

Savage's Theorem under changing awareness

Franz Dietrich¹

Final version of January 2018 for *Journal of Economic Theory*

Abstract

This paper proposes a simple unified framework of choice under changing awareness, addressing both *outcome* awareness and (*nature*) *state* awareness, and both how *fine* and how *exhaustive* the awareness is. Six axioms characterize an (essentially unique) expected-utility rationalization of preferences, in which utilities and probabilities are revised according to three revision rules when awareness changes: (R1) utilities of unaffected outcomes are transformed affinely; (R2) probabilities of unaffected events are transformed proportionally; (R3) enough probabilities ‘objectively’ never change (they represent revealed objective risk). Savage’s Theorem is a special case of the theorem, namely the special case of fixed awareness, in which our axioms reduce to Savage’s axioms while R1 and R2 hold trivially and R3 reduces to Savage’s requirement of atomless probabilities. Rule R2 parallels Karni and Viero’s (2013) ‘reverse Bayesianism’ and Ahn and Ergin’s (2010) ‘partition-dependence’. The theorem draws mathematically on Kopylov (2007), Niiniluoto (1972) and Wakker (1981). (*JEL codes*: D81, D83.)

Keywords: Decision under uncertainty, outcome unawareness versus state unawareness, non-fine versus non-exhaustive awareness, utility revision versus probability revision, small worlds versus grand worlds

1 Introduction

Savage’s (1954) expected-utility framework is the cornerstone of modern decision theory. A widely recognized problem is that Savage relies on sophisticated and stable concepts of outcomes and (nature) states: ideally, *outcomes* always capture everything that matters ultimately, and *states* always capture everything that influences outcomes of actions.² In real life, an agent’s concepts or ‘awareness’ can

¹Paris School of Economics & CNRS; fd@franzdietrich.net; www.franzdietrich.net.

²This ideal translates partly into Savage’s formal analysis: his axioms imply high state sophistication (i.e., infinitely many states), while permitting low outcome sophistication (i.e., possibly just two outcomes). So Savage’s formal model can handle an unsophisticated outcome concept, but neither an unsophisticated state concept, nor changing state or outcome concepts.

be limited at two levels, and in two ways. It can be limited at the outcome or state level, and it can be non-fine (coarse) or non-exhaustive (domain-restricted). A social planner deciding where to build a new nuclear power plant on his island 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 arriving from the east, west, north, or south. These are examples of *state* unawareness; analogous examples exist for *outcome* unawareness. Figure 1




Figure 1: An objective act (left) and its two subjective representations f (middle) and g (right) under two different awareness levels (indicated by circles)

shows a situation where from an omniscient third-person perspective there are four objective states a, b, c, d and four objective outcomes x, y, z, w . The left-hand plot shows an objective act whose outcome is x under state c and is z otherwise. The middle and right-hand plots show two subjective awareness levels of the agent who both times conceives only two outcomes and two states; each *subjective* outcome (or state) is given by a set of one or more *objective* outcomes (or states), indicated by a circle. In the middle plot, the agent lacks a fine awareness of outcomes and states: x and y are merged into the same subjective outcome, and a and b into the same subjective state. He also lacks an exhaustive awareness of outcomes and states: w and d are ignored, i.e., absent from all subjective outcomes or states. The mentioned objective act is conceived as an act f mapping $\{a, b\}$ to $\{z\}$, and $\{c\}$ to $\{x, y\}$. In the right-hand plot, awareness is still not fine, but it is exhaustive, both at the outcome and state level. The act is now reconceived as a constant act g which yields outcome $\{x, z\}$ at both subjective states $\{a, b\}$ and $\{c, d\}$. One might compare our objective and subjective states with Savage's (1954) grand-world and small-world states, respectively, although we allow changes in subjective states while Savage takes both types of states to be fixed.

An agent with an *expected-utility rationalization* does in each awareness state or context hold (i) a utility function over currently conceived subjective outcomes and (ii) a probability function over currently conceived subjective states. So in the awareness state shown in the middle of Figure 1 the agent assigns utilities only to the subjective outcomes $\{x, y\}$ and $\{z\}$, and probabilities only to the subjective states $\{a, b\}$ and $\{c\}$. I will consider revision rules governing the change in utilities and probabilities as the agent's awareness or concepts change. The first two rules

are:

R1: *Utilities of preserved subjective outcomes are transformed in an increasing affine way.*

R2: *Probabilities of preserved subjective events are transformed proportionally.*

These rules are vacuous when applied to the change from the middle awareness state in Figure 1 to the right one, since no subjective outcome and only one subjective state (i.e., $\{a, b\}$) is preserved. But now assume the middle awareness state changes differently: all existing subjective outcomes and states are preserved, and the new ones $\{w\}$ and $\{d\}$ are added. Then R2 requires the new probabilities of $\{a, b\}$ and $\{c\}$ to be proportional to the old ones, and R1 requires the new utilities of $\{x, y\}$ and $\{z\}$ to be an increasing transformation of the old ones ('affineness' is vacuous in case of only two preserved outcomes).

There is a clear need for a generalization of Savage's expected-utility theory so as to cope with changes in awareness of the various sorts. If such a generalization has not yet been offered, it is perhaps because of two obstacles. One is the literature's almost exclusive focus on *state* unawareness; I hope to raise 'awareness' of *outcome* unawareness. Another obstacle is Savage's high demands of state sophistication which go against the idea of state unawareness; we will find a way to require less state sophistication, allowing for finite state spaces.

I shall offer a Savagean expected-utility ('EU') theory under changing awareness, with 'rational' revision of utilities and probabilities. I take the agent to be classical in all respects except from changing awareness (future research might explore *non-EU* preferences under changing awareness and/or *boundedly* rational revision rules). I work within a simple unified model of changing awareness, capturing changes in outcome and state awareness, and in refinement and exhaustiveness. Six axioms are then introduced, and shown to characterize an EU agent who, under any change in awareness, updates his utilities and probabilities according to the rules R1 and R2 and a third rule stated later. Probabilities are unique, and utilities are unique up to increasing affine rescaling. Utility revision is a genuine necessity: utilities cannot generally be scaled such that R1's transformation is always the identity transformation. The theorem generalizes Savage's Theorem: it reduces to it in the limiting case of stable awareness, as our axioms then reduce to Savage's axioms, while rules R1 and R2 hold trivially and the third rule reduces to Savage's *atomlessness* condition on beliefs.

The framework allows for different interpretations. For instance, the agent's awareness level could have different sources; one of them is the framing of the decision problem. Also, unawareness could be of radical and non-radical type. Radical unawareness of X is an in-principle inability to imagine or represent X. As yet unexperienced dimensions of reality or undiscovered phenomena presumably

fall under radical unawareness. Non-radical unawareness of X means that we merely do not consider X , be it because we set X aside on purpose or overlook X by mistake. Although we are in principle able to consider or understand X , we leave X aside – either because X is not worth considering due to mental costs, or because X escapes our attention due to framing or other circumstances. For instance, in a cooking choice we ignore a coin toss just because nothing hinges on it, and forget to ask how salty the dish will taste out of distraction; but we are *radically* unaware of tastes and flavours we have never experienced.

Choice theorists have tackled unawareness in different ways. The agent’s (un)awareness level can be an input or an output of the analysis: it can be an exogenous starting point which is assumed, or a feature which should be revealed by observed behaviour. Recent examples of the ‘revealed (un)awareness’ approach are Schipper (2013) and Kochov (2016).³ My model follows the ‘exogenous awareness’ approach, just like Ahn-Ergin’s (2010) model of framed contingencies and Karni-Viero’s (2013) model of growing awareness. How does my model relate to these two seminal contributions? Working in an Anscombe-Aumann-type framework, Ahn-Ergin 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 EU representation with fixed utilities and partition-dependent probabilities. The systematic way in which these probabilities change with the partition implies our rule R2 (after suitable translations). Karni-Viero, by contrast, model the discovery of new acts, outcomes, and act-outcome links. They work in a non-standard framework which takes acts as primitive objects and states as functions from acts to outcomes (following Schmeidler and Wakker 1987 and Karni and Schmeidler 1991). 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 current analysis differs strongly from Ahn-Ergin’s and Karni-Viero’s. I now mention some differences. I analyse awareness change at 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.⁴ Ahn-Ergin and Karni-Viero find that only probabilities are revised, yet I find that also

³Schipper takes unawareness of an event X to be revealed via nullness of both X and X ’s negation. Kochov studies revealed unawareness of future contingencies in a dynamic setting.

⁴Karni-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. outcomes effectively renders states finer resp. more exhaustive.

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 that ‘state sophistication’ is reached *sometimes*, or *never*, or *always*; this flexibility is crucial for ‘generalizing Savage’.

In the background of the paper is a vast and active literature on unawareness (e.g., Dekel, Lipman and Rustichini 1998, Halpern 2001, Heifetz, Meier and Schipper 2006, 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. The theorem’s long proof, presented in different appendices, makes use of key theorems by Kopylov (2007), Niiniluoto (1972) and Wakker (1981).

2 A model of changing awareness

2.1 Variable Savage structures

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

Definition 1 A *Savage structure* is a triple (X, S, \succsim) of a non-empty finite⁵ 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 structures (X_t, S_t, \succsim_t) where t ranges over an arbitrary set of ‘contexts’. To be able to relate the subjective outcomes/states of each context to objective ones, I take each X_t to partition (coarsen) some underlying set of ‘objective’ outcomes, and each S_t to partition (coarsen) some underlying set of ‘objective’ states.⁶ This captures the ‘objective world’ and the agent’s (changing) awareness of it. Formally:

Definition 2 A *variable Savage structure* is a family of Savage structures $(X_t, S_t, \succsim_t)_{t \in T}$ indexed by some non-empty set T (of *contexts*), such that

- each X_t is a partition of some set (of *objective outcomes encompassed in context* t),
- each S_t is a partition of some set (of *objective states encompassed in context* t).

⁵Savage in fact did not impose finiteness. I add finiteness for simplicity.

⁶A **partition** of a set is a set of *non-empty*, pairwise exclusive and exhaustive subsets.

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

From now on, let $(X_t, S_t, \succsim_t)_{t \in T}$ be a given variable Savage structure. Let:

$$\begin{aligned}
F_t &:= X_t^{S_t} && \text{(acts conceived in context } t) \\
\mathbb{X}_t &:= \text{set partitioned by } X_t && \text{(objective outcomes encompassed in context } t) \\
\mathbb{S}_t &:= \text{set partitioned by } S_t && \text{(objective states encompassed in context } t) \\
\mathbb{X} &:= \cup_{t \in T} \mathbb{X}_t && \text{(objective outcomes)} \\
\mathbb{S} &:= \cup_{t \in T} \mathbb{S}_t && \text{(objective states)} \\
\mathbb{F} &:= \mathbb{X}^{\mathbb{S}} && \text{(objective acts).}
\end{aligned}$$

The spaces X_t and S_t could reflect the *framing* at work in context t . This framing renders certain outcome/state concepts salient, e.g., through a *mode of presentation* or a *level of descriptive detail*.⁷ The framing-based interpretation follows Ahn and Ergin (2010), extending their idea to outcome framing as well as state framing. Framing is only one of many possible sources/interpretations of the agent's concepts and hence of the spaces X_t and S_t ; see Section 2.3.

Here is a two-context example corresponding exactly to Figure 1: $T = \{t, t'\}$, $X_t = \{\{x, y\}, \{z\}\}$, $S_t = \{\{a, b\}, \{c\}\}$, $X_{t'} = \{\{x, z\}, \{y, w\}\}$, and $S_{t'} = \{\{a, b\}, \{c, d\}\}$; so $\mathbb{X}_t = \{x, y, z\}$, $\mathbb{S}_t = \{a, b, c\}$, $\mathbb{X} = \mathbb{X}_{t'} = \{x, y, z, w\}$, and $\mathbb{S} = \mathbb{S}_{t'} = \{a, b, c, d\}$. The agent's outcome/state spaces are non-fine in both contexts, and non-exhaustive only in context t . In general, the smaller the sets in X_t and S_t are, the finer the agent's outcome/state concepts are, up to the point of singleton sets (full refinement); and the larger the sets \mathbb{X}_t and \mathbb{S}_t are, the more exhaustive these concepts are, up to the entire sets \mathbb{X} and \mathbb{S} (full exhaustiveness).

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

Definition 3 *The variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ has*

- (a) **exhaustive outcomes** if in all contexts t we have $\mathbb{X}_t = \mathbb{X}$,
- (b) **exhaustive states** if in all contexts t we have $\mathbb{S}_t = \mathbb{S}$,
- (c) **fine outcomes** if in all contexts t each outcome $x \in X_t$ is singleton,
- (d) **fine states** if in all contexts t each state $a \in S_t$ is singleton.

Our theorem will simplify under exhaustive states, and simplify *differently* under fine states. Here are examples demonstrating the flexibility of our model:

- **Savage:** T contains a single context t_0 . Our variable structure reduces to a classic Savage structure $(X, S, \succsim) := (X_{t_0}, S_{t_0}, \succsim_{t_0})$. Objective outcomes and

⁷If 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_t contains the 11 events ' n accidents' for $n = 0, 1, \dots, 10$, and X_t contains the 11 net-benefit outcomes; another context t' with a different mode of presentation will induce different spaces $S_{t'}$ and $X_{t'}$.

states are not needed: w.l.o.g. we can, like Savage, let X and S be primitive sets rather than any partitions.

- **Stable outcome awareness:** All contexts t lead to the same outcome space $X_t = X$, which we may take as a primitive set rather than any partition. One might then identify contexts with state spaces (*state awareness levels*); so T contains the ‘possible/feasible’ state spaces and $S_t := t$ for all $t \in T$.
- **Stable state awareness:** All contexts t lead to the same state space $S_t = S$, which we may take to be a primitive set rather than a partition. One might then identify contexts with outcome spaces (*outcome awareness levels*), so that T contains the ‘possible/feasible’ outcome spaces and $X_t := t$ for all $t \in T$.
- **Fully variable awareness:** All logically possible awareness states occur: for all partitions X of \mathbb{X} and S of \mathbb{S} , where $|X| < \infty$, there is a context $t \in T$ in which $X_t = X$ and $S_t = S$. (To allow non-exhaustive awareness, one should extend the quantification to partitions of any non-empty subset of \mathbb{X} or \mathbb{S} , respectively.) This permits arbitrary ways to conceive the world. If one identifies contexts with awareness levels, T becomes the set of all pairs $t = (X, S)$ of a finite partition of \mathbb{X} and a partition of \mathbb{S} , where $X_t := X$ and $S_t = S$.
- **Finite awareness:** All spaces X_t and S_t (and so all act sets F_t) are finite. The agent can only conceive finitely many things at a time.

2.2 Excursion: preferences over subjective or objective acts?

Some might prefer working with preferences over objective rather than subjective acts (I do not, because the objects of preferences should be things the agent can conceive, in line with my mentalistic rather than behaviouristic position on economic modelling; see Dietrich and List 2016a, 2016b). One can indeed recast each preference relation \succsim_t in our variable Savage structure as one over objective acts, more precisely over objective acts that are representable (measurable) in context t : just identify a preference between two subjective acts in F_t with a preference between any two objective acts represented by them in context t , respectively. If wished, one could thus replace the relations \succsim_t on F_t ($t \in T$) by corresponding relations on $\{f \in \mathbb{F} : f \text{ is representable in } F_t\}$ ($t \in T$), or at least *interpret* the former relations as shorthands for the latter relations. But what exactly is the subjective representation of an objective act?

Definition 4 *Given a context $t \in T$,*

- *the **subjectivization** of an objective outcome $x \in \mathbb{X}_t$, denoted x_t , is the subjective outcome in X_t containing x (the assignment $x \mapsto x_t$ maps \mathbb{X}_t onto*

X_t),

- the **subjectivization** of an objective state $a \in \mathbb{S}_t$, denoted a_t , is the subjective state in S_t containing a (the assignment $a \mapsto a_t$ maps \mathbb{S}_t onto S_t),

Definition 5 The (**subjective**) **representation** of an objective act $f \in \mathbb{F}$ in a context $t \in T$ is the unique act $f_t \in F_t$ (if existent) which matches f modulo subjectivization, in that whenever $a' \in S_t$ subjectivizes $a \in \mathbb{S}_t$ (i.e., $a \in a'$), then $f_t(a')$ subjectivizes $f(a)$ (i.e., $f(a) \in f_t(a')$). An $f \in \mathbb{F}$ is (**subjectively**) **representable** in context t if its representation f_t exists.

Remark 1 (representability as measurability) Given a context $t \in T$,

- under exhaustive states and outcomes, an objective act is representable if and only if it is (S_t, X_t) -measurable,
- in general, an objective act is representable if and only if its restriction to \mathbb{S}_t maps into \mathbb{X}_t and is (S_t, X_t) -measurable.⁸




Figure 2: An objective act $f : \mathbb{S} \rightarrow \mathbb{X}$ and its representation $f_t : S_t \rightarrow X_t$ (note that f indeed maps \mathbb{S}_t into \mathbb{X}_t by Remark 2)

Remark 2 For any objective act $f \in \mathbb{F}$ representable in given a context t , the diagram in Figure 2 is

- (i) well-defined, i.e., $f(\mathbb{S}_t) \subseteq \mathbb{X}_t$ (a trivial condition under exhaustive outcomes),
- (ii) commutative, i.e., $[f(a)]_t = f_t(a_t)$ for all $a \in \mathbb{S}_t$.

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 t the agent conceives only ‘wealth outcomes’ and ‘coin states’: $X_t = \{rich, poor\}$ and $S_t = \{heads, tails\}$, where *rich* and *poor* are the outcomes (sets of objective outcomes) in which he is rich or poor respectively, and *heads* and *tails* are the states (sets of objective states) in which the coin lands heads or tails respectively. Then f is represented by the subjective act f_t that maps *heads* to *rich* and *tails* to *poor*. But if instead $X_t = \{rich, poor\}$ and $S_t = \{\mathbb{S}\}$, the state concept no longer captures the coin toss, and f is no longer representable.

⁸ (S_t, X_t) -measurability means that members of the same $a \in S_t$ are mapped into the same $x \in X_t$, or equivalently, that the inverse image of any $x \in X_t$ is a union of zero or more $a \in S_t$.

2.3 Three clarifying interpretive remarks

1. One can take the spaces X_t and S_t ($t \in T$) to represent the awareness/concepts/ontology *ascribed* to the agent *by the observer*: they reflect how we take him to perceive the world, and hence they embody our hypothesis (or theory, stipulation, conjecture etc.) about the agent's perception. This ascription or hypothesis could be guided by: (i) the framing of outcomes and states; or (ii) the awareness/concepts as *reported* by the agent when asked by the analyst about his current outcome/state concepts; or (iii) the modeller's own common sense or intuition; or (iv) neuro-physiological evidence about how the context affects the cognitive system; or (v) the sort of options that are *feasible* in context t (here X_t and S_t are constructed such that all feasible options become representable as subjective acts, in a sense made precise in Section 2.7); or (vi) *observed choices* (among objective acts) which *reveals* the agent's awareness, in a sense that can be made precise (here X_t and S_t are constructed so as to be fine enough to distinguish between those objective acts between which observed choices distinguish). There is also a completely different interpretation: the spaces X_t and S_t could represent the agent's *real* rather than *ascribed* concepts in context t , adopting a first-person rather than third-person perspective. The literature routinely uses (and switches between) both types of interpretation of Savage's outcomes and states – i.e., interpretations in terms of the agent's *ascribed* or *real* ontology, i.e., of either *our assumption* about his concepts or his *real* concepts possibly unknown to us. Savage himself had the second interpretation in mind; he focused on rationality from a first-person perspective. By contrast, mainstream behaviourist economics favours the first interpretation, to ensure observability of all primitives.

2. By modelling subjective outcomes/states as sets of objective ones, I by no means suggest that the agent conceives outcomes/states *in terms of* (complex) sets. He may conceive them as indecomposable primitives. He may for instance conceive the outcome 'having close friends' in complete unawareness of a huge (infinite) set of underlying objective outcomes. Only our third-person perspective identifies subjective outcomes/states with sets of objective ones.

3. One can think of a context $t \in T$ in broadly two ways. *Either* t represents the environment (or frame, time point, decision node in a decision/game tree, ...) which triggers or causes the agent's awareness state (X_t, S_t) and preference relation \succsim_t . *Or* t is 'only' an index; t can then be identified with the agent's awareness state itself (under a minor loss of generality⁹), so that t is a space pair (X, S) , where $X_t := X$ and $S_t := S$. In the first case $(X_t, S_t, \succsim_t)_{t \in T}$ captures how awareness and preferences *react* causally to the environment. In the second

⁹The loss of generality is that preference must then be determined (fully) by awareness: $(X_t, S_t) = (X_{t'}, S_{t'}) \Rightarrow \succsim_t = \succsim_{t'}$. This restriction is minor, since even without making it, it later follows from Axiom 2 (see Proposition 1).

case $(X_t, S_t, \succsim_t)_{t \in T}$ is simply a summary of all ‘possible’ or ‘feasible’ awareness-preference states (X, S, \succsim) of the agent; the structure $(X_t, S_t, \succsim_t)_{t \in T}$ could then be abbreviated as $(\succsim_t)_{t \in T}$, as each context $t = (X, S)$ already encodes the spaces $X_t := X$ and $S_t := S$.

2.4 The objective/subjective terminology and notation

We must carefully distinguish between objective and subjective descriptions. By default descriptions are subjective; so I often drop ‘subjective’. Formally:

- An **objective outcome/state/act/event** is a member of $\mathbb{X} / \mathbb{S} / \mathbb{F} = \mathbb{X}^{\mathbb{S}} / 2^{\mathbb{S}}$.
- A **(subjective) outcome/state/act/event conceived in context** t ($\in T$) is a member of $X_t / S_t / F_t = X_t^{S_t} / 2^{S_t}$.
- A **(subjective) outcome/state/act/event simpliciter** (without reference to a context) is a member of *some* $X_t / S_t / F_t / 2^{S_t}$ ($t \in T$).
- The **(subjective) outcome/state space in context** t is X_t / S_t .
- A **(subjective) outcome/state space simpliciter** is *some* X_t / S_t ($t \in T$).
- The **objectivization** of an event $A \subseteq S_t$ ($t \in T$), denoted A^* , is the objective event partitioned by A ; A and A^* are said to **correspond** to each other.
- Events A and B from possibly different contexts are **(objectively) equivalent** if $A^* = B^*$.
- Any act $f \in F_t$ ($t \in T$) induces a function on \mathbb{S}_t rather than S_t , denoted f^* and given by $f^*(a) := f(a_t)$ (where a_t is a ’s subjectivization, given by $a \in a_t \in S_t$); f and f^* are said to **correspond** to each other.

2.5 Expected-utility rationalizations and revision rules

We can apply the classical expected-utility paradigm *within* each context.

