{**} extends P_α^{**} , $S_\beta = S_\alpha \vee \{(A \cap \mathbf{S}_\Delta), \mathbf{S}_\Delta \setminus (A \cap \mathbf{S}_\Delta)\}$ and $P_\beta^{**}(A \cap \mathbf{S}_\Delta) = \rho_\Delta(A \cap \mathbf{S}_\Delta)$. Turning to the general framework, we must show that (i) P_β^* extends P_α^* , (ii) $S_\beta = S_\alpha \vee \{A, \mathbf{S} \setminus A\}$, and (iii) $P_\beta^*(A) = \pi_\Delta(A)$. Claim (i) holds as for all $B \subseteq \mathbf{S}$ in the domain of P_α^* , hence of P_β^* , $P_\alpha^*(B) = P_\alpha^{**}(B \cap \mathbf{S}_\Delta) = P_\beta^{**}(B \cap \mathbf{S}_\Delta) = P_\beta^*(B)$, where the

second equality holds as P_β^{**} extends P_α^{**} , while the first (resp. third) holds as B and $B \cap S_\Delta$ have same representation in context α (resp. β). Claim (ii) holds as $S_\beta = S_\alpha \vee \{(A \cap \mathbf{S}_\Delta), \mathbf{S}_\Delta \setminus (A \cap \mathbf{S}_\Delta)\} = S_\alpha \vee \{A, \mathbf{S} \setminus A\}$. Claim (iii) holds as $P_\beta^*(A) = P_\beta^{**}(A \cap \mathbf{S}_\Delta) = \rho_\Delta(A \cap \mathbf{S}_\Delta) = \pi_\Delta(A) = \rho(A)$. ■

Proof of Lem. 38. Assume Ax. 1–6. Let $\alpha, \beta, \mathcal{R}, f, g$ be as given. Note that $\mathcal{R}|_{\mathbf{S}_\alpha}$ is included in the extrapolated algebra \mathcal{E}_α , as by Lem. 46 $\mathcal{R}|_{\mathbf{S}_\alpha}$ consists of (w.r.t. the Δ_α -subframework) incorporable objective events. As f and g are \mathcal{R} -measurable, $f_{\mathbf{S}_\alpha}$ and $g_{\mathbf{S}_\alpha}$ are $\mathcal{R}|_{\mathbf{S}_\alpha}$ -measurable, so (as $\mathcal{R}|_{\mathbf{S}_\alpha} \subseteq \mathcal{E}_\alpha$) \mathcal{E}_α -measurable. Hence by Lem. 11 (applied to the subframework) we may pick an $\alpha' \in \Delta_\alpha$ such that $f_{\mathbf{S}_\alpha} = \hat{f}^*$ and $g_{\mathbf{S}_\alpha} = \hat{g}^*$ for certain $\hat{f}, \hat{g} \in F_{\alpha'}$ and $S_{\alpha'}$ harmlessly refines S_α ; so, by Lem. 9, $f_{\mathbf{S}_\alpha} \succsim_\alpha^+ g_{\mathbf{S}_\alpha} \Leftrightarrow \hat{f} \succ_{\alpha'} \hat{g}$. By analogous arguments, we may pick a $\beta' \in \Delta_\beta$ such that $f_{\mathbf{S}_\beta} = \tilde{f}^*$ and $g_{\mathbf{S}_\beta} = \tilde{g}^*$ for certain $\tilde{f}, \tilde{g} \in F_{\beta'}$ and $S_{\beta'}$ harmlessly refines S_β ; so, $f_{\mathbf{S}_\beta} \succsim_\beta^+ g_{\mathbf{S}_\beta} \Leftrightarrow \tilde{f} \succ_{\beta'} \tilde{g}$. As $f_{\mathbf{S}_\alpha} \succsim_\alpha^+ g_{\mathbf{S}_\alpha} \Leftrightarrow \hat{f} \succ_{\alpha'} \hat{g}$ and $f_{\mathbf{S}_\beta} \succsim_\beta^+ g_{\mathbf{S}_\beta} \Leftrightarrow \tilde{f} \succ_{\beta'} \tilde{g}$, it suffices to show that $\hat{f} \succ_{\alpha'} \hat{g} \Leftrightarrow \tilde{f} \succ_{\beta'} \tilde{g}$. This holds since \hat{f} and \tilde{f} are corresponding \mathcal{R} -measurable acts (as the \mathcal{R} -measurable function f equals \hat{f}^* on $\mathbf{S}_\alpha = \mathbf{S}_{\alpha'}$ and \tilde{f}^* on $\mathbf{S}_\beta = \mathbf{S}_{\beta'}$) and since also \hat{g} and \tilde{g} are corresponding \mathcal{R} -measurable acts (as the \mathcal{R} -measurable function g equals \hat{g}^* on $\mathbf{S}_\alpha = \mathbf{S}_{\alpha'}$ and \tilde{g}^* on $\mathbf{S}_\beta = \mathbf{S}_{\beta'}$). ■