Definition 6 *An expected-utility (‘EU’) rationalization/representation of the variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ is a system $(U_t, P_t)_{t \in T}$ of non-constant ‘utility’ functions $U_t : X_t \rightarrow \mathbb{R}$ and probability measures¹⁰ $P_t : 2^{S_t} \rightarrow [0, 1]$ such that*

$$f \succsim_t g \Leftrightarrow \mathbb{E}_{P_t} U_t(f) \geq \mathbb{E}_{P_t} U_t(g) \text{ for all contexts } t \in T \text{ and acts } f, g \in F_t.$$

Our axioms will imply existence of an (essentially unique) EU rationalization $(U_t, P_t)_{t \in T}$ whose utilities and probabilities obey three revision rules. The first two of them are easily stated:

¹⁰The term ‘probability measure’ is used throughout in its *finitely* additive sense.

R1: Any U_t is an increasing affine transformation of any $U_{t'}$ on the domain overlap $X_t \cap X_{t'}$.

R2: Any P_t is proportional to any $P_{t'}$ on the domain overlap $2^{S_t \cap S_{t'}}$.

Under R1 and R2, utilities are affinely rescaled and probabilities are proportionally rescaled where concepts are stable. So if the agent, say, splits an outcome $x \in X_t$ into y and z , resulting in a context t' with $X_{t'} = (X_t \setminus \{x\}) \cup \{y, z\}$ and $S_{t'} = S_t$, then $P_{t'} = P_t$ by R2, and utilities are essentially unchanged on $X_t \setminus \{x\}$ by R1.

R1 and R2 are in contrast with two stronger conditions that forbid revisions:

R1+: Any U_t equals any $U_{t'}$ on the domain overlap $X_t \cap X_{t'}$ ('stable utilities').

R2+: Any P_t equals any $P_{t'}$ on the domain overlap $2^{S_t \cap S_{t'}}$ ('stable probabilities').

Neither R2 nor even R2+ prevents the agent from attaching a different probability to an event $A \subseteq S_t$ conceived in a context t and an *objectively equivalent* event $A' \subseteq S_{t'}$ conceived in another context t' . This happens if beliefs are description-sensitive, i.e., dependent on how objective events are perceived subjectively. Imagine that in context t the agent conceives the fine states $\{a\}$ and $\{b\}$ and hence the event $A = \{\{a\}, \{b\}\}$, while in context t' he conceives the coarser state $\{a, b\}$ and hence the event $A' = \{\{a, b\}\}$. Although A and A' represent the same objective event $\{a, b\}$, the agent might in context t find A unlikely on the grounds that $\{a\}$ and $\{b\}$ each appear implausible, while in context t' finding A' likely because he fails to analyse this event in terms of its implausible subcases.¹¹ The following revision rule – a significant strengthening of R2 and R2+ – excludes such description-sensitivity:

R2++: If events $A \subseteq S_t$ and $A' \subseteq S_{t'}$ ($t, t' \in T$) are objectively equivalent, then they get the same probability, i.e., $A^* = A'^* \Rightarrow P_t(A) = P_{t'}(A')$ ('objectively stable probabilities').

I now give three examples. All of them assume an EU rationalization $(U_t, P_t)_{t \in T}$. The first two concern utilities, and the last one concerns probabilities.

Example 1: stable utilities. Objective outcomes are numbers between 0 and 100: $\mathbb{X} = [0, 100)$. The agent has a coarse conception of numbers, i.e., conceives 'vague numbers' in the form of intervals. So each outcome space X_t consists of pairwise disjoint intervals. If for instance $X_t = \{[k, k+1) : k = 0, \dots, 99\}$, then the agent effectively ignores decimals, i.e., identifies any numbers having the same non-decimal digits. How might he assign utilities? Suppose all utilities are reducible

¹¹Concretely, a could stand for country 1 attacking country 2, and b for 2 attacking 1. In context t the agent finds event $A = \{\{a\}, \{b\}\}$ unlikely: he reasons that $\{a\}$ and $\{b\}$ are each implausible as each country is unlikely to attack. In context t' , he finds event $A' = \{\{a, b\}\}$ likely on unsophisticated grounds: he treats A' as a primitive scenario of 'war', which seems likely to him, as he fails to realise that a war requires an (unlikely) attack by either country.

to a fixed utility function of *objective* states $U : \mathbb{X} \rightarrow \mathbb{R}$, in one of the following ways. For any context t and any outcome $I \in X_t$ (an interval),

- $U_t(I) = U(m)$ where m is a ‘representative’ or ‘rounded’ number defined, e.g., as I ’s lower boundary $\inf I$ or upper boundary $\sup I$ or midpoint $\frac{\inf I + \sup I}{2}$;
- $U_t(I) = u$ where u is a ‘representative’ or ‘rounded’ utility level defined, e.g., as $\inf_{z \in I} U(z)$ or $\sup_{z \in I} U(z)$ or $\frac{\inf_{z \in I} U(z) + \sup_{z \in I} U(z)}{2}$.

In all these cases utility revision satisfies not just R1, but even R1+ (stable utilities), since $U_t(I)$ depends only on I , not on the context t .¹²

Example 2: changing utilities. As in the previous example, let $\mathbb{X} = [0, 100)$, and let each X_t consist of pairwise disjoint intervals. But utilities are no longer reducible to a fixed utility function on \mathbb{X} . Instead they are formed as follows. In any context t , the outcomes (intervals) in X_t are put into a linear order I_1, I_2, \dots, I_n such that I_1 ’s members are smaller than I_2 ’s, which are smaller than I_3 ’s, etc., where $n = |X_t|$. Let $U_t(I_1) = 1$, $U_t(I_2) = 2$, ..., $U_t(I_n) = n$. This leads to utility revisions: for instance, the same outcome $[50, 100)$ gets utility $U_t([50, 100)) = 2$ when $X_t = \{[0, 50), [50, 100)\}$, but utility $U_{t'}([50, 100)) = 3$ when $X_{t'} = \{[0, 10), [10, 50), [50, 100)\}$. Neither R1+ nor even R1 needs to hold: two utility functions U_t and $U_{t'}$ need not be increasing affine transformations of one another on $X_t \cap X_{t'}$. However, R1 does hold if never more than 3 outcomes are conceived (i.e., $|X_t| \leq 3$ for all $t \in T$), or never more than 4 outcomes are conceived *and* outcome awareness is exhaustive (i.e., each X_t partitions the full interval $\mathbb{X} = [0, 100)$). The reason is that in these cases any two spaces $X_t, X_{t'}$ either share at most two elements (so that U_t is an increasing affine transformation of $U_{t'}$ on $X_t \cap X_{t'}$), or coincide (so that $U_t = U_{t'}$).

Example 3: objectively stable beliefs. Let the source of uncertainty be purely objective. That is, let \mathbb{S} contain the outcomes of some random experiment, e.g., coin tossing sequences or roulette outcomes. We capture the objective risk by a probability measure π on some algebra \mathcal{R} on \mathbb{S} . In each context t the agent conceives only finitely many states: S_t is a finite partition of \mathbb{S} into members of \mathcal{R} . (If \mathbb{S} contains coin tossing sequences and in context t only the first three tosses are conceived, then states in S_t correspond to triples of outcomes of the first three tosses.) Despite his limited state conception, let the agent give the true (objective) probability to those states he conceives: $P_t(s) = \pi(s)$ for all contexts t and states $s \in S_t$. Then clearly R2++ (and thus R2 and R2+) hold.

¹²Is it plausible to reduce all U_t s to a fixed function U on \mathbb{X} ? One might object that agents with limited awareness do not conceive the objects \mathbb{X} and U , and so cannot consciously calculate values like $U(\inf I)$ or $\inf_{z \in I} U(z)$; the utility model would thus have an ‘as if’ status, hence become questionable.

2.6 Tighter EU rationalizations and their revision rules

I now sketch tighter kinds of EU rationalizations with fewer degrees of freedom: ‘stable’, ‘unified’, and ‘classical’ EU rationalizations. In each case some of the revision rules R1, R1+, R2, R2+ and R2++ come for granted. The unified and classical cases can be linked to Ahn-Ergin’s (2010) central rationalization concepts.¹³

Stable and unified EU rationalizations. I begin with utilities, and then turn to probabilities.

Definition 7 A family $(U_t)_{t \in T}$ of utility functions U_t on X_t is **stable** if it is given by a single function U on $\cup_{t \in T} X_t$ in that each U_t matches U on X_t , i.e., $U_t = U|_{X_t}$.

Remark 3 Stability of $(U_t)_{t \in T}$ is equivalent to R1+, so implies R1.

Definition 8 A family $(P_t)_{t \in T}$ of probability functions P_t on 2^{S_t} is

- **stable** if it is given by a single function P on $\cup_{t \in T} 2^{S_t}$ in that each P_t is equal to P on 2^{S_t} , i.e., $P_t = P|_{2^{S_t}}$,
- **unified** if more generally it is given by a single function P on $\cup_{t \in T} 2^{S_t}$ in that each P_t is proportional to P on 2^{S_t} , i.e., $P_t \propto P|_{2^{S_t}}$.¹⁴

Remark 4 Stability of $(P_t)_{t \in T}$ is equivalent to R2+, which implies unification, which in turn implies R2.

Definition 9 If an EU rationalization $(U_t, P_t)_{t \in T}$ is stable in its utility functions (given by U) and stable or more generally unified in its probability functions (given by P), then $(U_t, P_t)_{t \in T}$ – or in short (U, P) – is a **stable EU rationalization** or more generally **unified EU rationalization**, respectively.¹⁵

Remark 5 An EU rationalization $(U_t, P_t)_{t \in T}$ is

- stable if and only if R1+ and R2+ hold,
- unified if R1+ and R2+ hold, and only if R1+ and R2 hold.

¹³If Ahn-Ergin’s framework is recast within ours, it corresponds to the special case of a fixed outcome space $X_t = X$ and exhaustive and finite state spaces S_t . In this case, our *unified* and *classical* EU rationalizations reduce essentially to their ‘partition-dependent’ resp. ‘partition-independent’ EU representations (partly because the conditions which unification and classicality impose on the functions U_t reduce to Ahn-Ergin’s assumption of a fixed utility function).

¹⁴Even if P_t is just proportional to P on 2^{S_t} , P fully determines P_t , given that $P_t(S_t) = 1$.

¹⁵One might also consider EU rationalizations which are stable only in utilities (short-hand: $(U, (P_t)_{t \in T})$), or stable or more generally unified only in probabilities (short-hand: $((U_t)_{t \in T}, P)$).

Classical EU rationalizations. *Stability* of an EU rationalization is a big (and questionable) step towards classical rationality. But it falls short of it, as the agent can display major forms of dynamic inconsistency. For instance, as the context changes from t to t' and two old outcomes $x, y \in X_t$ get merged into a single one $x \cup y \in X_{t'}$, the utility $U_{t'}(x \cup y)$ could exceed the utilities $U_t(x)$ and $U_t(y)$ of *both* suboutcomes – a ‘dynamic dominance violation’. Similarly, as two old states $a, b \in S_t$ get merged into a single one $a \cup b \in S_{t'}$, the probability $P_{t'}(a \cup b)$ can differ from $P_t(a) + P_t(b)$ – a ‘dynamic additivity violation’. The natural explanation of such violations is, of course, that in context t' the agent no longer conceives the ‘subcases’ of $x \cup y$ and $a \cup b$, as $x \cup y$ and $a \cup b$ appear as indecomposable primitives.

I now define ‘classical’ EU rationalizations. They exclude any such dynamic inconsistency, and render the agent classical in all respects – except from awareness changes (imposed exogenously via the variable Savage structure). Such an ‘almost classical’ agent ranks subjective acts *as if* he conceived the underlying objective outcomes and states: although he evaluates *subjective* entities (outcomes, states and acts), he does so *objectively*, by taking into account the underlying objective outcomes and states. One may legitimately question the plausibility of such a hybrid agent: why should someone who can come up with objective evaluations fail to come up with objective outcomes, states and acts in the first place? The point of defining classical EU rationalizations is not to defend ‘objective evaluations of subjective objects’ as genuinely realistic, but to spell out the classical benchmark from which our less classical rationalizations depart.

Let us start with classical probabilities, before turning to classical utilities. From an orthodox rationality perspective, it should not matter how objective events are subjectively represented (described, framed), in the following sense:

Definition 10 *A family $(P_t)_{t \in T}$ of probability functions P_t on 2^{S_t} is **classical** if the probability of any event A depends only on its objectivization A^* , i.e., if $(P_t)_{t \in T}$ is given by a single function π on the set $\cup_{t \in T} \{A^* : A \in 2^{S_t}\}$ of objectivized events:*

$$P_t(A) = \pi(A^*) \text{ for all contexts } t \in T \text{ and events } A \subseteq S_t.$$

Remark 6 *Classicality of $(P_t)_{t \in T}$ is equivalent to $R2++$ (objectively stable probabilities), hence implies $R2$ and $R2+$.*

In Example 3, $(P_t)_{t \in T}$ is classical, and generated precisely by the example’s true probability measure π (restricted to $\cup_{t \in T} \{A^* : A \in 2^{S_t}\}$).

Turning to classical utilities, imagine in a context t the agent conceives the outcome $\{x, y\}$, which has two underlying objective outcomes; $\{x, y\}$ might stand for ‘rich’, x for ‘very rich’, and y for ‘moderately rich’. Under the classical expected-utility paradigm, composite prospects are systematically evaluated in terms of the

expected/average utility of their subprospects. Accordingly, the prospect $\{x, y\}$ is to be evaluated in terms of the expected/average utility of x and y . This requires assigning utilities and probabilities to *objective* outcomes. Invoking probabilities at the outcome level is unconventional – but it is natural and necessary if coarse outcomes should be evaluated classically by the expected utility of the objective (sub)outcomes. This motivates the following definition:

Definition 11 A family $(U_t)_{t \in T}$ of utility functions U_t on X_t is **classical** if the utility of any outcome is the expected utility (given that outcome) of the objective outcome, i.e., if

$$U_t(x) = \mathbb{E}_\mu(V|x) \text{ for all contexts } t \in T \text{ and outcomes } x \in X_t, \quad (1)$$

relative to some fixed utility function of objective outcomes $V : \mathbb{X} \rightarrow \mathbb{R}$ and some fixed probability measure μ on (some algebra on) \mathbb{X} .¹⁶

Remark 7 If $(U_t)_{t \in T}$ is classical, then it is stable (so obeys R1+), as it is given by the single function U mapping $x \in \cup_{t \in T} X_t$ to $\mathbb{E}_\mu(V|x)$ (with V and μ as in Definition 11).

Classicality of utilities goes far beyond stability, so beyond the rule R1+. It implies a notable property (of which R1+ is the special case in which $n = 1$):

Remark 8 Classicality of $(U_t)_{t \in T}$ implies that whenever the context changes from t to t' and a new outcome $x \in X_{t'}$ is partitionable into old (sub)outcomes $x_1, \dots, x_n \in X_t$, then the new utility of x is a weighted average of the old utilities of x_1, \dots, x_n : $U_{t'}(x) = \sum_{i=1}^n \mu(x_i|x)U_t(x_i)$, with μ as in Definition 11.

Definition 12 If an EU rationalization $(U_t, P_t)_{t \in T}$ is classical both in its utility functions (given by V and μ) and its probability functions (given by π), then $(U_t, P_t)_{t \in T}$ – or in short (V, μ, π) – is a **classical EU rationalization**.¹⁷

Remark 9 Classical EU rationalizations $(U_t, P_t)_{t \in T}$ are stable (hence unified), by Remarks 5, 6 and 7.

Classical EU rationalizations perform a universal reduction to the objective level. Yet, as mentioned, it is hard to imagine how an agent with limited awareness could evaluate outcomes and states *as if* he were aware of underlying objective

¹⁶In particular, for each $x \in \cup_{t \in T} X_t$, $\mathbb{E}_\mu(V|x)$ is well-defined, meaning that (1) $\mu(\cdot|x)$ is well-defined, i.e., μ is defined and non-zero at x , and (2) V has a finite expectation w.r.t. $\mu(\cdot|x)$ (e.g., V is bounded on x and measurable w.r.t. the algebra on which μ is defined).

¹⁷One might consider EU rationalizations $(U_t, P_t)_{t \in T}$ which are classical only in utilities (short-hand: $(V, \mu, (P_t)_{t \in T})$) or only in probabilities (short-hand: $((U_t)_{t \in T}, \pi)$).

outcomes and states. In defence of classical EU rationalizations, one might try to interpret them in a more literal and less *as-if*-like way. For that purpose imagine all unawareness is non-radical: the agent has *in principle* mental access to objective outcomes and states, but for some reason (such as simplicity) conceives outcomes and states more coarsely than he could have. Yet when assigning utilities/probabilities to his coarse outcomes/events he suddenly becomes sophisticated and goes down to objective outcomes/states. I leave it open whether this back-and-forth between a coarse and a fine perspective is psychologically plausible – and if not, whether an ‘as if’ interpretation is plausible. Figure 3 summarizes the

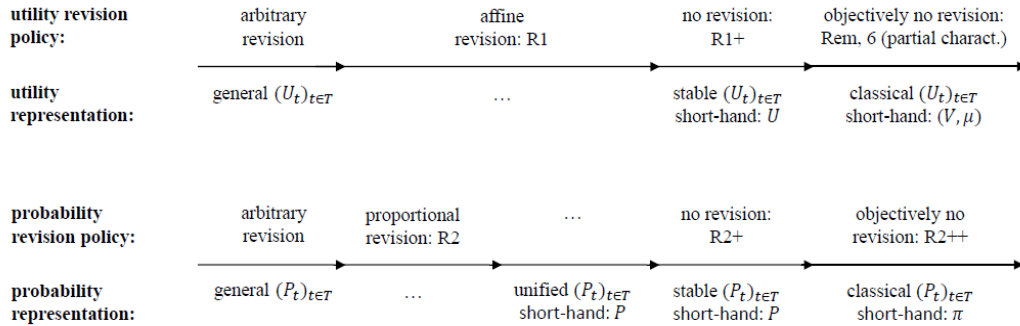


Figure 3: Revision policies for utilities/probabilities and corresponding representations of utilities/probabilities, from most general (left) to most specific (right)

various revision policies with their corresponding representations, in the order of increasing specificity. I insert ‘...’ where a revision policy has no corresponding representation or where a representation has no corresponding revision policy.

2.7 Excursion: unawareness, choice behaviour, and non-representable options

Our model is easily connected to choice behaviour. Assume the agent finds himself in a context $t \in T$ 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_t *or* as objective acts in \mathbb{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_t , and the prediction is simply that a most \succsim_t -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 \mathbb{F} , not F_t . Which choice does \succsim_t predict? It predicts that the agent chooses a feasible objective act whose *subjective representation* in F_t (see Definition 4) is most \succsim_t -preferred. More precisely, our

variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ predicts that whenever in a context $t \in T$ the agent has to choose from a set $\mathcal{F} \subseteq \mathbb{F}$ of representable objective acts, then he chooses an $f \in \mathcal{F}$ such that $f_t \succsim_t g_t$ for all $g \in \mathcal{F}$. (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 states $(X_t, S_t)_{t \in T}$ – **adaptive (to feasible options)** if for each context t all objective acts that can be feasible in t are representable in t .¹⁸ The idea is that the agent forms awareness of a coin toss *when and because* some feasible 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 concepts/awareness states.¹⁹

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’*).²⁰

3 Six Savage-inspired axioms

Sections 3–5 temporarily assume *exhaustive states* (see Definition 3). In fact, each ‘axiom’, ‘theorem’ or ‘proposition’, and most ‘definitions’ and ‘remarks’, also apply

¹⁸A full-fledged definition could state as follows. Let *choice situations* be pairs (\mathcal{F}, t) of a non-empty menu $\mathcal{F} \subseteq \mathbb{F}$ of (feasible) objective acts and a context or ‘frame’ $t \in T$ (in which the choice from \mathcal{F} 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 $(\mathcal{F}, t) \in \mathcal{CS}$ each $f \in \mathcal{F}$ is representable in context t .

¹⁹The agent’s awareness state (X_t, S_t) 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 example at the end of Section 2.2, health features are absent from $X_t = \{rich, poor\}$, though presumably relevant). (2) States may be too coarse (given the outcome space X_t) for all feasible objective acts to be representable (in the mentioned example, f is not representable if $S_t = \{\mathbb{S}\}$, given that $X_t = \{rich, poor\}$). In (1) and (2) I assumed exhaustive states and outcomes, but the idea can be generalized.

²⁰Under 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_t -preferred). Under the second hypothesis, a non-representable feasible option f in \mathbb{F} is not ignored, but (mis)perceived as some subjective act in F_t which fails to properly represent f . *Which* is this subjective act? Here one would need to develop a theory of misrepresentation.

to non-exhaustive states. The three exceptions – two definitions and one remark – will be marked by ‘exh’. So ‘Definition 20_{exh}’ applies only under exhaustive states, but ‘Definition 13’ applies generally. For each exception (marked by ‘exh’), a general re-statement is given in Section 6.

I now state six axioms which reduce to Savage’s axioms in the one-context case.

Standard notation: Let f_A be the restriction of function f to subdomain A . For any object x and set 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 $f_A g_{S_t \setminus A} \in F_t$, where $f, g \in F_t$ and $A \subseteq S_t$ ($t \in T$).

A background assumption: Henceforth let the structure $(X_t, S_t, \succsim_t)_{t \in T}$ satisfy *independence between outcome and state awareness*, so that the agent’s outcome awareness and state awareness do not constrain one another. Formally: any occurring outcome and state spaces X_t and $S_{t'}$ ($t, t' \in T$) can occur *jointly*, i.e., some context $t'' \in T$ has $X_{t''} = X_t$ and $S_{t''} = S_{t'}$.²¹

I begin with the analogue of Savage’s first axiom:

Axiom 1 (weak order): For all contexts $t \in T$, \succsim_t is a transitive and complete relation (on F_t).

Savage’s sure-thing principle requires that the preference between two acts only depends on the acts’ outcomes at those states where they differ. This famous postulate can be rendered in two ways in our setting, by applying sure-thing reasoning either only *within* each context, or even *across* contexts:

Axiom 2* (sure-thing principle, local version): For all contexts $t \in T$, acts $f, g, f', g' \in F_t$, and events $A \subseteq S_t$, if $f_A = f'_A$, $g_A = g'_A$, $f_{S_t \setminus A} = g_{S_t \setminus A}$ and $f'_{S_t \setminus A} = g'_{S_t \setminus A}$, then $f \succsim_t g \Leftrightarrow f' \succsim_t g'$.

Axiom 2 (sure-thing principle, global version): For all contexts $t, t' \in T$, acts $f, g \in F_t$ and $f', g' \in F_{t'}$, and events conceived in both contexts $A \subseteq S_t \cap S_{t'}$, if $f_A = f'_A$, $g_A = g'_A$, $f_{S_t \setminus A} = g_{S_t \setminus A}$ and $f'_{S_{t'} \setminus A} = g'_{S_{t'} \setminus A}$, then $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$.

Remark 10 *Axiom 2* is the restriction of Axiom 2 to the case $t = t'$.*

How does Axiom 2 go beyond Axiom 2*? The preference between the two acts is insensitive not just to the outcomes outside A , but also to the conception/awareness of states outside A , since $S_t \setminus A$ can differ from $S_{t'} \setminus A$. If two acts agree when it doesn’t rain, then it does not matter whether the agent conceives

²¹This excludes agents who conceive the outcome ‘I am popular’ only jointly with the state ‘I win in the lottery’, or who conceive fine states only jointly with coarse outcomes.

just one coarse ‘non-rainy state’ or 17 fine ‘non-rainy states’. Axiom 2 thus applies sure-thing reasoning all the way through, regardless of barriers of context, i.e., of the concept/awareness of irrelevant states.²²

Axiom 2 is decomposable into two axioms, namely Axiom 2* and a new axiom which requires the preference between acts f and g to be unchanged whenever the agent *reconceives* (e.g., refines or coarsens) states at which f and g coincide. The reconception of states of course leads f and g to be recast as acts f' and g' defined on the new state space. Formally:²³

Axiom 2:** For all contexts $t, t' \in T$ with set of common states $A := S_t \cap S_{t'}$, if two acts $f, g \in F_t$ coincide on $S_t \setminus A$ where they yield a constant outcome $x \in X_t \cap X_{t'}$, then $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$, where f' and g' denote the acts in $F_{t'}$ which respectively match f and g on A and both yield x on $S_{t'} \setminus A$.

To paraphrase Axiom 2**, the preference between f and g does not change as the states on which f and g coincide are reconceived, so that f and g become f' and g' . While in Axiom 2** f' and g' are the direct counterparts of f and g in context t' , in the sure-thing principle (Axiom 2 or 2*) f' and g' are by no means equivalent to f and g : their outcomes may have changed outside A .

Remark 11 *Axiom 2** is the special case of Axiom 2 in which $A = S_t \cap S_{t'}$ and in which $f_{S_t \setminus A}$ and $f'_{S_{t'} \setminus A}$ (and thus $g_{S_t \setminus A}$ and $g'_{S_{t'} \setminus A}$) all generate a same constant outcome.*

Compare the reasoning underlying Axiom 2** with the classical sure-thing reasoning underlying Axiom 2*. In both cases, states on which f and g coincide are deemed irrelevant to the preference between f and g , yet in two different senses: *either* the outcomes at these states do not matter (Axiom 2*) *or* the awareness of these states does not matter (Axiom 2**). The two modes of reasoning are complementary. Together they yield Axiom 2:

Remark 12 *Axioms 2* and 2** are jointly equivalent to Axiom 2.*²⁴

It is debatable whether Axiom 2 or 2* is the ‘right’ or ‘natural’ rendition of Savage’s sure-thing principle in our framework of changing awareness. Axiom 2 builds in additional ‘rationality’ in the form of evaluative consistency across contexts; if such consistency is not viewed as an integral part of sure-thing reasoning,

²²Replacing sure-thing reasoning by ambiguity aversion in our setting is an interesting avenue.

²³Axiom 2** is comparable to Karni-Viero’s (2013) *awareness consistency* axiom.

²⁴Why do Axioms 2* and 2** jointly imply Axiom 2? Let t, t', f, g, f', g' obey Axiom 2’s premises. To show that $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$, fix an $x \in X_t \cap X_{t'}$. Applying Axiom 2* on each side, the claimed equivalence reduces to $f_A x_{S_t \setminus A} \succsim_t g_A x_{S_t \setminus A} \Leftrightarrow f'_A x_{S_{t'} \setminus A} \succsim_{t'} g'_A x_{S_{t'} \setminus A}$, which holds by Axiom 2**. (Here A is a subset of $S_t \cap S_{t'}$, the set denoted ‘ A ’ in Axiom 2**.)

Axiom 2* is presumably the right rendition of sure-thing reasoning. By contrast, Axiom 2 is the right rendition if one construes sure-thing reasoning as reasoning which compares acts systematically and solely based on their outcomes *where they differ*, so that the preference between any acts f and g is *determined* by (or, as philosophers say, *supervenes on*) their restrictions f_A and g_A to the ‘disagreement domain’ $A := \{a \in S_t : f(a) \neq g(a)\}$.²⁵

I now extend four familiar Savagean notions to our setting:

Definition 13 (preferences over outcomes) *In a context $t \in T$, an outcome $x \in X_t$ is **weakly preferred** to another $y \in X_t$ – written $x \succsim_t y$ – if $x_{S_t} \succsim_t y_{S_t}$ (recall that x_{S_t} and y_{S_t} are constant acts defined on the state space S_t).*

Definition 14 (conditional preferences) *In a context $t \in T$, an act $f \in F_t$ is **weakly preferred** to another $g \in F_t$ **given** an event $A \subseteq S_t$ – written $f \succsim_{t,A} g$ – if $f' \succsim_t g'$ for some (hence under Axiom 2 any) acts $f', g' \in F_t$ which agree respectively with f and g on A and agree with each other on $S_t \setminus A$.*

Definition 15 (conditional preferences over outcomes) *In a context $t \in T$, an outcome $x \in X_t$ is **weakly preferred** to another $y \in F_t$ **given** an event $A \subseteq S_t$ – written $x \succsim_{t,A} y$ – if $x_{S_t} \succsim_{t,A} y_{S_t}$.*

Definition 16 (null events) *In a context $t \in T$, an event $A \subseteq S_t$ is **null** if it does not affect preferences, i.e., $f \sim_t g$ whenever acts $f, g \in F_t$ agree outside A .*

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

Axiom 3 (state independence): For all contexts $t \in T$, outcomes $x, y \in X_t$, and non-null events $A \subseteq S_t$, $x \succsim_{t,A} y \Leftrightarrow x \succsim_t y$.

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 ; the rationale is that such preferences are driven exclusively by the agent’s assessment of the relative likelihood of the events on which bets are taken. Savage’s axiom can again be rendered as an *intra-* or *inter-*context condition:

Axiom 4* (comparative probability, local version): For all contexts $t \in T$, events $A, B \subseteq S_t$, and outcomes $x \succ_t y$ and $x' \succ_t y'$ in X_t , $x_{A}y_{S_t \setminus A} \succsim_t x_{B}y_{S_t \setminus B} \Leftrightarrow x'_{A}y'_{S_t \setminus A} \succsim_t x'_{B}y'_{S_t \setminus B}$.

²⁵Such supervenience amounts to the existence of a fixed binary relation over ‘subacts’ \triangleright ($\subseteq \cup_{t \in T} \cup_{A \subseteq S_t} (X_t^A \times X_t^A)$) such that, for all contexts $t \in T$ and acts $f, g \in F_t$, $f \succsim_t g \Leftrightarrow f_A \triangleright g_A$, where $A := \{a \in S_t : f(a) \neq g(a)\}$. This is in turn equivalent to Axiom 2.

Axiom 4 (comparative probability, global version): For all contexts $t, t' \in T$ with same state space $S := S_t = S_{t'}$, events $A, B \subseteq S$, and outcomes $x \succ_t y$ in X_t and $x' \succ_{t'} y'$ in $X_{t'}$, $x_{A|S_t \setminus A} \succsim_t x_{B|S_t \setminus B} \Leftrightarrow x'_{A|S_{t'} \setminus A} \succsim_{t'} x'_{B|S_{t'} \setminus B}$.

Remark 13 *Axiom 4* is the restriction of Axiom 4 to the case that $t = t'$.*

I shall use Axiom 4 rather than 4*. Axiom 4 applies the reasoning underlying Savage’s fourth axiom across barriers of context. Yet Axiom 4 is only a ‘mildly global’ axiom, since its contexts t and t' share the same state concept; only the outcome concept may change. Under an interesting condition, Axioms 4 and 4* are actually equivalent (given other axioms of our key theorem); so we could work simply with Axiom 4* if we sacrifice some generality. The condition in question is, informally, that the variable Savage structure offers sufficient flexibility for conceiving/inventing new outcomes.²⁶

Like for the sure-thing principle, so for Savage’s comparative-probability postulate it is debatable whether Axiom 4 or 4* is the more faithful translation into our framework of changing awareness. Axiom 4* seems more faithful to Savage if any demand of cross-context consistency is viewed as ‘orthogonal’ to the principle underlying Savage’s postulate. By contrast Axiom 4 seems more faithful if one interprets Savage’s principle as requiring that the preference between two bets is fully determined by (‘supervenes on’) the pair of underlying events.²⁷

Another familiar notion can now be imported into our setting:

Definition 17 (comparative beliefs) *In a context $t \in T$, an event $A \subseteq S_t$ is **at least as probable as** another $B \subseteq S_t$ – written $A \succsim_t B$ – if $x_{A|S_t \setminus A} \succsim_t x_{B|S_t \setminus B}$ for some (hence under Axiom 4 any) outcomes $x \succ_t y$ in X_t .*

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

Axiom 5 (non-triviality): For all context $t \in T$, there are acts $f \succ_t g$ in F_t .

Axiom 6* (Archimedean, local version): For all contexts $t \in T$, acts $f \succ_t g$ in F_t , and outcomes $x \in X_t$, one can partition S_t into events A_1, \dots, A_n such that

$$f_{S_t \setminus A_i} x_{A_i} \succ_t g \text{ and } f \succ_t g_{S_t \setminus A_i} x_{A_i} \text{ for all } A_i.$$

²⁶Formally, the condition is this: for any two contexts $t, t' \in T$ there are contexts $\hat{t}, \hat{t}' \in T$ such that (i) $X_t \subseteq X_{\hat{t}}$, (ii) $X_{t'} \subseteq X_{\hat{t}'}$, and (iii) $X_{\hat{t}} \cap X_{\hat{t}'}$ contains at least two outcomes (which are non-indifferent in context \hat{t} and/or \hat{t}'). Think of \hat{t} and \hat{t}' as contexts in which additional outcomes have been invented/conceived compared to t resp. t' , such that \hat{t} and \hat{t}' share at least two (non-indifferent) outcomes. In Appendix B I show that under this condition Axioms 4 and 4* are equivalent given Axioms 1, 2 and 6.

²⁷Such supervenience amounts to the existence of a binary relation over subjective events \supseteq ($\subseteq \cup_{t \in T} (2^{S_t} \times 2^{S_t})$), interpretable as an ‘at least as likely’ relation, such that, for all contexts t , events $A, B \subseteq S_t$ and reference outcomes $x \succ_t y$ in X_t , betting on A is weakly preferred to betting on B (i.e., $x_{A|S_t \setminus A} \succsim_t x_{B|S_t \setminus B}$) if and only if $A \supseteq B$. This is equivalent to Axiom 4.

Just as Savage’s 6th postulate, Axiom 6* is very demanding. It forces the agent to conceive plenty of small events, ultimately forcing all state spaces S_t to be infinite (assuming Axiom 5 for non-triviality). I shall thus use a cognitively less demanding Archimedean axiom, which permits all state spaces S_t to be finite. To avoid ‘state-space explosion’, it allows the events A_1, \dots, A_n to be *not yet* conceived in context t : they are conceived in *some* possibly different context t' . So the agent can presently have limited state awareness, as long as states are refinable by moving to a new context/awareness. The slogan is: ‘*refinable* rather than (already) *refined* states’. To refine states, it suffices to incorporate new contingencies into states until a sufficiently fine partition exists; one might incorporate the results of three independent tosses of a fair dice.²⁸ The next axiom formalises the idea of refinable states.²⁹

Definition 18 Acts $f \in F_t$ and $g \in F_{t'}$ ($t, t' \in T$) are (**objectively**) **equivalent** if they induce the same function of objective states, i.e., $f^* = g^*$.

Definition 19 A partition S **refines** or is **at least as fine as** a partition S' if, for some equivalence relation on S , $S' = \{\cup_{a \in E} a : E \text{ is an equivalence class}\}$.³⁰

Axiom 6 (Archimedean, first global version):** For all contexts $t \in T$, acts $f \succ_t g$ in F_t , and outcomes $x \in X_t$, there is a context $t' \in T$ with a state space $S_{t'}$ at least as fine as S_t and an outcome space $X_{t'} \supseteq X_t$ (ensuring that $F_{t'}$ contains acts f' equivalent to f and g' equivalent to g) such that one can partition $S_{t'}$ into events A_1, \dots, A_n for which

$$f'_{S_{t'} \setminus A_i} x_{A_i} \succ_{t'} g' \text{ and } f' \succ_{t'} g'_{S_{t'} \setminus A_i} x_{A_i} \text{ for all } A_i.$$

Remark 14 Axiom 6* is the restriction of Axiom 6** to the case that $t = t'$.

Axiom 6** is not yet satisfactory. It fails to ensure any connection between $\succ_{t'}$ and \succ_t , allowing even that $g \succ_{t'} f$ although $f \succ_t g$. I thus use a variant

²⁸Here each refined state describes an ‘old’ state *and* a 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 $i, j, k \in \{1, 2, \dots, 6\}$.

²⁹‘Refinability’ of states in a context t is for us an existential notion: the occurrence of suitably refined states in *some* context $t' \in T$. Under a stronger reading, states are ‘refinable’ if *the agent can* refine them, i.e., bring about the finer state concept – an ability of which he makes no use in context t , but makes use in a context t' with finer states. Is the idea of *unrefined yet refinable* states coherent even if we read into ‘refinability’ an ability of the agent to refine? We seem to get dangerously close to ‘unawareness but awareness’. Recall however Section 1’s distinction between radical and non-radical unawareness. The idea that in a context t the agent ‘can’ (is ‘able’ to) refine his state space S_t , say into $S_{t'}$, is meaningful provided his initial unawareness of whatever he will bring to his awareness is of the non-radical kind: he ‘can’ bring a coin toss to his awareness only if initially he is *not radically* unaware it, but merely does *not consider* it.

³⁰In other words, S' **coarsens** or is **at least as coarse as** S .

of Axiom 6**, which indirectly guarantees a connection. It requires the objective events represented by A_1, \dots, A_n to be of a certain innocuous kind. Informally, they must belong to an algebra \mathcal{R} of *risky* objective events, e.g., roulette events or coin flipping events. Before turning to formal details, I give an example of the risky algebra \mathcal{R} , preceded by a natural definition.

Definition 20_{exh} *In a context $t \in T$, an objective event $A \subseteq \mathbb{S}$ is (**subjectively**) **representable** if it corresponds to some subjective event, which is then called its (**subjective**) **representation** and denoted $A_t (= \{a \in S_t : a \subseteq A\})$.*

The objective event $\{c, a, b\} \subseteq \mathbb{S}$ might be represented by $\{\{c, a\}, \{b\}\} \subseteq S_t$ in a context t , and by $\{\{c, a, b\}\} \subseteq S_{t'}$ in a context t' , while being non-representable in a context t'' in which the agent lacks appropriate state awareness.

Example of risky algebra \mathcal{R} . Let the objective state space be a Cartesian product $\mathbb{S} = \mathbb{S}_1 \times \mathbb{S}_2$ of a set \mathbb{S}_1 of ‘risky objective states’ (e.g., coin-tossing sequences) and a set \mathbb{S}_2 of ‘non-risky objective states’ (e.g., weather states). Let \mathcal{R} be an algebra on \mathbb{S} consisting of objective events *about the risky objective state* (e.g., about coin tosses). That is, $\mathcal{R} = \{A \times \mathbb{S}_2 : A \subseteq \mathbb{S}_1\}$ (or more generally $\mathcal{R} = \{A \times \mathbb{S}_2 : A \in \mathcal{R}'\}$ for some algebra \mathcal{R}' on \mathbb{S}_1). In each context t let the agent have a certain awareness of risky states given by a partition $S_{t,1}$ of \mathbb{S}_1 and a certain awareness of non-risky states given by a partition $S_{t,2}$ of \mathbb{S}_2 .³¹ The overall state space is $S_t = \{a \times b : a \in S_{t,1}, b \in S_{t,2}\}$. So states $a \times b$ in S_t have a risky part a and a non-risky part b .³² Let there be a ‘true’ objective probability measure π on $2^{\mathbb{S}_1}$ (or more generally on the above-mentioned algebra \mathcal{R}' on \mathbb{S}_1). I also regard π as a function on the ‘risky’ algebra \mathcal{R} , via $\pi(A \times \mathbb{S}_2) := \pi(A)$ for all $A \times \mathbb{S}_2 \in \mathcal{R}$. Let the agent admit an EU rationalization $(U_t, P_t)_{t \in T}$. Let his beliefs about *risky* events always be correct: if an $R = A \times \mathbb{S}_2 \in \mathcal{R}$ is representable in a context t , the representation $R_t \subseteq S_t$ gets the ‘true’ probability $P_t(R_t) = \pi(R)$ (note that $R_t = \{a \times b \in S_t : a \subseteq A\}$). Under this picture, the agent may have limited and changing awareness of risky objective events in \mathcal{R} , but whenever such an event is representable the agent assigns the true and stable probability. For instance, the agent needs not conceive the risky contingency that a fair coin lands heads in the first 10 tosses, but whenever he does, he assigns the true probability of $(\frac{1}{2})^{10}$.

Instead of exogenously fixing a ‘risky’ algebra \mathcal{R} , I use two criteria for when an algebra \mathcal{R} on \mathbb{S} can count as ‘risky’: \mathcal{R} is *robust* and consists of *incorporable* objective events. Before defining ‘robust’ and ‘incorporable’, I should anticipate the definitive statement of the sixth axiom:

³¹One might define contexts as *being* (not just inducing) pairs of finite partitions of \mathbb{S}_1 resp. \mathbb{S}_2 . Then T is some set of such partition pairs $t = (S_1, S_2)$, where $S_{t,1} := S_1$ and $S_{t,2} := S_2$.

³²If one wished to allow for non-exhaustive state awareness (a case currently set aside), then one should let the $S_{t,2}$ s be partitions of arbitrary non-empty subsets of \mathbb{S}_2 .

Axiom 6 (Archimedean, second global version): There is a robust algebra \mathcal{R} of incorporable objective events (a ‘risky’ algebra) such that, for all contexts $t \in T$, acts $f \succ_t g$ in F_t , and outcomes $x \in X_t$, there is a partition of \mathbb{S} into some $R_1, \dots, R_n \in \mathcal{R}$ and a context $t' \in T$ with state space $S_{t'} = S_t \vee \{R_1, \dots, R_n\}$ (ensuring that each R_i is representable by an $A_i \subseteq S_{t'}$) and an outcome space $X_{t'} \supseteq X_t$ (ensuring that $F_{t'}$ contains acts f' equivalent to f and g' equivalent to g), such that

$$f'_{S_{t'} \setminus A_i} x_{A_i} \succ_{t'} g' \text{ and } f' \succ_{t'} g'_{S_{t'} \setminus A_i} x_{A_i} \text{ for all } A_i.$$

Axiom 6 allows that $t = t'$; then R_1, \dots, R_n are already representable in context t . I now gradually build up the axiom’s terminology. I start with a familiar notion:

Definition 21 *If S and S' are partitions (in the generalized sense of possibly containing \emptyset), their join is $S \vee S' := \{a \cap a' : a \in S, a' \in S'\} \setminus \{\emptyset\}$.*³³

An algebra³⁴ \mathcal{R} on \mathbb{S} is *robust* if the ranking of \mathcal{R} -determined acts is stable:

Definition 22 *For an algebra \mathcal{R} on \mathbb{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} .*

Remark 15_{exh} *An act f is \mathcal{R} -determined (given an algebra \mathcal{R} on \mathbb{S}) if and only if f^* is \mathcal{R} -measurable.*³⁵

Definition 23_{exh} *An algebra \mathcal{R} on \mathbb{S} is **robust** if, for all contexts $t, t' \in T$, we have $f \succ_t g \Leftrightarrow f' \succ_{t'} g'$ whenever $f \in F_t$ and $f' \in F_{t'}$ are equivalent \mathcal{R} -determined acts and $g \in F_t$ and $g' \in F_{t'}$ are also equivalent \mathcal{R} -determined acts.*

Robustness is natural if \mathcal{R} contains *risky* objective events, so that \mathcal{R} -determined acts are *risky acts*, because the agent plausibly has fixed ‘preferences under risk’. The idea is that a *risky* objective event gets the same subjective probability regardless of the state space S_t in which it is represented, namely it gets the *objective* probability. 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 of \mathcal{R} .³⁶

³³For now we only need joins of two partitions of \mathbb{S} . But in principle S and S' could partition *different* sets – a case needed later when allowing non-exhaustive states.

³⁴ \mathcal{R} is an algebra on \mathbb{S} if (a) $\mathbb{S} \in \mathcal{R}$, (b) $A \in \mathcal{R} \Rightarrow \overline{A} \in \mathcal{R}$, and (c) $A, B \in \mathcal{R} \Rightarrow A \cup B \in \mathcal{R}$.

³⁵ \mathcal{R} -measurability of f^* means that $(f^*)^{-1}(x) \in \mathcal{R}$ for all outcomes x of f^* , i.e., of f .

³⁶The idea of stable preferences over risky acts parallels Karni-Viero’s (2013) *invariant risk preferences* axiom. The main difference is that our Savagean approach has no exogenous notion of risky acts. For us invariance of preferences is not an *axiomatic requirement* on risky acts, but a *criterion for interpreting* certain objective events (and the acts that depend on them) as ‘risky’.

Example of risky algebra \mathcal{R} – continued. As noted, members of this \mathcal{R} indeed get the same (true) probability *whenever representable*. This makes \mathcal{R} robust, assuming from now on that the utility functions U_t ($t \in T$) obey revision rule R1 (or even R1+) to guarantee cross-context consistency in outcome evaluations. To show robustness, consider contexts t, t' . By R1 we may assume w.l.o.g. that U_t and $U_{t'}$ coincide on their domain overlap $X_t \cap X_{t'}$ (otherwise rescale $U_{t'}$). It suffices to consider equivalent \mathcal{R} -determined acts $f \in F_t$ and $f' \in F_{t'}$ and show that $\mathbb{E}_{P_t} U_t(f) = \mathbb{E}_{P_{t'}} U_{t'}(f')$. The claimed equality reduces to $\mathbb{E}_\pi U_t(f^*) = \mathbb{E}_\pi U_{t'}(f'^*)$, where both expectations can refer to the measure π on \mathcal{R} as f and f' are \mathcal{R} -determined. This holds as $f^* = f'^*$ and as U_t and $U_{t'}$ coincide on $X_t \cap X_{t'}$.

Definition 24 A preference relation $\succsim_{t'}$ is **faithful to** another \succsim_t ($t, t' \in T$) if it preserves all comparisons made by \succsim_t : given any acts $f, g \in F_t$, we have $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$ for some (unique) acts $f' \in F_{t'}$ equivalent to f and $g' \in F_{t'}$ equivalent to g .