Proof of Lem. 40. Let $(U_\alpha, P_\alpha)_{\alpha \in \Gamma}$, ρ and \mathcal{R} be as assumed. Let $\Delta \in \Pi$. w.r.t. the Δ -subframework, the subsystem $(U_\alpha, P_\alpha)_{\alpha \in \Delta}$ is still a variable expected-utility representation satisfying R1 and R2, as all this is inherited from the full system. It suffices to show R3. We have

$$B \cap \mathbf{S}_\Delta = C \cap \mathbf{S}_\Delta \Rightarrow \rho(B) = \rho(C) \text{ for all } B, C \in \mathcal{R}, \quad (1)$$

because any $B, C \in \mathcal{R}$ are (by ρ 's uncontroversialness) representable in some context $\alpha \in \Delta$, for which $\rho(B) = P_\alpha^*(B) = P_\alpha^*(B \cap \mathbf{S}_\Delta)$ (the last equality holds as B and $B \cap \mathbf{S}_\Delta$ are represented by the same subjective event) and similarly $\rho(C) = P_\alpha^*(C) = P_\alpha^*(C \cap \mathbf{S}_\Delta)$. Now the function ρ induces a function $\rho_\Delta : \mathcal{R}|_{\mathbf{S}_\Delta} \rightarrow [0, 1]$ by defining, for any $A \in \mathcal{R}|_{\mathbf{S}_\Delta}$, $\rho_\Delta(A) := \rho(B)$, where B is some (hence by (1) any) member of \mathcal{R} such that $B \cap \mathbf{S}_\Delta = A$. By construction, $\rho_\Delta(B \cap \mathbf{S}_\Delta) = \rho(B)$ for all $B \in \mathcal{R}$. So the following two observations complete the proof.

Claim 1: ρ_Δ is a fine probability measure. ρ_Δ inherits these properties from ρ . Indeed, firstly, ρ_Δ is a probability measure, since $\rho_\Delta(\mathbf{S}_\Delta) = \rho(\mathbf{S}) = 1$, and since any disjoint $A, A' \in \mathcal{R}|_{\mathbf{S}_\Delta}$ may be written as $A = B \cap \mathbf{S}_\Delta$ and $A' = B' \cap \mathbf{S}_\Delta$ for some (w.l.o.g.) disjoint sets $B, B' \in \mathcal{R}$, so that

$$\rho_\Delta(A \cup A') = \rho(B \cup B') = \rho(B) + \rho(B') = \rho(A) + \rho(A').$$

Secondly, ρ_Δ is fine, since for each positive $\epsilon > 0$ we may (by ρ 's fineness) partition \mathbf{S} into $B_1, \dots, B_n \in \mathcal{R}$ such that $\rho(B_i) < \epsilon$ for all $i = 1, \dots, n$, and consequently \mathbf{S}_Δ is partitioned into $B_1 \cap \mathbf{S}_\Delta, \dots, B_n \cap \mathbf{S}_\Delta \in \mathcal{R}|_{\mathbf{S}_\Delta}$ (in the broad sense of ‘partitioned’ that

allows some of $B_1 \cap \mathbf{S}_\Delta, \dots, B_n \cap \mathbf{S}_\Delta$ to be empty), where $\rho_\Delta(B_i \cap \mathbf{S}_\Delta) = \rho(B_i) < \epsilon$ for all $i = 1, \dots, n$.