If $\succsim_{t'}$ is faithful to \succsim_t , any act in F_t is equivalent to one in $F_{t'}$. So in context t' the agent must conceive all previous outcomes and at least as fine states:

Remark 16 If $\succsim_{t'}$ is faithful to \succsim_t , then (a) $S_{t'}$ is at least as fine as S_t (assuming $|X_t| > 1$), and (b) $X_{t'} \supseteq X_t$ (hence $X_{t'} = X_t$ 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 25 An objective event $A \subseteq \mathbb{S}$ is **incorporable** if it is always representable after (if needed) a preference-neutral state refinement: for every context $t \in T$ there is a context $t' \in T$ (possibly equal to t) such that $S_{t'}$ refines S_t to make A representable, i.e., $S_{t'} = S_t \vee \{A, \bar{A}\}$, and $\succsim_{t'}$ is faithful to \succsim_t .

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.

Example of risky algebra \mathcal{R} – continued. Here \mathcal{R} indeed consists of incorporable objective events in two stylised cases:

1. *Stable and full risk awareness*: the agent always has full awareness of the risky state, i.e., each risky space $S_{t,1}$ ($t \in T$) is the finest partition $\{\{a\} : a \in \mathbb{S}_1\}$. Then all $A \times \mathbb{S}_2 \in \mathcal{R}$ are automatically *representable in all contexts* – a special case of incorporability.

2. *Unstable and finite risk awareness*: the agent can have *any finite* risk awareness, i.e., the risky spaces $S_{t,1}$ ($t \in T$) are *all the finite* partitions of \mathbb{S}_1 .³⁷ I add four auxiliary assumptions (which could be weakened). First, all non-risky spaces $S_{t,2}$ are also finite. Second, risky and non-risky states are probabilistically independent: for all $a \times b \in S_t$ ($t \in T$), $P_t(a \times b) = P_t(a)P_t(b)$, where $P_t(a)$ and $P_t(b)$ denote a 's and b 's marginal probabilities, i.e., $P_t(a) = \sum_{b' \in S_{t,2}} P_t(a \times b')$ and $P_t(b) = \sum_{a' \in S_{t,1}} P_t(a' \times b)$. The motivation is that risky contingencies like tossing outcomes are unrelated with non-risky ones like the weather. Third, I assume independence between risky and non-risky state awareness: any risky space $S_{t,1}$ and non-risky space $S_{t',2}$ ($t, t' \in T$) can occur *jointly*, i.e., some context t'' has $S_{t'',1} = S_{t,1}$ and $S_{t'',2} = S_{t',2}$. Fourth, probabilities of non-risky states depend only on awareness of *non-risky states*: whenever $S_{t,2} = S_{t',2}$ then $P_t(b) = P_{t'}(b)$ for all $b \in S_{t,2} = S_{t',2}$. Then each risky objective event $R = A \times \mathbb{S}_2 \in \mathcal{R}$, whilst not always representable, is *incorporable*. Why? Consider a context t . By assumption we may pick a context t' such that $S_{t',1} = S_{t,1} \vee \{A, \bar{A}\}$ and $S_{t',2} = S_{t,2}$. Clearly, $S_{t'} = S_t \vee \{R, \bar{R}\}$, so that R is representable in context t' . We may assume that $X_t = X_{t'}$ by independence between outcome and state awareness, and that $U_t = U_{t'} =: U$ by R1 (otherwise rescale $U_{t'}$). To show that $\succsim_{t'}$ is faithful to \succsim_t , take any act $f \in F_t$. Clearly, there is an equivalent act $f' \in F_{t'}$. It suffices to show that $\mathbb{E}_{P_t} U(f) = \mathbb{E}_{P_{t'}} U(f')$. This claim holds because, when moving from t to t' , (i) each state $a \times b$ in S_t splits into ‘substates’ in $S_{t'}$ which partition $a \times b$, namely into all $a' \times b \in S_{t'}$ such that $a' \subseteq a$ (at most two substates exist as a' must be $a \cap A$ or $a \cap \bar{A}$ and cannot be empty), (ii) f' ‘essentially extends’ f , i.e., f' 's value at a state $a \times b \in S_t$ matches f' 's value at each substate $a' \times b \in S_{t'}$, and (iii) $P_{t'}$ ‘essentially extends’ P_t , i.e., P_t assigns to each state $a \times b \in S_t$ the same probability as $P_{t'}$ does to the set of substates of $a \times b$, namely the probability $\pi(a)P_t(b) = \pi(a)P_{t'}(b)$.

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

Remark 17 *In the single-context case $T = \{t_0\}$, the variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ is equivalent to an ordinary Savage structure $(X, S, \succsim) = (X_{t_0}, S_{t_0}, \succsim_{t_0})$, and our axioms reduce to Savage’s axioms, i.e.,*

- (a) *Axiom 1 is equivalent to Savage’s Axiom P1,*
- (b) *Axioms 2 and 2* are each equivalent to Savage’s Axiom P2,*
- (c) *Axiom 3 is equivalent to Savage’s Axiom P3,*
- (d) *Axioms 4 and 4* are each equivalent to Savage’s Axiom P4,*
- (e) *Axiom 5 is equivalent to Savage’s Axiom P5,*

³⁷If \mathcal{R} takes the general form $\{A \times \mathbb{S}_2 : A \in \mathcal{R}'\}$ for an algebra \mathcal{R}' on \mathbb{S}_1 , not the simple form $\{A \times \mathbb{S}_2 : A \subseteq \mathbb{S}_1\}$, then the $S_{t,1}$ s are the finite partitions of \mathbb{S}_1 into members of \mathcal{R}' .

(f) *Axioms 6, 6* and 6** are each equivalent to Savage’s Axiom P6.*³⁸

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, due to framing effects and other empirically studied reasons. Just imagine that in two contexts $t, t' \in T$ the agent chooses between the same objective acts $f, g \in \mathbb{F}$, which he subjectively represents as $f_t, g_t \in F_t$ in context t , and as $f_{t'}, g_{t'} \in F_{t'}$ in context t' (see Definition 5). Then he will choose f in context t if $f_t \succ_t g_t$, but g in context t' if $g_{t'} \succ_{t'} f_{t'}$. 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 26 (*effective preference over objective acts*) *In a context $t \in T$, an objective act $f \in \mathbb{F}$ is (effectively) weakly preferred to $g \in \mathbb{F}$ – written $f \succsim_t g$ – if f and g are representable and the representations satisfy $f_t \succsim_t g_t$.*

The effective preference between $f, g \in \mathbb{F}$ is reversible, as possibly $f \succsim_t g$ but $g \succ_{t'} f$. Similar reversals can happen for *beliefs*. The agent may attach high probability to the event $\{\{a\}, \{b\}\}$ (when conceived), but low probability to the objectively equivalent event $\{\{a, b\}\}$ (when conceived); both events represent the same objective event $\{a, b\}$. In the experimental literature this sort of phenomenon is known under names such as ‘packing/unpacking events’ (e.g., Tversky and Koehler 1994). Formally, we may lift the agent’s comparative beliefs to the objective level, talking then of ‘effective’ beliefs:

Definition 27 (*effective comparative belief about objective events*) *In a context $t \in T$, an objective event $A \subseteq \mathbb{S}$ is (effectively) at least as probable as another one $B \subseteq \mathbb{S}$ – written $A \succsim_t B$ – if A and B are representable and the representations satisfy $A_t \succsim_t B_t$.*

Nothing in Axioms 1–6 prevents an effective belief $A \succsim_t B$ from reversing into $B \succ_{t'} A$.

³⁸Axioms 6* and 6** imply Axioms 6 by letting \mathcal{R} contain *all* representable objective events.

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 1 (preference stability) *Under Axiom 2, acts are ranked the same way wherever conceived: $f \succsim_t g \Leftrightarrow f \succsim_{t'} g$ for all contexts $t, t' \in T$ and acts conceived in both contexts $f, g \in F_t \cap F_{t'}$.*

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 1 has a bite only for those pairs of contexts $t, t' \in T$ for which $F_t \cap F_{t'} \neq \emptyset$, i.e., for which $S_t = S_{t'}$ and $X_t \cap X_{t'} \neq \emptyset$. If awareness varies so drastically that no distinct contexts share any acts, then Proposition 1 is vacuous.

Proposition 2 (outcome-preference stability) *Under Axiom 6, outcomes are ranked the same way wherever conceived: $x \succsim_t y \Leftrightarrow x \succsim_{t'} y$ for all contexts $t, t' \in T$ and outcomes conceived in both contexts $x, y \in X_t \cap X_{t'}$.*

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 $x \in X_t \cap X_{t'}$ is identified with *distinct* constant acts $x_{S_t} \in F_t$ and $x_{S_{t'}} \in F_{t'}$ if $S_t \neq S_{t'}$.

Proposition 3 (comparative-belief stability) *Under Axioms 2, 4, 5 and 6, events are ranked the same way wherever conceived: $A \succsim_t B \Leftrightarrow A \succsim_{t'} B$ for all contexts $t, t' \in T$ and events conceived in both contexts $A, B \subseteq S_t \cap S_{t'}$.*

By contrast, our axioms permit changes in *quantitative* rather than comparative beliefs: the same event can get different probability in different contexts in which it is conceived. This is clear from Theorem 1 below, in which probabilities must obey the revision rule R2, but not the rule R2+ of stable probabilities. An extreme type of probability – *zero probability* – is however stable, for the following reason:

Proposition 4 (stability of nullness) *Under Axiom 2, any event A is null in all or none of the contexts t where it is conceived, i.e., where $A \subseteq S_t$.*

Finally, the comparative-belief stability extends even to certain *objective* events, namely those taken from a robust algebra:

Proposition 5 (stability of comparative belief on robust algebras) *Under Axioms 2, 4 and 5, objective events from a robust algebra \mathcal{R} on \mathbb{S} are ranked the same way wherever representable: $A \succsim_t B \Leftrightarrow A \succsim_{t'} B$ for all contexts $t, t' \in T$ and objective events $A, B \in \mathcal{R}$ representable in both contexts.*

5 The characterization theorem

After introducing the third revision rule (Section 5.1), I state and discuss the main theorem (Section 5.2), and give examples of belief revision (Section 5.3)

5.1 The theorem's third revision rule: enough objectively stable probabilities

While the revision rules R1 and R2 concern utilities/probabilities of *subjective* outcomes/events, the third rule concerns (induced) probabilities of *objective* events. Informally, the third rule requires that enough objective events have stable, i.e., representation-invariant, probabilities; interpretatively, there must be enough (revealed) objective risk. Before defining relevant terms, I anticipate this rule's statement, which refers to the agent's family of probability measures P_t on 2^{S_t} ($t \in T$):

R3: *The objective agreement of the functions P_t ($t \in T$) is fine.*

I now define 'objective agreement' and 'fine'. The P_t s 'agree' on a probability x of an *objective* event A if, roughly speaking, each P_t assigns probability x to A 's representation. The precise definition is more general: it allows A to be not (yet) representable in a context t , in which case the probability x is derived not from P_t itself, but from a version of P_t defined on a refined state space in which A becomes representable. Formally:

Definition 28 *Given a context $t \in T$, any function P on 2^{S_t} induces a function, denoted P^* , on the set of representable objective events, via $P^*(A) := P(A_t)$.*

Definition 29 *Functions P_t on 2^{S_t} ($t \in T$) **agree on value x at objective event A** ($A \subseteq \mathbb{S}$) – or simply **agree at A** – if each P_t induces the value x at A 'modulo extension', i.e., if each induced function P_t^* has an extension $P_{t'}^*$ ($t' \in T$) such that $S_{t'} = S_t \vee \{A, \bar{A}\}$ (so that $P_{t'}^*(A)$ is defined) and $P_{t'}^*(A) = x$. (So $P_t^*(A) = x$ for those contexts t where A is already representable, i.e., for which t' could equal t .)*

Definition 30 *The **objective agreement** of functions P_t on 2^{S_t} ($t \in T$) is the function O on the **objective agreement domain** $\mathcal{A} = \{A \subseteq \mathbb{S} : \text{the } P_t\text{s agree at } A\}$ assigning to each $A \in \mathcal{A}$ the agreed value at A .*

In both cases 1 and 2 of Section 3's continued 'Example of risky algebra \mathcal{R} ', the functions P_t agree on value $\pi(R)$ at each $R \in \mathcal{R}$, and so their objective agreement domain is (or includes) \mathcal{R} and their objective agreement is (or extends) π .

In general, the objective agreement domain \mathcal{A} of subjective probability functions P_t can be as small as $\{\emptyset, \mathbb{S}\}$ or as large as $2^{\mathbb{S}}$. The larger this domain is, the more objective events are assessed context-invariantly.

Remark 18 *The objective agreement O of probability measures P_t on 2^{S_t} ($t \in T$) is a probability measure in the wide sense that its domain \mathcal{A} need not be an algebra (in fact it need not be closed under union, but must be closed under complement and contain \mathbb{S}). That is, $O(\mathbb{S}) = 1$ and O is additive, i.e., $O(A \cup B) = O(A) + O(B)$ if $A, B, A \cup B \in \mathcal{A}$ and $A \cap B = \emptyset$.*

Definition 31 *A probability measure on an algebra \mathcal{A} on \mathbb{S} is **fine** if for all $\epsilon > 0$ one can partition \mathbb{S} into some $A_1, \dots, A_n \in \mathcal{A}$ of probabilities at most ϵ . More generally, a probability measure in the wide sense of an arbitrary domain $\mathcal{A} \subseteq 2^{\mathbb{S}}$ is **fine** if on some algebra $\mathcal{B} \subseteq \mathcal{A}$ the induced ordinary probability measure is fine.³⁹*

5.2 The theorem

The following theorem states the implications of our axioms (see Appendix A for results on tighter rationalizations and revision rules):

Theorem 1 *The variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ satisfies Axioms 1–6 if and only if it has an EU rationalization $(U_t, P_t)_{t \in T}$ satisfying the revision rules R1, R2 and R3. Each P_t is unique and each U_t is unique up to increasing affine transformation.⁴⁰*

Remark 19 *In the single-context case $T = \{t_0\}$, Theorem 1 reduces to Savage’s Theorem for the ordinary Savage structure $(X_{t_0}, S_{t_0}, \succsim_{t_0})$, as Axioms 1–6 reduce to Savage’s Axioms P1–P6 (by Remark 17), rules R1 and R2 hold trivially, and R3 reduces to Savage’s requirement that P_{t_0} be atomless.⁴¹*

Remark 20 *Unlike in Savage’s Theorem, Theorem 1’s representation allows that all state spaces S_t are finite; but it forces the objective state space \mathbb{S} to be infinite (as by R3 the objective agreement domain is infinite). \mathbb{X} can be finite or infinite.⁴²*

Remark 21 *In Theorem 1’s representation, probabilities depend only on state awareness, and utilities only on outcome awareness: if $S_t = S_{t'}$ then $P_t = P_{t'}$, and if $X_t = X_{t'}$ then $U_t = U_{t'}$ up to normalisation. In particular:*

- *if all state spaces coincide, then all P_t coincide,*

³⁹Fineness implies Savage’s atomlessness if \mathcal{A} is an algebra, and is equivalent to atomlessness if \mathcal{A} is even a σ -algebra.

⁴⁰Formally, if $(U_t, P_t)_{t \in T}$ is an EU rationalization satisfying R1–R3, then $(U'_t, P'_t)_{t \in T}$ is also one if and only if, for all $t \in T$, $P'_t = P_t$ and $U'_t = a_t U_t + b_t$ for some $a_t > 0$ and $b_t \in \mathbb{R}$.

⁴¹Indeed, R3 reduces to fineness of P_{t_0} , and so to atomlessness of P_{t_0} by footnote 39.

⁴²As the objective agreement domain is infinite, infinitely many objective events are incorporable. So, although states can be unrefined in a given context, they must be ‘indefinitely refinable’.

- if all outcome spaces coincide, then all U_t coincide up to normalisation.

Revealed objective probabilities. I propose to interpret the objective agreement of the P_t s as capturing (revealed) *objective* probabilities, i.e., *risk*. Accordingly, the objective agreement domain $\mathcal{A} \subseteq 2^{\mathbb{S}}$ contains (revealed) *risky* objective events. Seen this way, R3 requires a sufficient amount of objective rather than subjective uncertainty. Interpreting members of the objective agreement domain \mathcal{A} as ‘risky’ is in line with our earlier move to interpret members of the algebra \mathcal{R} in Axiom 6 as ‘risky’, because $\mathcal{R} \subseteq \mathcal{A}$, as the proof will reveal.

The need to revise both probabilities and utilities. Under our framework and axioms, revision both of probabilities and utilities is an unavoidable consequence of changing awareness. The rules R1–R3 ‘discipline’ the revision. By R2, probability *ratios* are stable, an interesting analogy to Ahn-Ergin’s (2010) ‘partition-dependent probabilities’ and Karni-Viero’s (2013) ‘reverse Bayesianism’. But nothing so far guarantees existence of Section 2.6’s tighter rationalizations with less or no revision. For one, the P_t s need not be of unified (‘Ahn-Ergin’) type⁴³ – let alone stable (rule R2+) or even classical, i.e., objectively stable (rule R2++). Secondly, the functions U_t might have to be non-stable, i.e., violate R1+, regardless of how one scales them via increasing affine transformations. To see why, consider contexts t, t', t'' . After scaling $U_{t'}$ to match U_t on $X_t \cap X_{t'}$, and then scaling $U_{t''}$ to match $U_{t'}$ on $X_{t'} \cap X_{t''}$, $U_{t''}$ might fail to match U_t on $X_t \cap X_{t''}$. This shows the genuine need for utility revision over and above probability revision. However, under certain conditions tighter rationalizations do exist, as shown in Appendix A.2.

How do the revision rules follow from the axioms? Let me give an early hunch for interested readers; details appear in the Appendix. I here take for granted that Axioms 1–6 jointly guarantee an EU rationalization $(U_t, P_t)_{t \in T}$.

- *How does R1 follow?* As in R1, consider contexts t and t' , for simplicity with same state space $S_t = S_{t'}$. By preference stability (Proposition 1), \succsim_t and $\succsim_{t'}$ coincide on $F_t \cap F_{t'}$, the set of functions from $S_t = S_{t'}$ into $X_t \cap X_{t'}$. Meanwhile the (common) restriction of \succsim_t and $\succsim_{t'}$ to $F_t \cap F_{t'}$ admits two expected-utility rationalizations, i.e., $(U_t|_{X_t \cap X_{t'}}, P_t)$ and $(U_{t'}|_{X_t \cap X_{t'}}, P_{t'})$, obtained from the initial rationalizations (U_t, P_t) of \succsim_t and $(U_{t'}, P_{t'})$ of $\succsim_{t'}$ by restricting the utility functions. Since $(U_t|_{X_t \cap X_{t'}}, P_t)$ and $(U_{t'}|_{X_t \cap X_{t'}}, P_{t'})$ represent the same relation, a uniqueness result – shown in the appendix – implies that $U_t|_{X_t \cap X_{t'}}$ must be an increasing affine transformation of $U_{t'}|_{X_t \cap X_{t'}}$, as in R1.

⁴³Indeed, there may exist contexts $t, t', t'' \in T$ and states $a \in S_t \cap S_{t'}, b \in S_{t'} \cap S_{t''}, c \in S_{t''} \cap S_t$ such that $P_t(a) = P_t(c), P_{t'}(a) = P_{t'}(b), P_{t''}(b) \neq P_{t''}(c)$; a unifying function P would then have to satisfy $P(a) = P(c), P(a) = P(b), P(b) \neq P(c)$, a contradiction.

- *Axiom 2 (or just 2**) is crucial for R2.* Let us see why. As in R2 consider contexts t and t' , for simplicity with the same outcome space $X_t = X_{t'}$. Ordinal probability comparisons between events conceived in both states are stable: for all $B, C \in 2^{S_t \cap S_{t'}}$, $B \succsim_t C \Leftrightarrow B \succsim_{t'} C$. This is because bets on B and C (of type $x_B y_{S_t \setminus B}, x_C y_{S_t \setminus C}$ for ‘win’ and ‘lose’ outcomes $x \succ_t y$ in X_t) yield the same ‘lose’ outcome outside $S_t \cap S_{t'}$, so that by Axiom 2** the preference between these bets does not change as we move to context t' and accordingly recast both bets as functions on $S_{t'}$. Thus the probability measures P_t and $P_{t'}$ are *ordinally equivalent* on $2^{S_t \cap S_{t'}}$: for all $B, C \in 2^{S_t \cap S_{t'}}$, $P_t(B) \geq P_t(C) \Leftrightarrow P_{t'}(B) \geq P_{t'}(C)$. Getting from ordinal equivalence to *proportionality* uses further arguments and axioms.
- *Axiom 6 is crucial for R3.* Just as in Savage’s Theorem the 6th axiom is the main key to the richness in small-probability events, i.e., to an atomless probability measure, so in our theorem Axiom 6 is the main key to the richness in small-probability objective events, i.e., to R3.

5.3 Examples of belief revision

Complementing Examples 1–3 about utility and probability revision, I now work out two concrete examples of *probability* revision.

Example: objectively correct probabilities. Consider again Example 3, in which all uncertainty is objective and each probability measure P_t assigns the objectively right probabilities to all currently conceived events; as noted, this implies R2++ (objectively stable beliefs), and so R2 and R2+. How about R3? For concreteness, let the uncertainty come from an infinite sequence of tosses of a fair coin. So $\mathbb{S} = \{0, 1\}^{\mathbb{N}}$, where the i^{th} entry of an objective state $(a_i)_{i \in \mathbb{N}} \in \mathbb{S}$ is 1/0 if the i^{th} toss results in heads/tails. An objective event $A \subseteq \mathbb{S}$ is *finitely complex* if it concerns only a finite subsequence of tosses, i.e., $A = \{(a_i)_{i \in \mathbb{N}} \in \mathbb{S} : (a_i)_{i \in I} \in D\}$ for some finite $I \subseteq \mathbb{N}$ and some set of I -subsequences $D \subseteq \{0, 1\}^I$; so A says that the I -subsequence lies in D (A might say that the 1st and 4th toss coincide, so that $I = \{1, 4\}$ and D contains the 1-1 and 0-0 subsequences). Let \mathcal{R} be the algebra of finitely complex objective events, and π the true probability measure on \mathcal{R} . So, for all $A = \{(a_i)_{i \in \mathbb{N}} \in \mathbb{S} : (a_i)_{i \in I} \in D\}$ in \mathcal{R} , $\pi(A) := \frac{|D|}{2^{|I|}} = \frac{\text{number of } I\text{-subsequences in } D}{\text{total number of } I\text{-subsequences}}$, since all subsequences in $\{0, 1\}^I$ are equally likely.⁴⁴ Let the agent only ever conceive finitely many states, each of which is finitely complex. Formally, the set of state spaces $\{S_t : t \in T\}$ consists of all finite partitions of \mathbb{S} into members of \mathcal{R} . An example is $S_t = \{\text{‘1st toss heads’}, \text{‘1st toss tails \& 2nd toss heads’}, \text{‘1st toss tails \& 2nd toss tails’}\}$. As subjective probabilities match objective ones, $P_t(s) = \pi(s)$ for

⁴⁴The value $\frac{|D|}{2^{|I|}}$ does not depend on the pair (I, D) chosen to represent A (the canonical choice has minimal I).

all contexts t and states $s \in S_t$. Then R3 holds, because the objective agreement of the P_t s is the true probability measure π , which is indeed fine. The objective agreement domain is thus the algebra \mathcal{R} of finitely complex objective events.⁴⁵

Example: both objective and subjective uncertainty. I now take up Section 3's 'Example', with its objective state space $\mathbb{S} = \mathbb{S}_1 \times \mathbb{S}_2$ built from a risky space \mathbb{S}_1 and a non-risky space \mathbb{S}_2 . \mathbb{S}_1 might be the coin-tossing space $\{0, 1\}^{\mathbb{N}}$ of the previous example. Recall that in each context t the agent has some risk awareness given by a partition $S_{t,1}$ of \mathbb{S}_1 and some awareness of non-risky contingencies given by a partition $S_{t,2}$ of \mathbb{S}_2 ; this results in the state space $S_t = \{a \times b : a \in S_{t,1}, b \in S_{t,2}\}$. Recall also that probabilities of *risky* events follow a 'true' measure π defined on some algebra \mathcal{R}' on \mathbb{S}_1 , possibly the maximal one $2^{\mathbb{S}_1}$. So whenever an objective event of risky type $R = A \times \mathbb{S}_2$ is representable in a context t , its representation $R_t (= \{a \times b \in S_t : a \subseteq A\})$ gets the *true* probability: $P_t(R_t) = \pi(A)$. If \mathbb{S}_1 is the coin-tossing space $\{0, 1\}^{\mathbb{N}}$ of the previous example, then π could be that example's true probability measure on the algebra of finitely complex subsets of $\{0, 1\}^{\mathbb{N}}$, or an extension of it to $2^{\mathbb{S}_1} = 2^{\{0,1\}^{\mathbb{N}}}$.⁴⁶ To obtain the belief revision rules R2 and R3, I assume that probabilities of non-risky states are of the following unified (Ahn-Ergin-inspired) type. For any context t I denote the marginal probability of a non-risky state $b \in S_{t,2}$ by $P_{t,2}(b)$ ($= P_t(\{a \times b : a \in S_{t,1}\})$). Let any ever conceived non-risky state $b \in \cup_{t \in T} S_{t,2}$ have a 'plausibility' $\mu(b) > 0$, and let each $P_{t,2}$ ($t \in T$) be μ 's restriction $S_{t,2}$, rescaled proportionally such that the values sum to one. I also make three auxiliary assumptions familiar from Section 3: finiteness of all non-risky spaces $S_{t,2}$, probabilistic independence between risky and non-risky states, and independence between risky and non-risky state awareness.⁴⁷

To show R2, consider contexts t, t' . For simplicity the proof takes the overlap $S_t \cap S_{t'}$ to be finite, so that it suffices to show that P_t and $P_{t'}$ are proportional as functions of *states* in $S_t \cap S_{t'}$ rather than *events* in $2^{S_t \cap S_{t'}}$. This holds because, as functions of states $a \times b \in S_t \cap S_{t'}$, both P_t and $P_{t'}$ are proportional to $\pi(a)\mu(b)$. As for R3, I restrict attention to two cases invoked already in Section 3:

1. *Stable and full risk awareness:* each $S_{t,1}$ is the finest partition $\{\{a\} : a \in \mathbb{S}_1\}$. Here, the P_t s agree at each objective event of the risky type $A \times \mathbb{S}_2$, namely on the value $\pi(A)$. This is obvious, partly because such objective events are

⁴⁵ π is in fact the restriction to \mathcal{R} of the infinite independent product of the uniform Bernoulli measure. This product measure, sometimes denoted $\bigotimes_{i=1}^{\infty} \text{Bernoulli}(\frac{1}{2})$, is defined on a σ -algebra, which is obtainable as the σ -algebra generated by \mathcal{R} , i.e., the closure of \mathcal{R} under *countable* union and complement. This closure contains 'infinitely complex' objective events.

⁴⁶Such an extension exists (via the Hahn-Banach theorem), but is only finitely additive.

⁴⁷The general statement of the probabilistic independence assumption, valid even when S_t is infinite (because $S_{1,t}$ is infinite), is this: for all contexts t and sets $A \subseteq S_{t,1}$ and $B \subseteq S_{t,2}$, $P_t(\{a \times b \in S_t : a \in A, b \in B\}) = P_t(\{a \times b \in S_t : a \in A\})P_t(\{a \times b \in S_t : b \in B\})$.

representable in all contexts.

2. *Instable and finite risk awareness*: the agent can have *any finite* risk awareness, i.e., the risky spaces $S_{t,1}$ are *all finite* partitions of \mathbb{S}_1 (into members of \mathcal{R}'). Then the P_t s agree at each objective event of the risky type $A \times \mathbb{S}_2$ ($A \in \mathcal{R}'$), namely on the value $\pi(A)$. The proof is easy; the additional twist compared to stable and full risk awareness is that such objective events are not always representable, but merely incorporable.

So in both cases R3 holds as long as π is fine, as is for instance the case if \mathbb{S}_1 is the coin-tossing space $\{0, 1\}^{\mathbb{N}}$ and π reflects true tossing probabilities. The objective agreement domain is or includes the algebra $\{A \times \mathbb{S}_2 : A \in \mathcal{R}'\}$, i.e., the ‘risky’ algebra \mathcal{R} of Section 3’s ‘Example’.

6 The general case

I now lift the temporary restriction to exhaustive states. Recall that each above ‘axiom’, ‘theorem’ or ‘proposition’, and most ‘definitions’ and ‘remarks’ continue to apply *as stated*. The three exceptions, namely Definitions 20_{exh} and 23_{exh} and Remark 15_{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. First, if one partition refines another, it follows that both are partitions of the *same* set by Definition 19 (so if S_t refines $S_{t'}$, then $\mathbb{S}_t = \mathbb{S}_{t'}$). Second, the (unchanged) Definitions 25 and 29 and Axiom 6, when applied with non-exhaustive states, require forming the join of partitions of *possibly distinct* sets (i.e., \mathbb{S}_t and \mathbb{S}). This join then partitions the intersection of the two sets (i.e., \mathbb{S}_t), by Definition 21. I now state the three generalizations.

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

Definition 20 *In a context $t \in T$, an objective event $A \subseteq \mathbb{S}$ is (**subjectively**) **representable** if its encompassed part $A \cap \mathbb{S}_t$ corresponds to a subjective event, called then A ’s (**subjective**) **representation**, denoted A_t ($= \{a \in S_t : a \subseteq A\}$).*

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

Remark 15 *A subjective act $f \in F_t$ ($t \in T$) is \mathcal{R} -determined (given an algebra \mathcal{R} on \mathbb{S}) if and only if f^* is \mathcal{R}' -measurable where $\mathcal{R}' = \{A \cap \mathbb{S}_t : A \in \mathcal{R}\}$ is the trace of \mathcal{R} in \mathbb{S}_t .*

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 32 *Two acts $f \in F_t$ and $f' \in F_{t'}$ (where $t, t' \in T$) are **corresponding \mathcal{R} -determined acts** (for an algebra \mathcal{R} on \mathbb{S}) if both are given by an identical \mathcal{R} -measurable function, i.e., there is an \mathcal{R} -measurable function g on \mathbb{S} such that $g(a) = f(\tilde{a})$ whenever $a \in \tilde{a} \in S_t$ and $g(a) = f'(\tilde{a})$ whenever $a \in \tilde{a} \in S_{t'}$ (i.e., such that, $g_{S_t} = f^*$ and $g_{S_{t'}} = f'^*$).*

Definition 23 *An algebra \mathcal{R} on \mathbb{S} is **robust** if, for all contexts $t, t' \in T$, we have $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$ whenever $f \in F_t$ and $f' \in F_{t'}$ are corresponding \mathcal{R} -determined acts, and $g \in F_t$ and $g' \in F_{t'}$ are also corresponding \mathcal{R} -determined acts.*

Definition 23 indeed generalizes Definition 23_{exh}, for the following reason:

Remark 22 *Under exhaustive states, two acts $f \in F_t$ and $f' \in F_{t'}$ ($t, t' \in T$) 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 23 *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 $a \in S_t$ are singleton. So just one kind of state awareness can vary: the level of state exhaustiveness.

Remark 24 *Under 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 \mathbb{S} such that, for all contexts $t \in T$, acts $f \succ_t g$ in F_t , and outcomes $x \in X_t$, one can partition S_t into events $A_1, \dots, A_n \subseteq S_t$ representing objective events from \mathcal{R} such that $f_{S_t \setminus A_i} x_{A_i} \succ_t g$ and $f \succ_t g_{S_t \setminus A_i} x_{A_i}$ for all A_i .

We can also work with a more basic notion than ‘objective agreement’:

Definition 33 *The **objective overlap** of functions P_t on 2^{S_t} ($t \in T$) is the meet (greatest common subfunction) of the functions of objective events P_t^* ($t \in T$).*

Remark 25 *The objective overlap of functions P_t on 2^{S_t} ($t \in T$) is a subfunction (restriction) of their objective agreement, and coincides with it under fine states.*

Remark 26 *The objective overlap of probability measures P_t on 2^{S_t} ($t \in T$) is itself a probability measure, namely the restriction of each P_t^* to the algebra $\{A \subseteq \mathbb{S} : P_t^*(A) \text{ is defined}\}^{48}$ and identical for all $t \in T$.*

Theorem 1’s fine-state corollary follows via Remark 25 and a trivial lemma:

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

Corollary 1 *Under fine states, Axioms 1–5 and $\tilde{6}$ hold if and only if there is an EU rationalization $(U_t, P_t)_{t \in T}$ satisfying R1, R2, and a third revision rule: the objective overlap of the functions P_t is fine.⁴⁹*

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). To achieve this, I introduce an exogenous algebra \mathcal{R} (on \mathbb{S}) of ‘*risky*’ objective events, and replace Axiom 6 by three \mathcal{R} -specific axioms:

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 \mathbb{S} , the variable Savage structure $(X_t, S_t, \succsim_t)_{t \in T}$ satisfies Axioms 1–5 and $6_{\mathcal{R}}-8_{\mathcal{R}}$ if and only if it has an EU rationalization $(U_t, P_t)_{t \in T}$ satisfying R1, R2, and a third revision rule: the objective agreement of the functions P_t includes⁵⁰ a fine (‘objective’) probability measure on \mathcal{R} . Each P_t is unique and each U_t is unique up to increasing affine transformation.*

⁴⁸Definedness is equivalent to representability of A and comes for free under fine states.

⁴⁹Fine states are essentially objective states. So, had this paper focused exclusively on fine states, we could have introduced each S_t as a primitive set (not a partition), redefined the ‘objective state space’ as $\cup_{t \in T} S_t$, and redefined accordingly all concepts that refer to objective states (such as ‘robust algebras’ and ‘objective agreement/overlap’).

⁵⁰A function *includes* another if it is an extension of the other.

Remarks 20 and 21 apply analogously to Theorem 2. To obtain Theorem 2’s fine-state corollary, I drop Axiom $8_{\mathcal{R}}$ (which comes for free), simplify the third revision rule (using Remark 25’s reduction of ‘objective agreement’ to ‘objective overlap’), and simplify Axiom $6_{\mathcal{R}}$ (using a lemma)

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’.

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

Corollary 2 *Under fine states and an exogenous (risky) algebra \mathcal{R} on \mathbb{S} , Axioms 1–5, $\tilde{6}_{\mathcal{R}}$ and $7_{\mathcal{R}}$ hold if and only if there is an EU rationalization $(U_t, P_t)_{t \in T}$ satisfying R1, R2, and a third revision rule: the objective overlap of the functions P_t includes a fine (‘objective’) probability measure on \mathcal{R} .*

9 Concluding remarks

I have presented a unified framework and theorem for preferences under uncertainty and changing awareness. Preferences are governed by expected utility with three rules for revising utilities and probabilities as awareness changes. The theorem has many special cases, including (i) fixed awareness, where we recover Savage’s Theorem, (ii) fixed outcome awareness, where utilities are stable, (iii) fixed state awareness, where probabilities are stable, (iv) exhaustive state awareness, where some key definitions simplify, and (v) fine state awareness, where Axiom 6 and the third revision rule simplify. Just as Savage’s axioms have been weakened over time, giving rise to ‘non-EU’ theories, it would be interesting to relax the current axioms and explore ‘less rational’ representations and revision rules.

In our analysis, the agent appears as either stable or unstable in his preferences and beliefs, depending on whether a subjective or objective level of description is chosen. Our analysis suggests that the instability and context-dependence which agents display from an objective perspective are driven by a changing subjective perception of the objective world.

A When do tighter EU rationalizations exist?

I give partial answers to when *unified* or even *classical* EU rationalizations exist.

A.1 When do unified EU rationalizations exist?

Call an EU rationalization $(U_t, P_t)_{t \in T}$ (or just a family $(U_t)_{t \in T}$ of utility functions) *equivalent* to another if the latter arises from the former via increasing affine transformations of all U_t . Equivalent rationalizations are interchangeable in that they imply the same preferences. Two conditions on stability in concepts/awareness will do the intended job:

PSO (Partially Stable Outcome Concept): The contexts share at least two outcomes, i.e., $|\cap_{t \in T} X_t| \geq 2$; or more generally, one may partition T such that the contexts of any partition cell share at least two outcomes and no two contexts from distinct partition cells share any outcomes.

PSS (Partially Stable State Concept): The contexts share a non-null event, i.e., $\cap_{t \in T} S_t$ is a non-null event in each context in T ; or more generally, one may partition T such that the contexts of any partition cell share a non-null event and no two contexts from distinct partition cells share a non-null event.

Proposition 6 *Given any EU rationalization $(U_t, P_t)_{t \in T}$,*

- (i) *under PSO, $(U_t)_{t \in T}$ obeys R1 if and only if some equivalent family is stable,*
- (ii) *under PSS, $(P_t)_{t \in T}$ obeys R2 if and only if it is unified,*
- (iii) *so, under PSO and PSS, $(U_t, P_t)_{t \in T}$ obeys R1 and R2 if and only if some equivalent EU rationalization is unified.*

So Theorem 1 implies axiomatizations of some tighter rationalizations:

Corollary 3 *Under PSO, the variable Savage structure obeys Axioms 1–6 if and only if it has an EU rationalization which obeys R2 and R3 and is stable in utilities (i.e., obeys R1+).*

Corollary 4 *Under PSS, the variable Savage structure obeys Axioms 1–6 if and only if it has an EU rationalization which obeys R1 and R3 and is unified in probabilities (so obeys R2).*

Corollary 5 *Under PSO and PSS, the variable Savage structure obeys Axioms 1–6 if and only if it has a unified EU rationalization obeying R3 (and by unification R1+ and R2).*

Proof of Proposition 6. (i): Assume PSO. and let $(U_t)_{t \in T}$ obey R1. I define a function U on $\cup_{t \in T} X_t$ of which each U_t is an increasing affine transformation on X_t . It suffices to define U on $\cup_{t \in T'} X_t$ for any fixed cell T' of the partition in PSO. For scaling purposes, fix a $t^* \in T'$. For each $t \in T'$ let U'_t be the (by R1 existing) increasing affine transformation of U_t which coincides with U_{t^*} on

$X_t \cap X_{t^*}$. Any two functions U'_{t_1}, U'_{t_2} ($t_1, t_2 \in T'$) coincide on their domain overlap $X_{t_1} \cap X_{t_2}$, because they are increasing affine transformations of one another and coincide at two or more points (as $|X_{t_1} \cap X_{t_2}| \geq |\cap_{t \in T'} X_t| \geq 2$). So we may define U on $\cup_{t \in T'} X_t$ so as to match each U'_t ($t \in T'$) on X_t .

(ii): Assume PSS. Let $(P_t)_{t \in T}$ obey R2. I define a unifying function P on $\cup_{t \in T} 2^{S_t}$. It suffices to define P on $\cup_{t \in T'} 2^{S_t}$ for any fixed cell T' of the partition in PSS.⁵¹ Consider the non-null overlap $A := \cap_{t \in T'} S_t$. For scaling purposes fix $t^* \in T'$. For all $t \in T'$, let P'_t be the (by R2 existing) multiple of P_t which matches P_{t^*} on the domain overlap $2^{S_t \cap S_{t^*}}$, hence on A (P'_t is not a probability measure as $P'_t(S_t) \neq 1$, unless $P'_t = P_t$). Any two of the P'_t s ($t \in T'$) coincide on their domain overlap, since on the overlap they are multiples of each other and agree at A . So we may define P on $\cup_{t \in T'} 2^{S_t}$ as matching all P'_t ($t \in T'$) on 2^{S_t} . ■

A.2 When do classical EU rationalizations exist?

Classical probabilities. Recall that the agent's system probability measures $(P_t)_{t \in T}$ is 'classical' just in case it obeys the condition R2++ of *objective* stability: any *objective* event receives the same probability whenever it is representable. In Theorem 1 we can strengthen R2 to R2++ if we add a very strong axiom and meanwhile restrict attention to exhaustive states:

Axiom DIP (description-invariant preferences): For all contexts $t, t' \in T$, we have $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$ whenever acts $f, g \in F_t$ are objectively equivalent to acts $f', g' \in F_{t'}$, respectively.

Under this axiom, the preference between getting \$50 for sure and getting \$100 only in the objective event $A = \{a, b, c\}$ does not depend on whether A is subjectively represented as $\{\{a\}, \{b\}, \{c\}\}$, $\{\{a, b\}, \{c\}\}$, $\{\{a\}, \{b, c\}\}$, or $\{\{a, b, c\}\}$.

Corollary 6 *Under exhaustive states, the variable Savage structure satisfies Axioms 1–6 and DIP if and only if it has an EU rationalization satisfying R1, R2++ and R3.*⁵²

Proof. Assume exhaustive states. First, assume the stated representation. Ax. 1–6 hold by Theorem 1. Regarding Ax. DIP, consider $t, t' \in T$; w.l.o.g. let U_t and $U_{t'}$ coincide on $X_t \cap X_{t'}$. It suffices to consider $f \in F_t$ and $f' \in F_{t'}$ with $f^* = f'^*$ and show that $\mathbb{E}_{P_t}(U_t \circ f) = \mathbb{E}_{P_{t'}}(U_{t'} \circ f')$, or equivalently, that $\mathbb{E}_{P_t^*}(U_t \circ f^*) = \mathbb{E}_{P_{t'}^*}(U_{t'} \circ f'^*)$. This holds because (i) $U_t \circ f^* = U_{t'} \circ f'^*$, (ii) $f^* = f'^*$

⁵¹Conflicts across partition cells cannot arise since events shared by contexts from different partition cells will always be assigned the same value of zero

⁵²Theorem 2 has an analogous corollary.

is measurable w.r.t. the algebra of objective events representable in both contexts, and (iii) P_t^* and $P_{t'}^*$ agree on that algebra.

Conversely, assume Ax. 1–6 and DIP. By Theorem 1, there is a rationalization $(U_t, P_t)_{t \in T}$ obeying R1, R2 and R3. Let $O : \mathcal{A} \rightarrow [0, 1]$ be the (fine) objective agreement of the P_t s. To show R2++, let $A \subseteq \mathbb{S}$ be representable in contexts t and r . For a contradiction, let $P_t^*(A) \neq P_r^*(A)$, say $P_t^*(A) < P_r^*(A)$. As O is fine, its range is dense in $[0, 1]$. So we can pick a $B \in \mathcal{A}$ such that $P_t^*(A) < O(B) < P_r^*(A)$. As P_t and P_r agree on value $O(B)$ at B , P_t^* extends to some $P_{t'}^*$ where $S_{t'} = S_t \vee \{B, \overline{B}\}$ and $P_{t'}^*(B) = O(B)$; and P_r^* extends to some $P_{r'}^*$ where $S_{r'} = S_r \vee \{B, \overline{B}\}$ and $P_{r'}^*(B) = O(B)$. The inequalities $P_t^*(A) < O(B)$ and $O(B) < P_r^*(A)$ now reduce to $P_{t'}^*(A) < P_{r'}^*(B)$ and $P_{r'}^*(B) < P_{r'}^*(A)$. This contradicts belief-stability on robust algebras (Prop. 5), since A and B belong to $2^{\mathbb{S}}$, a robust algebra by Ax. DIP. ■

If we take limited awareness seriously, Axiom DIP (with its implication R2++) is questionable, as it operates *as if* the agent had awareness of objective states and could thus recognize the objective equivalence between acts or events.

Classical utilities. It would go far beyond this paper to try to provide axiomatic foundations for classicality in utilities. But here is a hunch of how one might proceed. Assume an EU rationalization is already stable in utilities (Corollary 3 tells us how to get up to this point). Let U be the stable utility function on $\cup_{t \in T} X_t$. Bolker’s (1966) powerful representation theorem tells us when a function f on a system Σ of subsets of an underlying set Ω has an ‘expectational representation’, so that any $f(A)$ is the *expectation conditional on A* of a some fixed function V on Ω w.r.t. some fixed probability measure μ on Σ . (Bolker did not have EU theory in mind, but his result became the basis of Jeffrey’s 1983 far too unknown alternative EU theory.) If we now think of Ω as \mathbb{X} , of Σ as $\cup_{t \in T} X_t$, and of f as U , then Bolker’s expectational representation of f renders U ($= f$) ‘classical’ in our sense. Of course, Bolker’s expectational representation applies under conditions that do not yet follow from our axioms; he for instance takes Σ to be an algebra. One might thus try to strengthen our axioms such that they imply (an EU rationalization with) a stable utility function U which obeys Bolker-type conditions that guarantee an expectational (hence, ‘classical’) representation.

B Proof of the stability propositions and the claim in footnote 26

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 and some technical lemmas which will be stated in due course and proved *en bloc* at the very end to avoid distraction (App. F). Proofs use the following notation:

- Recall the notation ‘ x_t ’ and ‘ a_t ’ (Def. 4), ‘ A^* ’ and ‘ f^* ’ (Sect. 2.1), ‘ A_t ’ (Def. 20_{exh} or 20), and P^* (Def. 28).
- For any set E of events, define the set of objective events $E^* := \{A^* : A \in E\}$.
- For any set F of acts, define the set of functions $F^* := \{f^* : f \in F\}$.
- For any $t \in T$ and $f \in F_t^*$, let $f_t \in F_t$ be the act given by $(f_t)^* = f$. (‘ f_t ’ was also used for the representation of an objective act $f \in \mathbb{F}$; see Def. 5.)