Claim 2: ρ_Δ is uncontroversial (w.r.t. the Δ -subframework). For any $\gamma \in \Delta$, let P_γ^* be (as usual) the function of representable objective events induced by P_γ , and let P_γ^{**} be the analogous function defined w.r.t. the Δ -subframework; so P_γ^* is a function of (representable) subsets of \mathbf{S} , whereas P_γ^{**} is a function of (representable) subsets of \mathbf{S}_Δ . Now consider an $\alpha \in \Delta$ and an $A \in \mathcal{R}|_{\mathbf{S}_\Delta}$. We need to show that there is a $\beta \in \Delta$ such that (a) P_β^{**} extends P_α^{**} , (b) $S_\beta = S_\alpha \vee \{A, \mathbf{S}_\Delta \setminus A\}$, and (c) $P_\beta^{**}(A) = \rho_\Delta(A)$. Write A as $B \cap \mathbf{S}_\Delta$ for some $B \in \mathcal{R}$. As ρ is uncontroversial w.r.t. the general framework, there is a $\beta \in \Gamma$ such that P_β^* extends P_α^* , $S_\beta = S_\alpha \vee \{B, \overline{B}\}$, and $P_\beta^*(B) = \rho(B)$. We may assume w.l.o.g. that $\beta \in \Delta$, as one may verify using independence between outcome and state awareness and Rem. 11. Condition (a) holds because, when restricted to subsets of \mathbf{S}_Δ , P_β^* coincides with P_β^{**} and P_α^* coincides with P_α^{**} . Condition (b) holds because $S_\alpha \vee \{B, \overline{B}\} = S_\alpha \vee \{A, \mathbf{S}_\Delta \setminus A\}$. Condition (c) holds because, as $A \subseteq \mathbf{S}_\Delta$, we have $P_\beta^{**}(A) = P_\beta^*(A)$ and $\rho(A) = \rho_\Delta(A)$. ■

References

- Ahn, D., Ergin, H. (2010) Framing Contingencies, *Econometrica* 78: 655–695
- Anscombe, F. J., Aumann, R. J. (1963) A Definition of Subjective Probability, *Annals of Mathematical Statistics* 34 (1): 199–205
- Dekel, E., Lipman, B. L., Rustichini, A. (1998) Standard state-space models preclude unawareness, *Econometrica* 66: 159–73
- Halpern, J. Y. (2001) Alternative Semantics for Unawareness, *Games and Economic Behavior* 37: 321–39
- Halpern, J. Y., Rego, L. C. (2008) Interactive Unawareness Revisited, *Games and Economic Behavior* 62: 232–62
- Hill, B. (2010) Awareness Dynamics, *Journal of Philosophical Logic* 39: 113–37
- Karni, E., Schmeidler, D. (1991) Utility Theory with Uncertainty. In: *Handbook of Mathematical Economics*, Vol. 4, edited by Werner Hildenbrand and Hugo Sonnenschein, 1763–1831, New York: Elsevier Science
- Karni, E., Viero, M. (2013) Reverse Bayesianism: a choice-based theory of growing awareness, *American Economic Review* 103: 2790–2810
- Karni, E., Viero, M. (2015) Awareness of unawareness: a theory of decision making in the face of ignorance, working paper, Johns Hopkins University
- Kopylov, I. (2007) Subjective probabilities on “small” domains, *Journal of Economic Theory* 133: 236–265
- Niiniluoto, I. (1972) A note on fine and tight qualitative probabilities, *Annals of Mathematical Statistics* 43: 1581–91
- Pivato, M., Vergopoulos, V. (2015) Categorical decision theory, working paper,

University Cergy-Pontoise

Savage, L. J. (1954) *The Foundations of Statistics*, New York: Wiley

Schmeidler, D., Wakker, P. (1987) Expected Utility and Mathematical Expectation. In: *The New Palgrave: A Dictionary of Economics*, first edition, edited by J. Eatwell, M. Milgate, and P. Newman, New York: Macmillan Press

Wakker, P. (1981) Agreeing probability measures for comparative probability structures, *Annals of Statistics* 9: 658-62