Proof of Prop. 1. Just take $A = S_t = S_{t'}$, $f = f'$ and $g = g'$ in Ax. 2. ■

Proof of Prop. 2. Ax. 6 implies existence of a robust algebra \mathcal{R} on \mathbb{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. 3. Assume Ax. 2, 4, 5 and 6, and let $t, t' \in T$ and $A, B \subseteq S_t \cap S_{t'}$. I suppose $A \succsim_t B$ and show that $A \succsim_{t'} B$ (the converse is analogous). Using Ax. 5, pick outcomes $x \succ_t y$ in X_t and $x' \succ_{t'} y'$ in $X_{t'}$. Using independence between outcome and state awareness, pick a $r \in T$ with $X_r = X_t$ and $S_r = S_{t'}$. As $x \succ_t y$ we have $x \succ_r y$ by Prop. 2. As $A \succsim_t B$ and $x \succ_t y$, by Ax. 4 $x_{AY_{S_t \setminus A}} \succsim_t x_{BY_{S_t \setminus B}}$. So $x_{AY_{S_r \setminus A}} \succsim_r x_{BY_{S_r \setminus B}}$ by Ax. 2 applied to the event $A \cup B \subseteq S_t \cap S_r$. Hence $x'_{AY'_{S_{t'} \setminus A}} \succsim_{t'} x'_{BY'_{S_{t'} \setminus B}}$ by Ax. 4. So $A \succsim_{t'} B$. ■

Proof of Prop. 5. Assume Ax. 2, 4 and 5. Consider a robust algebra \mathcal{R} on \mathbb{S} , $t, t' \in T$, $A, B \subseteq S_t$, and $A', B' \subseteq S_{t'}$, such that A and A' represent an identical objective event $\tilde{A} \in \mathcal{R}$, and B and B' also represent an identical $\tilde{B} \in \mathcal{R}$. I show that $A \succsim_t B \Rightarrow A' \succsim_{t'} B'$ (the converse direction ‘ \Leftarrow ’ holds analogously). Let $A \succsim_t B$. Using Ax. 5, pick outcomes $x \succ_t y$ in X_t and $x' \succ_{t'} y'$ in $X_{t'}$. Using independence between outcome and state awareness, pick an $r \in T$ with $X_r = X_t$ and $S_r = S_{t'}$. As $x \succ_t y$ we have $x \succ_r y$, by Prop. 2 (more exactly, a version of Prop. 2 based not on Ax. 6, but only on the existence of a robust algebra). Also, as $A \succsim_t B$ and $x \succ_t y$, by Ax. 4 $x_{AY_{S_t \setminus A}} \succsim_t x_{BY_{S_t \setminus B}}$. So $x_{A'Y_{S_r \setminus A'}} \succsim_r x_{B'Y_{S_r \setminus B'}}$, because \mathcal{R} is robust, $x_{AY_{S_t \setminus A}}$ and $x_{A'Y_{S_r \setminus A'}}$ are corresponding \mathcal{R} -determined acts (as both stem from the \mathcal{R} -measurable function $x_{\tilde{A}Y_{\mathbb{S} \setminus \tilde{A}}}$), and $x_{BY_{S_t \setminus B}}$ and $x_{B'Y_{S_r \setminus B'}}$ are also corresponding \mathcal{R} -determined acts (as both stem from the \mathcal{R} -measurable function $x_{\tilde{B}Y_{\mathbb{S} \setminus \tilde{B}}}$). As $x_{A'Y_{S_r \setminus A'}} \succsim_r x_{B'Y_{S_r \setminus B'}}$ and $x \succ_r y$, by Ax. 4 $A' \succsim_r B'$. So $A' \succsim_{t'} B'$, by Prop. 3 (more precisely, a version of Prop. 3, like before). ■

Proof of Prop. 4. Assume Ax. 2. Let $t, t' \in T$ and $A \subseteq S_t \cap S_{t'}$. We let A be non-null in t and prove non-nullness in t' . By assumption, there are $f, g \in F_t$ such that $f_{S_t \setminus A} = g_{S_t \setminus A}$ and $f \not\sim_t g$. Pick $f', g' \in F_{t'}$ such that $f_A = f'_A$, $g_A = g'_A$, and $f'_{S_{t'} \setminus A} = g'_{S_{t'} \setminus A}$. As $f \not\sim_t g$, by Ax. 2 $f' \not\sim_{t'} g'$. So A is non-null in t' . ■

Proof of the claim in footnote 26. Assume the condition in footnote 26 and Ax. 1, 2 and 6. Ax. 4* clearly implies Ax. 4. Now assume Ax. 4, and let $t, t', S_t, S_{t'}, S, A, B, x, y, x', y'$ satisfy the premises of Ax. 4*. We must show that $x_{AY_{S \setminus A}} \succsim_t x_{BY_{S \setminus B}} \Leftrightarrow x'_{AY'_{S \setminus A}} \succsim_{t'} x'_{BY'_{S \setminus B}}$. By assumption we may pick $\hat{t}, \hat{t}' \in T$ and $z, w \in X_{\hat{t}} \cap X_{\hat{t}'}$ such that $X_t \subseteq X_{\hat{t}}, X_{t'} \subseteq X_{\hat{t}'}$, and $z \not\sim_{\hat{t}} w$. So by Ax. 1 $z \succ_{\hat{t}} w$ or $w \succ_{\hat{t}} z$; w.l.o.g. let $z \succ_{\hat{t}} w$. Hence also $z \succ_{\hat{t}'} w$ by outcome-preference stability (see Prop. 2 which uses Ax. 6). Now $x_{AY_{S \setminus A}} \succsim_t x_{BY_{S \setminus B}}$ is equivalent to $x_{AY_{S \setminus A}} \succsim_{\hat{t}} x_{BY_{S \setminus B}}$ by preference stability (see Prop. 1 which uses Ax. 2), hence is equivalent to $z_{AW_{S \setminus A}} \succsim_{\hat{t}} z_{BW_{S \setminus B}}$ by Ax. 4. Similarly, $x'_{AY'_{S \setminus A}} \succsim_{t'} x'_{BY'_{S \setminus B}}$ is equivalent to $x'_{AY'_{S \setminus A}} \succsim_{\hat{t}'} x'_{BY'_{S \setminus B}}$, hence to $z_{AW_{S \setminus A}} \succsim_{\hat{t}'} z_{BW_{S \setminus B}}$. So it remains to show $z_{AW_{S \setminus A}} \succsim_{\hat{t}} z_{BW_{S \setminus B}} \Leftrightarrow z_{AW_{S \setminus A}} \succsim_{\hat{t}'} z_{BW_{S \setminus B}}$. This holds by preference stability. ■

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_t ($t \in T$) 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 $I_1, I_2, \dots \subseteq \mathbb{S}$, we can successively refine the state space S_t to $S_{t_1} = S_t \vee \{I_1, \bar{I}_1\}$ (for a context t_1), then to $S_{t_2} = S_{t_1} \vee \{I_2, \bar{I}_2\}$ (for a context t_2), and so on. In each step another I_i becomes representable, and the new relation \succsim_{t_i} remains faithful to the earlier ones if it has the same outcome space. These refinements do not lead far enough: if S_t was finite, then all S_{t_i} are finite, hence still too small ‘for Savage’. We will thus go further: we will faithfully extrapolate each \succsim_t to a relation whose state space incorporates infinitely many and indeed *all* incorporable $I \subseteq \mathbb{S}$. This high state sophistication is purely hypothetical: it might never be reached by the agent in any context in T . 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_t^+ , satisfies Savage’s axioms, (ii) deducing a EU representation of each \succsim_t^+ via Savage’s Theorem in Kopylov’s (2007) version, and (iii) deducing suitable EU representations (U_t, P_t) of the original relations \succsim_t ($t \in T$) jointly 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.

Reformulating R3. Hereafter whenever we refer to R3, we will take this revision rule to be stated in the following (equivalent) way, which draws on the notion of

an ‘agreed function’:

Definition 34 Functions P_t on 2^{S_t} ($t \in T$) **agree on** a given function ρ on a set of objective events $\mathcal{R} \subseteq 2^{\mathbb{S}}$ – and ρ is **agreed among** the P_t s – if the P_t s agree on ρ ’s value at each $A \in \mathcal{R}$ (cf. Def. 29), or equivalently, if ρ is a possibly domain-restricted version of the objective agreement of the P_t s (cf. Def. 30).

R3 (equivalent restatement): The P_t s agree (at least) on a fine probability measure on an algebra on \mathbb{S} .

C.1 Definition of extrapolated preferences

As mentioned, we extrapolate each relation \succsim_t 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 35 An objective event $A \subseteq \mathbb{S}$ is **weakly incorporable** if there is a finite partition \mathcal{P} of \mathbb{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 $t \in T$ there is a $t' \in T$ (possibly equal to t) with $S_{t'} = S_t \vee \mathcal{P}$ and with $\succsim_{t'}$ faithful to \succsim_t . Let $\mathcal{I} := \{A \subseteq \mathbb{S} : A \text{ is weakly incorporable}\}$.

Remark 27 Incorporability implies weak incorporability: here $\mathcal{P} = \{A, \bar{A}\}$.

Remark 28 The set \mathcal{I} of weakly incorporable objective events is an algebra on \mathbb{S} : (i) $\mathbb{S} \in \mathcal{I}$; (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} and \mathcal{P}' respectively) 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_t to a partition S' of \mathbb{S} (a hypothetical subjective state space) in which all $I \in \mathcal{I}$ are representable, and to extrapolate the relation \succsim_t to one on $X_t^{S'}$. It will in fact be easier to work not with a (hypothetical) subjective state space S' , but with the *objective* state space \mathbb{S} . So I will extrapolate \succsim_t to a relation on the set $X_t^{\mathbb{S}}$ of ‘semi-objective acts’, which map objective states to subjective outcomes.

Definition 36 A partition of \mathbb{S} **harmlessly refines** another one S if it is the join of S and some finite partition of \mathbb{S} into weakly incorporable objective events.

Definition 37 For a contexts $t \in T$, the **extrapolated relation** \succsim_t^+ on $X_t^{\mathbb{S}}$ is given as follows: $f \succsim_t^+ g$ if and only if $f_\tau \succsim_\tau g_\tau$ for some context $\tau \in T$ such that (i) $f, g \in F_\tau^*$ (so f_τ and g_τ are defined) and (ii) S_τ harmlessly refines S_t .⁵³

⁵³Clause (ii) ensures that \succsim_t^+ is intimately linked to (i.e., ‘extrapolates’) \succsim_t .

C.2 Sufficiency of the axioms

Using extrapolated preferences, I now gradually prove sufficiency.

Definition 38 *The join $\mathcal{R} \vee \mathcal{R}'$ of algebras \mathcal{R} and \mathcal{R}' on \mathbb{S} is the smallest algebra $\mathcal{A} \supseteq \mathcal{R} \cup \mathcal{R}'$ on \mathbb{S} , i.e., the closure of $\mathcal{R} \cup \mathcal{R}'$ under complement and finite union.*

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

Definition 39 *The **extrapolated algebra** for context $t \in T$ is the set \mathcal{E}_t of objective events that are representable after a harmless state refinement: $\mathcal{E}_t := \{A^* : A \subseteq S \text{ for some harmless refinement } S \text{ of } S_t\}$.*

Lemma 3 *For all contexts $t \in T$, \mathcal{E}_t is an algebra on \mathbb{S} , characterizable as*

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

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

An ordinary Savage structure (X, S, \succsim) is identified with the generalized one $(X, (S, 2^S), \succsim)$. In a generalized Savage structure $(X, (S, \mathcal{A}), \succsim)$ 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 $A \in \mathcal{A}$, if $f_A = f'_A$, $g_A = g'_A$, $f_{S \setminus A} = g_{S \setminus A}$ and $f'_{S \setminus A} = g'_{S \setminus A}$, then $f \succsim g \Leftrightarrow f' \succsim g'$.
- P3:** For all $x, y \in X$ and non-null $A \in \mathcal{A}$, $x \succsim_A y \Leftrightarrow x \succsim y$.⁵⁴
- P4:** For all $A, B \in \mathcal{A}$ and all $x \succ y$ and $x' \succ y'$ in X , $x_A y_{S \setminus A} \succsim x_B y'_{S \setminus B} \Leftrightarrow x'_A y'_{S \setminus A} \succsim_t x'_B y'_{S \setminus B}$.

⁵⁴Elements of X 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 $A \in \mathcal{A}$ – written $f \succsim_A g$ – if $f' \succsim g'$ for some acts f' and g' such that $f_A = f'_A$, $g_A = g'_A$ and $f'_{S \setminus A} = g'_{S \setminus A}$.

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 some $A_1, \dots, A_n \in \mathcal{A}$ such that $f_{S \setminus A_i} x_{A_i} \succ g$ and $f \succ g_{S \setminus A_i} x_{A_i}$ for all A_i .

Lemma 4 (*Savage's Theorem for arbitrary event algebras; see Kopylov 2007*) *A generalized Savage structure $(X, (S, \mathcal{A}), \succsim)$ 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 $P : \mathcal{A} \rightarrow [0, 1]$ 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.⁵⁵*

Lemma 5 *If Ax. 1–6 hold, then for each context $t \in T$ Ax. P1–P6 hold for the generalized Savage structure $(X_t, (\mathbb{S}, \mathcal{A}), \succsim)$ in which (i) \mathcal{A} is \mathcal{E}_t or more generally any algebra such that $\mathcal{R} \subseteq \mathcal{A} \subseteq \mathcal{E}_t$ for some algebra \mathcal{R} as in Ax. 6, and (ii) \succsim is \succsim_t^+ restricted to the set of acts $F = \{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{A}\text{-measurable}\}$.*

Lem. 5's proof rests on some technical lemmas (shown in App. F):

Lemma 6 *Under Ax. 2, a relation \succsim_τ is faithful to another \succsim_t if $X_\tau \supseteq X_t$ and S_τ harmlessly refines S_t .*

Lemma 7 *Under Ax. 2, whenever $f \succsim_t^+ g$ (where $t \in T$ and $f, g \in X_t^{\mathbb{S}}$), then*

- (a) $f_\tau \succsim_\tau g_\tau$ for all (not just some) $\tau \in T$ satisfying (i)–(ii) in Def. 37,
- (b) $f_\tau \succsim_\tau g_\tau$ for some $\tau \in T$ such that (i)–(ii) in Def. 37 hold and \succsim_τ is faithful to \succsim_t (in particular, $X_\tau \supseteq X_t$).

Lemma 8 *For any context $t \in T$ and finite set $\mathcal{B} \subseteq \mathcal{E}_t$, there is a context $\tau \in T$ such that (i) all $B \in \mathcal{B}$ are representable (i.e., $\mathcal{B} \subseteq (2^{S_\tau})^*$), (ii) S_τ harmlessly refines S_t , and (iii) \succsim_τ is faithful to \succsim_t .*

Lemma 9 *For all contexts $t \in T$ and finite sets \mathcal{G} of \mathcal{E}_t -measurable functions from \mathbb{S} to X_t , there is a context $\tau \in T$ such that (i) $\mathcal{G} \subseteq F_\tau^*$, (ii) S_τ harmlessly refines S_t , and (iii) \succsim_τ is faithful to \succsim_t .*

Lemma 10 *Assume Ax. 2 and 5 and let $t \in T$. For all acts f, g and events A of Lem. 5's generalized Savage structure, the conditional preference $f \succsim_A g$, i.e., $f \succsim_{t,A}^+ g$, holds if and only if $f_\tau \succsim_{\tau, A_\tau} g_\tau$ holds for some $\tau \in T$ such that $f, g \in F_\tau^*$, A is representable in context τ , and S_τ harmlessly refines S_t . The equivalence remains true when also requiring that \succsim_τ is faithful to \succsim_t .*

⁵⁵Kopylov proves this theorem for the case that \mathcal{A} is a *mosaic*, a more general structure than an algebra. My statement uses the condition that P is *fine*, which in the algebra case is equivalent to his condition that P is *finely ranged*. In Savage's special case $\mathcal{A} = 2^S$, a probability measure P on \mathcal{A} is fine if and only if it is atomless, and if and only if for all $A \in \mathcal{A}$ and $0 < \lambda < 1$ there is a $B \subseteq A$ in \mathcal{A} such that $P(B) = \lambda P(A)$. In general, fineness is weaker than atomlessness.

Lemma 11 *Assume Ax. 2 and let $t \in T$. An event A in Lem. 5's generalized Savage structure is non-null if and only if A_τ is a non-null event in some context $\tau \in T$ such that A is representable (i.e., A_τ is defined) and S_τ harmlessly refines S_t . The equivalence remains true when also requiring that \succsim_τ is faithful to \succsim_t .*

Proof of Lem. 5. Assume Ax. 1–6. Let $t, \mathcal{R}, \mathcal{A}$ be as specified. I show P1–P6 for the extrapolated relation \succsim_t^+ restricted to $F := \{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{A}\text{-measurable}\}$.

Claim 1: P1 holds. To show completeness, let $f, g \in F$. Using Lem. 9, pick a $\tau \in T$ such that $f, g \in F_\tau^*$ and S_τ harmlessly refines S_t . By Ax. 1, $f_\tau \succsim_\tau g_\tau$ or $g_\tau \succsim_\tau f_\tau$. In the first case $f \succsim_t^+ g$, in the second $g \succsim_t^+ f$. To show transitivity, let $f, g, h \in F$ such that $f \succsim_t^+ g$ and $g \succsim_t^+ h$. Using Lem. 9, pick a $\tau \in T$ such that $f, g, h \in F_\tau^*$ and S_τ harmlessly refines S_t . So, as $f \succsim_t^+ g$ and $g \succsim_t^+ h$, we have $f_\tau \succsim_\tau g_\tau$ and $g_\tau \succsim_\tau h_\tau$ by Lem. 7. Hence, $f_\tau \succsim_\tau h_\tau$ by Ax. 1, and so $f \succsim_t^+ h$.

Claim 2: P2 holds. Consider $f, g, f', g' \in F$ and $A \in \mathcal{A}$ such that $f_A = f'_A$, $g_A = g'_A$, $f_{S \setminus A} = g_{S \setminus A}$ and $f'_{S \setminus A} = g'_{S \setminus A}$. Pick an $h \in F$ taking one value on A and another on \bar{A} (h exists as $|X_t| \geq 2$ by Ax. 5). Using Lem. 9, pick a $\tau \in T$ such that $f, g, f', g', h \in F_\tau^*$ and S_τ harmlessly refines S_t . As $f, g, g', g' \in F_\tau^*$, the acts $f_\tau, g_\tau, f'_\tau, g'_\tau \in F_\tau$ are defined; and as $h \in F_\tau^*$, the event A is representable in context τ , so that A_τ is defined (the sole purpose of introducing h was indeed to ensure representability of A). Note that $(f_\tau)_{A_\tau} = (f'_\tau)_{A_\tau}$, $(g_\tau)_{A_\tau} = (g'_\tau)_{A_\tau}$, $(f_\tau)_{S \setminus A_\tau} = (g_\tau)_{S \setminus A_\tau}$ and $(f'_\tau)_{S \setminus A_\tau} = (g'_\tau)_{S \setminus A_\tau}$. So, by Ax. 2 (or just 2*), $f_\tau \succsim_\tau g_\tau \Leftrightarrow f'_\tau \succsim_\tau g'_\tau$. This equivalence reduces to $f \succsim_t^+ g \Leftrightarrow f' \succsim_t^+ g'$ by Lem. 7.

Claim 3: P3 holds. Let $x, y \in X_t$. Let $A \in \mathcal{A}$ be non-null. I show $x \succsim_{t,A}^+ y \Leftrightarrow x \succsim_t^+ y$. By Lem. 11, A_τ is non-null for a $\tau \in T$ such that A is representable, $S_\tau = S_t \vee \mathcal{P}$ with a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbb{S} , and \succsim_τ is faithful to \succsim_t . First, if $x \succsim_t^+ y$, then $x \succsim_\tau y$ by Lem. 7, so $x \succsim_{\tau, A_\tau} y$ by Ax. 3 and A_τ 's non-nullness, hence $x \succsim_{t,A}^+ y$ by Lem. 10. Now let $x \succsim_{t,A}^+ y$. By Lem. 10, $x \succsim_{\sigma, A_\tau} y$ for a $\sigma \in T$ such that A is representable, $S_\sigma = S_t \vee \mathcal{Q}$ with a finite partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbb{S} , and \succsim_σ is faithful to \succsim_t . Using Lem. 8, pick $\delta \in T$ such that $\mathcal{P} \cup \mathcal{Q} \subseteq (2^{S_\delta})^*$, $S_\delta = S_t \vee \mathcal{P}'$ with a finite partition $\mathcal{P}' \subseteq \mathcal{I}$ of \mathbb{S} , and \succsim_δ is faithful to \succsim_t . W.l.o.g. X_τ and X_σ equal X_t (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_τ and \succsim_σ , each time by Lem. 6, using that $X_\delta \supseteq X_\tau = X_\sigma (= X_t)$ and that $S_\delta = S_\tau \vee \mathcal{P}' = S_\sigma \vee \mathcal{P}'$ (since each set equals $S_t \vee \mathcal{P}'$ as \mathcal{P}' refines \mathcal{P} and \mathcal{Q}). As $A_\tau (\subseteq S_\tau)$ is non-null and \succsim_δ is faithful to \succsim_τ , $A_\delta (\subseteq S_\delta)$ is non-null. As $x \succsim_{\sigma, A_\sigma} y$ and \succsim_δ is faithful to \succsim_σ , $x \succsim_{\delta, A_\delta} y$. So $x \succsim_\delta y$ by Ax. 3. Thus $x \succsim_t^+ y$.

Claim 4: P4 holds. Let $A, B \in \mathcal{A}$ and $x, y, x', y' \in X_t$ such that $x \succ_t^+ y$ and $x' \succ_t^+ y'$. I show $x_{A Y_{\mathbb{S} \setminus A}} \succsim_t^+ x_{B Y_{\mathbb{S} \setminus B}} \Leftrightarrow x'_{A Y_{\mathbb{S} \setminus A}} \succsim_t^+ x'_{B Y_{\mathbb{S} \setminus B}}$. Via Lem. 9, pick a $\tau \in T$ such that $x_{A Y_{\mathbb{S} \setminus A}}, x_{B Y_{\mathbb{S} \setminus B}}, x'_{A Y_{\mathbb{S} \setminus A}}, x'_{B Y_{\mathbb{S} \setminus B}} \in F_\tau^*$ and S_τ harmlessly refines S_t . By Lem. 7, $x \succ_\tau y$ and $x' \succ_\tau y'$. So the claimed equivalence reduces to $(x_{A Y_{\mathbb{S} \setminus A}})_\tau \succsim_\tau (x_{B Y_{\mathbb{S} \setminus B}})_\tau \Leftrightarrow (x'_{A Y_{\mathbb{S} \setminus A}})_\tau \succsim_\tau (x'_{B Y_{\mathbb{S} \setminus B}})_\tau$, i.e., $x_{A_\tau Y_{S_\tau \setminus A_\tau}} \succsim_\tau$

$x_{B_\tau} y_{S_\tau \setminus B_\tau} \Leftrightarrow x'_{A_\tau} y'_{S_\tau \setminus A_\tau} \succsim_\tau x'_{B_\tau} y'_{S \setminus B_\tau}$. This holds by Ax. 4.

Claim 5: P5 holds. Using Ax. 5, pick $f \succ_t g$ in F_t . Clearly, $f^* \succ_t^+ g^*$.

Claim 6: P6 holds. Let $f \succ_t^+ g$ in F and $x \in X_t$. As $f \succsim_t^+ g$, we have $f_\tau \succsim_\tau g_\tau$ for a $\tau \in T$ such that $f, g \in F_\tau^*$, $S_\tau = S_t \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbb{S} , and \succsim_τ is faithful to \succsim_t . Note $x \in X_\tau$; and $g_\tau \not\prec_\tau f_\tau$ as $g \not\prec_t^+ f$. So $f_\tau \succ_\tau g_\tau$. As $\mathcal{R} (\subseteq \mathcal{A})$ is as in Ax. 6, one can partition \mathbb{S} into $R_1, \dots, R_n \in \mathcal{R} \subseteq \mathcal{A}$ such that, for a $t' \in T$ with $S_{t'} = S_\tau \vee \{R_1, \dots, R_n\}$ and $X_{t'} \supseteq X_\tau$, $(f_{t'})_{S_{t'} \setminus (R_i)_{t'}} \succ_{t'} g_{t'}$ and $f_{t'} \succ_{t'} (g_{t'})_{S_{t'} \setminus (R_i)_{t'}} x_{(R_i)_{t'}}$ for all i , i.e., $(f_{\mathbb{S} \setminus R_i} x_{R_i})_{t'} \succ_{t'} g_{t'}$ and $f_{t'} \succ_{t'} (g_{\mathbb{S} \setminus R_i} x_{R_i})_{t'}$ for all i . So (as $S_{t'}$ harmlessly refines S_t , being the join of S_t and $\mathcal{P} \vee \{R_1, \dots, R_n\} \subseteq \mathcal{I}$), $f_{\mathbb{S} \setminus R_i} x_{R_i} \succ_t^+ g$ and $f \succ_t^+ g_{\mathbb{S} \setminus R_i} x_{R_i}$ for all i .⁵⁶ ■

Given Ax. 1–6, for each $t \in T$ we now use Lem. 4 and 5 to pick a utility function U_t on X_t and a fine probability measure P_t^+ on \mathcal{E}_t which represent the extrapolated relation \succsim_t^+ on $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{E}_t\text{-measurable}\}$:

$$f \succsim_t^+ g \Leftrightarrow \mathbb{E}_{P_t^+}(U_t \circ f) \geq \mathbb{E}_{P_t^+}(U_t \circ g) \text{ for all } \mathcal{E}_t\text{-measurable } f, g \in X_t^{\mathbb{S}}.$$

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

$$P_t(A) := P_t^+(A^*) \text{ for all } A \subseteq S_t.$$

The next four lemmas complete the sufficiency proof by establishing that the functions P_t and U_t ($t \in T$) have all properties required in Thm. 1.

Lemma 12 *Under Ax. 1–6, the above-defined $(U_t, P_t)_{t \in T}$ is an EU rationalization.*

Lemma 13 *Under Ax. 1–6, the above-defined functions U_t satisfy R1.*

Lemma 14 *Under Ax. 1–6, the above-defined functions P_t satisfy R2.*

Lemma 15 *Under Ax. 1–6, for each algebra \mathcal{R} as in Ax. 6,*

- (a) *all above-defined measures P_t^+ have identical restriction $\rho := P_t^+|_{\mathcal{R}}$,*
- (b) *the above-defined measures P_t satisfy R3 in virtue of ρ as the fine agreed probability measure.*

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

Proof of Lem. 12. Assume Ax. 1–6. Let $t \in T$ and $f, g \in F_t$. Let U_t, P_t and P_t^+ be as above. I show $f \succsim_t g \Leftrightarrow \mathbb{E}_{P_t}(U_t \circ f) \geq \mathbb{E}_{P_t}(U_t \circ g)$. The left side

⁵⁶To make the last step, one needs to first decompose each strict preference ($\succ_{t'}$) into a weak preference ($\succsim_{t'}$) without weak dispreference ($\not\prec_{t'}$), then infer corresponding extended weak preferences (\succsim_t^+) without weak dispreference ($\not\prec_t^+$) using Lem. 7, which implies extended strict preferences (\succ_t^+).

reduces to $f^* \succsim_t^+ g^*$ by Lem. 7, and the right side to $\mathbb{E}_{P_t^+}(U_t \circ f^*) \geq \mathbb{E}_{P_t^+}(U_t \circ g^*)$ because, letting $\sigma : \mathbb{S} \rightarrow S_t$ map any $a \in \mathbb{S}$ to its subjectivization $\sigma(a) = a_t$, we have $f^* = f \circ \sigma$, $g^* = g \circ \sigma$, and P_t is P_t^+ 's image under σ . To complete the proof, note $f^* \succsim_t^+ g^* \Leftrightarrow \mathbb{E}_{P_t^+}(U_t \circ f^*) \geq \mathbb{E}_{P_t^+}(U_t \circ g^*)$ by definition of U_t and P_t^+ . ■

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

Lemma 16 (*Niiniluoto 1972, Wakker 1981*) *Every fine and tight qualitative probability relation on an algebra \mathcal{A} on \mathbb{S} (not necessarily a σ -algebra) is uniquely representable by a probability measure on \mathcal{A} .*

Lemma 17 (*Wakker 1981, Kopylov 2007⁵⁷*) *A probability measure on an algebra \mathcal{A} on \mathbb{S} (not necessarily a σ -algebra) is fine if and only if the represented qualitative probability relation is fine and tight.*

I also need 5 technical lemmas (proved in App. F), the first 3 about the *extrapolated belief* relation over objective events induced by \succsim_t^+ and still labelled ‘ \succsim_t^+ ’.

Lemma 18 (*extrapolated comparative beliefs*) *Under Ax. 2, 4 and 5, for all $t \in T$ and $A, B \subseteq \mathbb{S}$, $A \succsim_t^+ B$ if and only if $A_\tau \succsim_\tau B_\tau$ for some $\tau \in T$ such that A and B are representable (i.e., A_τ and B_τ are defined) and S_τ harmlessly refines S_t . The equivalence remains true when also requiring \succsim_τ to be faithful to \succsim_t .*

Lemma 19 *Under Ax. 2, 4, 5 and 6, whenever $A \succsim_t^+ B$ (where $t \in T$ and $A, B \subseteq \mathbb{S}$), then $A_\tau \succsim_\tau B_\tau$ for each (not just some) context $\tau \in T$ in which A and B are representable (so that A_τ and B_τ are defined) and S_τ harmlessly refines S_t .*

Lemma 20 *Under Ax. 2, 4 and 5, the extrapolated relations \succsim_t^+ ($t \in T$) agree (as belief relations on $2^{\mathbb{S}}$) on each robust algebra \mathcal{R} of incorporable objective events.*

Lemma 21 *Under Ax. 1–6, the restriction of the above-defined measure P_t^+ to an algebra \mathcal{R} of type in Ax. 6 is (a) fine, and (b) the same for all $t \in T$.*

Lemma 22 *Given Ax. 1 and 2, for any contexts $t, \tau \in T$, if S_τ harmlessly refines S_t , then $\mathcal{E}_t = \mathcal{E}_\tau$, and if moreover \succsim_τ is faithful to \succsim_t then $\succsim_\tau^+ = \succsim_t^+$.*

Proof of Lem. 15. Assume Ax. 1–6. Let \mathcal{R} be as in Ax. 6, and P_t and P_t^+ ($t \in T$) as above. By Lem. 21, $\rho := P_t^+|_{\mathcal{R}}$ is fine and independent of $t \in T$. I show ρ is agreed. Let $A \in \mathcal{R}$ and $t \in T$. I must show existence of a $\tau \in T$

⁵⁷Lem. 17 is implicit in Wakker (1981) and a special case of Kopylov’s (2007) Thm. A.1.

such that $S_\tau = S_t \vee \{A, \bar{A}\}$, P_τ^* extends P_t^* , and $P_\tau^*(A) = \rho(A)$. As $A \in \mathcal{R}$, A is incorporable; so pick a $\tau \in T$ such that $S_\tau = S_t \vee \{A, \bar{A}\}$ and \succsim_τ is faithful to \succsim_t . By Lem. 22, $\mathcal{E}_t = \mathcal{E}_\tau$ and $\succsim_t^+ = \succsim_\tau^+$. So $P_t^+ = P_\tau^+$. Thus $P_\tau^* (= P_\tau^+|_{(2^{S_\tau})^*})$ extends $P_t^* (= P_t^+|_{(2^{S_t})^*})$. Finally, $P_\tau^*(A) = P_\tau^+(A) = \rho(A)$. ■

Proof of Lem. 14. Assume Ax. 1–6. Let P_t, P_t^+ ($t \in T$) be as above, $t, \tau \in T$, and $S := S_t \cap S_\tau$. If S is null in both contexts, P_t and P_τ are zero, so proportional, on 2^S . Now let S be non-null in one, hence by Prop. 4 both, contexts. Let \mathcal{R} be as in Ax. 6. Put $\mathcal{A} := \{A^* : A \subseteq S \vee \mathcal{P} \text{ for a finite partition } \mathcal{P} \subseteq \mathcal{R} \text{ of } \mathbb{S}\}$. Here $S \vee \mathcal{P}$ joins partitions of *distinct* sets S^* and \mathbb{S} ; Def. 21 still applies.

Claim 1: The measures P_t^+ and P_τ^+ are ordinally equivalent on \mathcal{A} . Note \mathcal{A} is an algebra on S^* , not \mathbb{S} .⁵⁸ Let $A, B \in \mathcal{A}$. I show $P_t^+(A) \geq P_t^+(B) \Leftrightarrow P_\tau^+(A) \geq P_\tau^+(B)$, or equivalently (as P_t^+ represents \succsim_t^+ and P_τ^+ represents \succsim_τ^+) $A \succsim_t^+ B \Leftrightarrow A \succsim_\tau^+ B$. As $A, B \in \mathcal{A}$, we may pick finite partitions $\mathcal{P}_A, \mathcal{P}_B \subseteq \mathcal{R}$ of \mathbb{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 \mathbb{S} . Using that all $X \in \mathcal{P}$ are incorporable (as $\mathcal{P} \subseteq \mathcal{R}$), pick $t', \tau' \in T$ such that $S_{t'} = S_t \vee \mathcal{P}$ and $S_{\tau'} = S_\tau \vee \mathcal{P}$. Now A and B are representable in context t' (as $S_{t'}$ refines $S_t \vee \mathcal{P}_A$); so $A \succsim_t^+ B \Leftrightarrow A_{t'} \succsim_{t'} B_{t'}$ by Lem. 18 and 19. Similarly, A and B are representable in context τ' ; so $A \succsim_\tau^+ B \Leftrightarrow A_{\tau'} \succsim_{\tau'} B_{\tau'}$. It remains to show $A_{t'} \succsim_{t'} B_{t'} \Leftrightarrow A_{\tau'} \succsim_{\tau'} B_{\tau'}$. This holds by comparative-belief stability (Prop. 3), since $A_{t'} = A_{\tau'}$ and $B_{t'} = B_{\tau'}$ as $S_{t'}$ and $S_{\tau'}$ agree within S^* ($\supseteq A, B$).

Claim 2: P_t and P_τ are proportional on 2^S . By Claim 1, the conditional probability measures $P_t^+(\cdot|S^*)$ and $P_\tau^+(\cdot|S^*)$ are ordinally equivalent on \mathcal{A} . Their restrictions $P_t^+(\cdot|S^*)|_{\mathcal{A}}$ and $P_\tau^+(\cdot|S^*)|_{\mathcal{A}}$ are probability measures on \mathcal{A} (an algebra on S^*), which are fine as P_t^+ and P_τ^+ are fine. So $P_t^+(\cdot|S^*)|_{\mathcal{A}} = P_\tau^+(\cdot|S^*)|_{\mathcal{A}}$ by Lem. 16 and 17. Hence, P_t^+ is proportional to P_τ^+ on \mathcal{A} , and thus on $(2^S)^*$ ($\subseteq \mathcal{A}$). So, P_t is proportional to P_τ on 2^S . ■

Proof of Lem. 13. Assume Ax. 1–6. Let U_t, P_t^+ ($t \in T$) be as above. Fix $t, \tau \in T$. Put $X := X_t \cap X_\tau$. For all $x, y \in X$, $x \succsim_t y \Leftrightarrow x \succsim_\tau y$ by outcome-preference stability (Prop. 2); so $U_t(x) \geq U_t(y) \Leftrightarrow U_\tau(x) \geq U_\tau(y)$ by Lem. 12. If U_t (and so U_τ) is constant on X , then U_t is an increasing affine transformation of U_τ on X . Now let U_t (and so U_τ) be non-constant on X . Let \mathcal{R} be as in Ax. 6. As (U_t, P_t^+) represents \succsim_t^+ restricted to the \mathcal{E}_t -measurable acts, $(U_t|_X, P_t^+|_{\mathcal{R}})$ represents \succsim_t^+ restricted further to \mathcal{R} -measurable acts mapping into X , i.e., to $F := \{f \in X^{\mathbb{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$. For analogous reasons, $(U_\tau|_X, P_\tau^+|_{\mathcal{R}})$ represents \succsim_τ^+ restricted to F . (In fact $P_t^+|_{\mathcal{R}} = P_\tau^+|_{\mathcal{R}}$ by Lem. 15.) Next I show that \succsim_t^+ and \succsim_τ^+ coincide on F . Let $f, g \in F$. As $f, g \in \mathcal{E}_t$ we may by Lem. 9 pick a $t' \in T$ such that $f, g \in F_{t'}$ and $S_{t'}$ harmlessly refines S_t . Analogously, as $f, g \in \mathcal{E}_\tau$ we may pick a $\tau' \in T$ such

⁵⁸ \mathcal{A} is the join of two algebras on S^* : $\{A^* : A \subseteq S\}$ and $\{A \cap S^* : A \in \mathcal{R}\}$ (\mathcal{R} 's trace in S^*).

that $f, g \in F_{\tau'}^*$ and $S_{\tau'}$ harmlessly refines S_{τ} . Now $f \succsim_t^+ g \Leftrightarrow f \succsim_{\tau'}^+ g$, as by Lem. 7 this reduces to $f_{t'} \succsim_{t'} g_{t'} \Leftrightarrow f_{\tau'} \succsim_{\tau'} g_{\tau'}$, which holds since $f (= (f_{t'})^* = (f_{\tau'})^*)$ and $g (= (g_{t'})^* = (g_{\tau'})^*)$ are measurable w.r.t. a robust algebra. As just shown, $(U_t|_X, P_t^+|_{\mathcal{R}})$ and $(U_{\tau}|_X, P_{\tau}^+|_{\mathcal{R}})$ represent the *same* relation on F ; note also that $U_t|_X$ and $U_{\tau}|_X$ are non-constant and $P_t^+|_{\mathcal{R}}$ and $P_{\tau}^+|_A$ are fine. So $U_t|_X$ is an increasing affine transformation of $U_{\tau}|_X$ by Lem. 5. ■

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 23 *Given an EU representation, Ax. 1, 3 and 5 hold.*

Proof. If $(U_t, P_t)_{t \in T}$ is such a representation, then Ax. 1 holds trivially, Ax. 5 holds by non-constancy of all U_t , and Ax. 3 holds by definition of conditional preference (using that non-null events $A \subseteq S_t$ have probabilities $P_t(A) \neq 0$). ■

Lemma 24 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1’s sense, Ax. 2 holds.*

Proof. Let $(U_t, P_t)_{t \in T}$ be a representation. Let $t, t' \in T$, $f, g \in F_t$, $f', g' \in F_{t'}$, and $A \subseteq S_t \cap S_{t'}$, such that $f_A = f'_A$, $g_A = g'_A$, $f_{S_t \setminus A} = g_{S_t \setminus A}$ and $f'_{S_{t'} \setminus A} = g'_{S_{t'} \setminus A}$. I must show $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$, i.e., $\mathbb{E}_{P_t}(U_t \circ f) \geq \mathbb{E}_{P_t}(U_t \circ g) \Leftrightarrow \mathbb{E}_{P_{t'}}(U_{t'} \circ f') \geq \mathbb{E}_{P_{t'}}(U_{t'} \circ g')$, or $\int_A U_t \circ f \, dP_t \geq \int_A U_t \circ g \, dP_t \Leftrightarrow \int_A U_{t'} \circ f' \, dP_{t'} \geq \int_A U_{t'} \circ g' \, dP_{t'}$ as $f_{S_t \setminus A} = g_{S_t \setminus A}$ and $f'_{S_{t'} \setminus A} = g'_{S_{t'} \setminus A}$. The latter holds as (i) P_t is proportional to $P_{t'}$ within A , (ii) $f_A = f'_A$ and $g_A = g'_A$, and (iii) U_t is an increasing affine transformation of $U_{t'}$ on $X_t \cap X_{t'}$ (where by (ii)–(iii) $U_t \circ f$ is an increasing affine transformation of $U_{t'} \circ f'$ on A , and $U_t \circ g$ is one of $U_{t'} \circ g'$ on A). ■

Lemma 25 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1’s sense, Ax. 4 holds.*

Proof. Assume $(U_t, P_t)_{t \in T}$ is a representation. Let $t, t' \in T$ such that $S := S_t = S_{t'}$, let $A, B \subseteq S$, and consider $x \succ_t y$ in X_t and $x' \succ_{t'} y'$ in $X_{t'}$. I claim that $x_A y_{S \setminus A} \succsim_t x_B y_{S \setminus B} \Leftrightarrow x'_A y'_{S \setminus A} \succsim_{t'} x'_B y'_{S \setminus B}$. Noting that $U_t(x) > U_t(y)$ (as $x \succ_t y$) and $U_{t'}(x') > U_{t'}(y')$ (as $x' \succ_{t'} y'$), the claimed equivalence reduces to the equivalence $P_t(A) \geq P_t(B) \Leftrightarrow P_{t'}(A) \geq P_{t'}(B)$, which holds as P_t is proportional (in fact, identical) to $P_{t'}$ on the full domain $2^S (= 2^{S_t} = 2^{S_{t'}})$. ■

Lemma 26 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1’s sense, with a fine agreed $\rho : \mathcal{R} \rightarrow [0, 1]$ in R3, then Ax. 6 holds in virtue of algebra \mathcal{R} .*

Proof. Let $(U_t, P_t)_{t \in T}$, ρ and \mathcal{R} be as assumed.

Claim 1: All $A \in \mathcal{R}$ are incorporable. Let $A \in \mathcal{R}$ and $t \in T$. As ρ is agreed, there is a $\tau \in T$ where $S_\tau = S_t \vee \{A, \bar{A}\}$ and P_τ^* extends P_t^* . W.l.o.g. $X_\tau = X_t$ (by independence between outcome and state awareness); so \succsim_τ is faithful to \succsim_t , using R1 and the fact that P_τ^* extends P_t^* .

Claim 2: \mathcal{R} is robust. Let $t_1, t_2 \in T$. Consider \mathcal{R} -determined acts $f_1, g_1 \in F_{t_1}$ and $f_2, g_2 \in F_{t_2}$ such that $f_1^* = f_2^* =: f$ and $g_1^* = g_2^* =: g$. I show that $f_1 \succsim_{t_1} g_1 \Leftrightarrow f_2 \succsim_{t_2} g_2$, i.e., $\mathbb{E}_{P_{t_1}}(U_{t_1} \circ f_1) \geq \mathbb{E}_{P_{t_1}}(U_{t_1} \circ g_1) \Leftrightarrow \mathbb{E}_{P_{t_2}}(U_{t_2} \circ f_2) \geq \mathbb{E}_{P_{t_2}}(U_{t_2} \circ g_2)$. As f and g are \mathcal{R} -measurable and $P_{t_i}^*(A) = \rho(A)$ for all $i \in \{1, 2\}$ and all $A \in \mathcal{R} \cap (2^{S_{t_i}})^*$, the desired equivalence reduces to $\mathbb{E}_\rho(U_{t_1} \circ f) \geq \mathbb{E}_\rho(U_{t_1} \circ g) \Leftrightarrow \mathbb{E}_\rho(U_{t_2} \circ f) \geq \mathbb{E}_\rho(U_{t_2} \circ g)$, which holds by R1 and the fact that $X_{t_1} \cap X_{t_2} \supseteq f(S_{t_1}), g(S_{t_2})$.

Claim 3: \mathcal{R} has the additional property required in Ax. 6. Let $t \in T$, $f \succ_t g$ in F_t , and $x \in X_t$. For any $\epsilon > 0$, pick (i) a finite partition $\mathcal{P}_\epsilon \subseteq \mathcal{R}$ of \mathbb{S} such that $\rho(A) \leq \epsilon$ for all $A \in \mathcal{P}_\epsilon$ (using ρ 's fineness) and (ii) a $t_\epsilon \in T$ such that $S_{t_\epsilon} = S_t \vee \mathcal{P}_\epsilon$, $X_{t_\epsilon} = X_t$, and $P_{t_\epsilon}^*$ extends P_t^* (using ρ 's agreedness and the independence between outcome and state awareness); let $f^\epsilon, g^\epsilon \in F_{t_\epsilon}$ be the acts equivalent to f and g respectively. It suffices to show that (*) for small enough $\epsilon > 0$, $\mathbb{E}_{P_{t_\epsilon}}(U \circ ((f^\epsilon)_{S_{t_\epsilon} \setminus A_{t_\epsilon}} x_{A_{t_\epsilon}})) > \mathbb{E}_{P_{t_\epsilon}}(U \circ g^\epsilon)$ for all $A \in \mathcal{P}_\epsilon$, and (**) for small enough $\epsilon > 0$, $\mathbb{E}_{P_{t_\epsilon}}(U \circ f^\epsilon) > \mathbb{E}_{P_{t_\epsilon}}(U \circ ((g^\epsilon)_{S_{t_\epsilon} \setminus A_{t_\epsilon}} x_{A_{t_\epsilon}}))$ for all $A \in \mathcal{P}_\epsilon$. As all U_{t_ϵ} have same domain as U_t , they are increasing affine transformations of U_t . W.l.o.g. let $U_{t_\epsilon} = U_t =: U$ for all $\epsilon > 0$. Given $\epsilon > 0$, each $(f^\epsilon)_{S_{t_\epsilon} \setminus A_{t_\epsilon}} x_{A_{t_\epsilon}}$ ($A \in \mathcal{P}_\epsilon$) differs from f^ϵ at most on A_{t_ϵ} , hence at most with $(P_{t_\epsilon}$ -)probability ϵ . Put $\Delta := \max_{x, y \in X_t} |U(x) - U(y)|$. Now $|\mathbb{E}_{P_{t_\epsilon}}(U \circ ((f^\epsilon)_{S_{t_\epsilon} \setminus A_{t_\epsilon}} x_{A_{t_\epsilon}})) - \mathbb{E}_{P_{t_\epsilon}}(U \circ f^\epsilon)| \leq \epsilon \Delta$ for all $A \in \mathcal{P}_\epsilon$. This implies (*) since $\mathbb{E}_{P_{t_\epsilon}}(U \circ f^\epsilon) = \mathbb{E}_{P_t}(U \circ f) > \mathbb{E}_{P_t}(U \circ g) = \mathbb{E}_{P_{t_\epsilon}}(U \circ g^\epsilon)$, where the ' $>$ ' holds as $f \succ_t g$, and both '=' hold as f^ϵ (g^ϵ) is equivalent to f (g) and $P_{t_\epsilon}^*$ extends P_t^* . An analogous argument shows (**). ■

C.4 Uniqueness of the representation

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

Lemma 27 *If a probability measure ρ on an algebra \mathcal{R} is agreed among probability measures P_t on 2^{S_t} ($t \in T$) which satisfy R2, then for each $t \in T$ there is a probability measure ρ_t on $\mathcal{R}_t := \mathcal{R} \vee (2^{S_t})^*$ which extends all P_τ^* for which S_τ is the join on S_t and a finite partition $\mathcal{P} \subseteq \mathcal{R}$ (so ρ_t extends ρ as ρ is agreed).*

Lemma 28 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1's sense with a fine agreed measure ρ on \mathcal{R} , and ρ_t and \mathcal{R}_t ($t \in T$) are as in Lem. 27, then, for all $t \in T$, (U_t, ρ_t) represents the restriction of \succsim_t^+ to $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{R}_t\text{-measurable}\}$ in Lem. 4's sense.*

Lemma 29 *If $(U_t, P_t)_{t \in T}$ and $(U'_t, P'_t)_{t \in T}$ are representations in Thm. 1's sense, then each P_t equals P'_t and each U_t is an increasing affine transformation of U'_t .*

Proof. Let $(U_t, P_t)_{t \in T}$ be a representation, with a fine agreed measure $\rho : \mathcal{R} \rightarrow [0, 1]$; so Ax. 1–6 hold. Let ρ_t and \mathcal{R}_t ($t \in T$) be as in Lem. 27 and 28. Let $(U'_t, P'_t)_{t \in T}$ be the representation defined in in App. C.2 under Ax. 1–6 using the objects \succsim_t^+ and P_t^+ (it was formerly denoted ‘ $(U_t, P_t)_{t \in T}$ ’). Fix $t \in T$. I show $P'_t = P_t$ and $U'_t = a_t U_t + b_t$ with $t_t > 0$ and $b_t \in \mathbb{R}$. As (U'_t, P_t^+) represents \succsim_t^+ on $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{E}_t\text{-measurable}\}$ (in Lem. 4's sense), $(U'_t, P_t^+|_{\mathcal{R}_t})$ represents the same relation as (U_t, ρ_t) on $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{R}_t\text{-measurable}\}$ by Lem. 28. So by Lem. 4 $\rho_t = P_t^+|_{\mathcal{R}_t}$ and U_t is an increasing affine transformation of U'_t . Finally, $P_t = P'_t$: for all $A \subseteq S_t$, $P_t(A) = P_t^*(A^*) = \rho_t(A^*) = P_t^+(A^*) = P'_t(A)$. ■

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 T into non-empty sets of contexts such that $t, t' \in T$ belong to the same set in Π if and only if $\mathbb{S}_t = \mathbb{S}_{t'}$. This yields for each $\Delta \in \Pi$ a (sub)structure $(X_t, S_t, \succsim_t)_{t \in \Delta}$ with exhaustive states, called the Δ -**substructure**, to which we may apply Thm. 1; let \mathbb{S}_Δ be its set of objective states. For all $t \in T$, let Δ_t be the member of Π containing t . Generalizing earlier objects, let \succsim_t^+ , \mathcal{E}_t and P_t^+ be defined as in App. C, but *w.r.t. the Δ_t -substructure* (which has exhaustive states, ensuring well-definedness). So \succsim_t^+ is a relation on $X_t^{\mathbb{S}_t} = X_t^{\mathbb{S}_{\Delta_t}}$ (not $X_t^{\mathbb{S}}$), and \mathcal{E}_t is an algebra on $\mathbb{S}_t = \mathbb{S}_{\Delta_t}$ (not \mathbb{S}). The *trace in \mathbb{S}'* ($\subseteq \mathbb{S}$) of an algebra \mathcal{R} on \mathbb{S} is the algebra on \mathbb{S}' given by $\mathcal{R}|_{\mathbb{S}'} := \{A \cap \mathbb{S}' : A \in \mathcal{R}\}$.⁵⁹

D.1 Sufficiency of the axioms

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

Lemma 30 *If Ax. 1–6 hold, then they hold for each Δ -substructure ($\Delta \in \Pi$).*

Now assume Ax. 1–6. By Lem. 30, each Δ -substructure ($\Delta \in \Pi$) satisfies Ax. 1–6. So by Thm. 1 each Δ -substructure ($\Delta \in \Pi$) has a representation $(U_t, P_t)_{t \in \Delta}$ in Thm. 1's sense. Joining these representations together, we obtain a grand system $(U_t, P_t)_{t \in T}$, which is now shown to represent the general structure.

⁵⁹A direct, non-reductive proof of Thm. 1 would also work, by generalizing App. C's proof strategy and defining the objects \mathcal{I} , \succsim_t^+ , \mathcal{E}_t and P_t^+ ($t \in T$) directly relative to the *general* structure; so \succsim_t^+ would be a relation on $X_t^{\mathbb{S}}$ (not $X_t^{\mathbb{S}_t}$), and \mathcal{E}_t an algebra on \mathbb{S} (not \mathbb{S}_t).

Lemma 31 Under Ax. 1–6, the above-defined $(U_t, P_t)_{t \in T}$ is an EU rationalization.

Proof. This property is inherited from the subsystems $(U_t, P_t)_{t \in \Delta}$ ($\Delta \in \Pi$). ■

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

Lemma 32 Given Ax. 1–6 and the above-defined functions P_t , if \mathcal{R} is an algebra as in Ax. 6 and for each Δ -substructure ($\Delta \in \Pi$) ρ_Δ is a fine probability measure on $\mathcal{R}|_{\mathbb{S}_\Delta}$ agreed among $(P_t)_{t \in \Delta}$, then the assignment $A \mapsto \rho_\Delta(A \cap \mathbb{S}_\Delta)$ defines a fine probability measure on \mathcal{R} independent of $\Delta \in \Pi$ and agreed among $(P_t)_{t \in T}$.

Lemma 33 Under Ax. 1–6, the above-defined measures P_t satisfy R3.

Proof. Assume Ax. 1–6, with \mathcal{R} as in Ax. 6. By Lem. 32 it suffices to show that for each $\Delta \in \Pi$ there is a fine probability measure ρ_Δ on $\mathcal{R}|_{\mathbb{S}_\Delta}$ agreed among the above-defined (Δ -)family $(P_t)_{t \in \Delta}$. Let $\Delta \in \Pi$. As $(P_t)_{t \in \Delta}$ satisfies R3, some fine measure ρ_Δ is agreed among $(P_t)_{t \in \Delta}$. By Lem. 30's proof, the Δ -substructure satisfies Ax. 6 in virtue of the trace algebra $\mathcal{R}|_{\mathbb{S}_\Delta}$. So by Thm. 1's proof we may w.l.o.g. let ρ_Δ have domain $\mathcal{R}|_{\mathbb{S}_\Delta}$. ■

Lemma 34 Under Ax. 1–6, the above-defined functions P_t satisfy R2.

Proof. The proof states literally like that of Lem. 14, the corresponding lemma under exhaustive states.⁶⁰ ■

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

Lemma 35 Under Ax. 1–6, for any contexts $t, t' \in T$, algebra \mathcal{R} as in Ax. 6 and \mathcal{R} -measurable functions $f, g : \mathbb{S} \rightarrow X_t \cap X_{t'}$, $f_{\mathbb{S}_t} \succsim_t^+ g_{\mathbb{S}_t} \Leftrightarrow f_{\mathbb{S}_{t'}} \succsim_{t'}^+ g_{\mathbb{S}_{t'}}$.

Lemma 36 Under Ax. 1–6, the above-defined functions U_t satisfy R1.

Proof. Suppose Ax. 1–6. Let $t, t' \in T$. Put $X := X_t \cap X_{t'}$. Let U_t and $U_{t'}$ be the above-defined functions. W.l.o.g. they are both non-constant on X .⁶¹ Let \mathcal{R} be as in Ax. 6, and ρ a (by Lem. 33 and its proof existing) fine agreed measure on \mathcal{R} . Let \geq be the relation on $F := \{f \in X^{\mathbb{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$ given by $f \geq g \Leftrightarrow f_{\mathbb{S}_\tau} \succsim_\tau^+ g_{\mathbb{S}_\tau}$ for some (hence by Lem. 35 any) $\tau \in \{t, t'\}$. To show that $U_t|_X$ is an

⁶⁰Two remarks are due. First, in Claim 1 t' (τ') automatically belongs to Δ_t , ($\Delta_{\tau'}$), and the Lemmas 18 and 19 (which had been stated under exhaustive states) are applied to the Δ_t - ($\Delta_{\tau'}$ -)substructure (with indeed have exhaustive states). Second, Lem. 4 is applied to the general structure; although it had been stated under exhaustive states, it holds in general, as one easily checks.

⁶¹The argument is like in the proof of Lemma 13, but using Lem. 31 rather than 12.

increasing affine transformations of $U_{t'}|_X$, I prove that $(U_t|_X, \rho)$ and $(U_{t'}|_X, \rho)$ both represent $(X, (S, R), \geq)$ in Lem. 4's sense. Let $\tau \in \{t, t'\}$. I show $\mathbb{E}_\rho(U_\tau \circ f) \geq \mathbb{E}_\rho(U_\tau \circ g) \Leftrightarrow f \geq g$ for all $f, g \in F$. As $f \geq_\tau g$ reduces to $f_{\mathbb{S}_\tau} \succsim_\tau^+ g_{\mathbb{S}_\tau}$, hence to $\mathbb{E}_{P_\tau^+}(U_\tau \circ f_{\mathbb{S}_\tau}) \geq \mathbb{E}_{P_\tau^+}(U_\tau \circ g_{\mathbb{S}_\tau})$, it suffices to show that $\mathbb{E}_\rho(U_\tau \circ f) = \mathbb{E}_{P_\tau^+}(U_\tau \circ f_{\mathbb{S}_\tau})$ for all $f \in F$. Let $f \in F$; I prove $\rho(f^{-1}(x)) = P_\tau^+(f_{\mathbb{S}_\tau}^{-1}(x))$ for all $x \in X$. By Lem. 32 (and Lem. 33's proof), we may write $\rho = \rho_{\Delta_\tau}(\cdot \cap \mathbb{S}_\tau)$ for a fine measure ρ_{Δ_τ} on $\mathcal{R}|_{\mathbb{S}_\tau}$ agreed among $(P_\delta)_{\delta \in \Delta_\tau}$. Not only ρ_{Δ_τ} , but also $P_\tau^+|_{\mathcal{R}|_{\mathbb{S}_\tau}}$ is agreed among $(P_\delta)_{\delta \in \Delta_\tau}$, by Lem. 4 (applied to the Δ_τ -substructure). So $\rho_{\Delta_\tau} = P_\tau^+|_{\mathcal{R}|_{\mathbb{S}_\tau}}$. For any $x \in X$, $\rho(f^{-1}(x)) = \rho_{\Delta_\tau}(f^{-1}(x) \cap \mathbb{S}_\tau) = \rho_{\Delta_\tau}(f_{\mathbb{S}_\tau}^{-1}(x)) = P_\tau^+(f_{\mathbb{S}_\tau}^{-1}(x))$, where the first equality holds as $\rho = \rho_{\Delta_\tau}(\cdot \cap \mathbb{S}_\tau)$, and the last one as $\rho_{\Delta_\tau} = P_\tau^+|_{\mathcal{R}|_{\mathbb{S}_\tau}}$. ■

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 substructures via this technical lemma (shown in App. F):

Lemma 37 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1's sense (with a fine agreed measure ρ on an algebra \mathcal{R}), then each subsystem $(U_t, P_t)_{t \in \Delta}$ ($\Delta \in \Pi$) represents the Δ -substructure in Thm. 1's sense (with a fine agreed measure ρ_Δ on $\mathcal{R}|_{\mathbb{S}_\Delta}$ given by $\rho(\cdot) = \rho_\Delta(\cdot \cap \mathbb{S}_\Delta)$).*

Lemma 38 *If $(U_t, P_t)_{t \in T}$ and $(U'_t, P'_t)_{t \in T}$ are representations in Thm. 1's sense, then any P_t equals P'_t and any U_t is an increasing affine transformation of U'_t .*

Proof. This property follows via Lem. 37 from the uniqueness property for substructures, which is guaranteed by Thm. 1 applied to substructures. ■

Lemma 39 *If $(U_t, P_t)_{t \in T}$ is a representation in Thm. 1's sense, with a fine agreed measure on an algebra \mathcal{R} , then Ax. 6 holds in virtue of algebra \mathcal{R} .*

Proof. Let $(U_t, P_t)_{t \in T}$ and \mathcal{R} be as specified. Let $t \in T$, $f \succ_t g$ in F_t , and $x \in X_t$. Put $\Delta := \Delta_t$. By Lem. 37, $(U_t, P_t)_{t \in \Delta}$ represents the Δ -substructure, with a fine agreed measure on $\mathcal{R}|_{\mathbb{S}_\Delta}$. So by Lem. 26, Ax. 6 holds for this substructure in virtue of algebra $\mathcal{R}|_{\mathbb{S}_\Delta}$. Hence one can partition \mathbb{S}_Δ into $R_1, \dots, R_n \in \mathcal{R}|_{\mathbb{S}_\Delta}$ and pick a $t' \in T$ where $S_{t'} = S_t \vee \{R_1, \dots, R_n\}$ (so any R_i is representable by an $A_i \subseteq S_{t'}$), $X_{t'} \supseteq X_t$ (so $F_{t'}$ contains acts f' equivalent to f and g' equivalent to g), and $f'_{S_{t'} \setminus A_i} x_{A_i} \succ_{t'} g'$ and $f' \succ_{t'} g'_{S_{t'} \setminus A_i} x_{A_i}$ for all A_i . Each R_i is in $\mathcal{R}|_{\mathbb{S}_\Delta}$; so $R_i = B_i \cap \mathbb{S}_\Delta$ for a $B_i \in \mathcal{R}$. W.l.o.g. B_1, \dots, B_n partition \mathbb{S} .⁶² Ax. 6 for the general structure follows since $S_{t'} = S_t \vee \{B_1, \dots, B_n\}$ (as $S_{t'} = S_t \vee \{R_1, \dots, R_n\}$ and each R_i matches B_i within \mathbb{S}_Δ) and any B_i is, like R_i , represented by A_i . ■

⁶²Otherwise 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 ' $\cup (\overline{\bigcup_{i=1}^n B_i})$ ').

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 \mathbb{S} be given. When referring to Thm. 2's third revision rule, we take it to be stated in the following (equivalent) way: the functions P_t agree (at least) on some fine probability measure on \mathcal{R} .

First assume Ax. 1–5 and $6_{\mathcal{R}}-8_{\mathcal{R}}$. As Ax. $6_{\mathcal{R}}-8_{\mathcal{R}}$ imply Ax. 6, Thm. 1's representation $(U_t, P_t)_{t \in T}$ exists. This representation satisfies even Thm. 2's modified third rule, as the agreed measure can be defined on any algebra as in Ax. 6, e.g., on the risky algebra \mathcal{R} , using Lem. 15 (under exhaustive states) or more generally Lem. 33.

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. 26 (under exhaustive states) or more generally Lem. 39. This implies Ax. $6_{\mathcal{R}}-8_{\mathcal{R}}$. ■

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 $t = t'$, because incorporability of all $A \in \mathcal{R}$ comes for free (see Rem. 24) and whenever $A_1, \dots, A_n \subseteq S_t$ partition S_t and represent some $R_1, \dots, R_n \in \mathcal{R}$, then we may choose R_1, \dots, R_n such as to partition \mathbb{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 $t \in T$, acts $f \succ_t g$ in F_t , and an outcome $x \in X_t$. Pick $t' \in T$, $A_1, \dots, A_n \subseteq S_{t'}$ and $f', g' \in F_{t'}$ as given by Ax. 6; so $f'_{S_{t'} \setminus A_i} x_{A_i} \succ_{t'} g'$ and $f' \succ_{t'} g'_{S_{t'} \setminus A_i} x_{A_i}$ for all A_i . By state fineness, $S_{t'} = S_t$; so $A_1, \dots, A_n \subseteq S_t$ and (as also $X_{t'} \supseteq X_t$) $f' = f$ and $g' = g$. So, by preference stability (see Prop. 1, which uses Ax. 2), $f_{S_t \setminus A_i} x_{A_i} \succ_t g$ and $f \succ_t g_{S_t \setminus A_i} x_{A_i}$ for all A_i . ■

Lem. 2 is provable analogously to Lem. 1.

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

Proof of Lem. 3. Let $t \in T$. Let \mathcal{R}_1 and \mathcal{R}_2 be the sets in (1) and (2) respectively. Since \mathcal{R}_1 is obviously an algebra, it suffices to show that $\mathcal{E}_t = \mathcal{R}_1 = \mathcal{R}_2$.

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

Claim 2: $\mathcal{R}_2 \subseteq \mathcal{E}_t$. Let $A \in \mathcal{R}_2$. Then we may pick a context $\tau \in T$ such that

S_τ harmlessly refines S_t and $A \in (2^{S_\tau})^*$. So $A = B^*$ for some $B \subseteq S_\tau$. Hence, $A \in \mathcal{E}_t$.

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

$$A = \cup_{I \in \mathcal{P}} \cup_{a \in S_t: a \cap I \in B} (a \cap I) = \cup_{I \in \mathcal{P}} (I \cap (\cup_{a \in S_t: a \cap I \in B} a)).$$

So A is a Boolean combination of members of \mathcal{I} and $(2^{S_t})^*$. Hence, $A \in \mathcal{R}_1$. ■

Lemma 40 *Given any finite set $\mathcal{J} \subseteq \mathcal{I}$, there is a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbb{S} refining each $\{J, \bar{J}\}$ ($J \in \mathcal{J}$) such that for all contexts $t \in T$ there is a $\tau \in T$ for which $S_\tau = S_t \vee \mathcal{P}$ and \succsim_τ is faithful to \succsim_t .*

Proof. 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 $\mathcal{P} = \{\mathbb{S}\}$. 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 and \mathcal{P}_2 respectively. 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 $t \in T$ we may first pick a context $\tau' \in T$ such that $S_{\tau'} = S_t \vee \mathcal{P}_1$ and $\succsim_{\tau'}$ is faithful to \succsim_t , and then pick a context $\tau \in T$ such that $S_\tau = S_{\tau'} \vee \mathcal{P}_2 = S_t \vee \mathcal{P}$ and \succsim_τ is faithful to $\succsim_{\tau'}$, hence to \succsim_t . ■

Proof of Lem. 6. Assume Ax. 2, and let $t, \tau \in T$ such that $X_t \subseteq X_\tau$ and $S_\tau = S_t \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbb{S} . Using Lem. 40, pick a finite partition $\mathcal{P}' \subseteq \mathcal{I}$ of \mathbb{S} refining \mathcal{P} such that there are $t', \tau' \in T$ where $S_{t'} = S_t \vee \mathcal{P}$, $S_{\tau'} = S_\tau \vee \mathcal{P}$, $\succsim_{t'}$ is faithful to \succsim_t , and $\succsim_{\tau'}$ is faithful to \succsim_τ . Now $S_{t'} = S_{\tau'}$ (as $S_\tau = S_t \vee \mathcal{P}$) and $X_{t'}, X_{\tau'} \supseteq X_t$ (as by faithfulness $X_{t'} \supseteq X_t$ and $X_{\tau'} \supseteq X_\tau$, and as $X_\tau \supseteq X_t$). So $F_{t'} \cap F_{\tau'} \supseteq X_t^{S_{t'}}$. Hence, by Ax. 2, $\succsim_{\tau'}$ matches $\succsim_{t'}$ on $X_t^{S_{t'}}$, hence is (like $\succsim_{t'}$) faithful to \succsim_t . As $\succsim_{\tau'}$ is faithful to \succsim_t and \succsim_τ (and as each $f \in F_t$ is objectively equivalent to some $g \in F_\tau$), \succsim_τ is faithful to \succsim_t . ■

Proof of Lem. 7. Assume Ax. 2 and $f \succsim_t^+ g$, where $t \in T$ and $f, g \in X_t^\mathbb{S}$.

(a) Let $\tau \in T$ satisfy the conditions (i)–(ii) in Def. 37. I show that $f_\tau \succsim_\tau g_\tau$. As $f \succsim_t^+ g$, we have $f_{\tau'} \succsim_{\tau'} g_{\tau'}$ for some $\tau' \in T$ satisfying these conditions. As S_τ and $S_{\tau'}$ harmlessly refine S_t , we may pick finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{I}$ of \mathbb{S} such that $S_\tau = S_t \vee \mathcal{P}$ and $S_{\tau'} = S_t \vee \mathcal{P}'$. Using Lem. 40, there is a partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbb{S} which refines \mathcal{P} and \mathcal{P}' and contexts $\sigma, \sigma' \in T$ such that $S_\sigma = S_\tau \vee \mathcal{Q}$, $S_{\sigma'} = S_{\tau'} \vee \mathcal{Q}$, \succsim_σ is faithful to \succsim_τ , and $\succsim_{\sigma'}$ to $\succsim_{\tau'}$. Note that $S_\sigma = S_{\sigma'} = S_t \vee \mathcal{Q}$, so that $f_\sigma = f_{\sigma'}$ and $g_\sigma = g_{\sigma'}$. Hence, by preference stability (Prop. 1), $f_\sigma \succsim_\sigma g_\sigma \Leftrightarrow f_{\sigma'} \succsim_{\sigma'} g_{\sigma'}$. This equivalence reduces to $f_\tau \succsim_\tau g_\tau \Leftrightarrow f_{\tau'} \succsim_{\tau'} g_{\tau'}$ by faithfulness of \succsim_σ to \succsim_τ and of $\succsim_{\sigma'}$ to $\succsim_{\tau'}$. As $f_{\tau'} \succsim_{\tau'} g_{\tau'}$, it follows that $f_\tau \succsim_\tau g_\tau$.

(b) Pick any $\tau \in T$ satisfying (i)–(ii) in Def. 37 hold. Pick a context $\tau' \in T$ such that $S_{\tau'} = S_\tau$ and $X_{\tau'} = X_t$. Clearly, also $\tau' \in T$ satisfies (i)–(ii) in Def. 37. Moreover, $\succsim_{\tau'}$ is faithful to \succsim_t by Lem. 6. ■

Proof of Lem. 8. Consider a $t \in T$ and a finite $\mathcal{B} \subseteq \mathcal{E}_t$. For each $B \in \mathcal{B}$, pick a partition \mathcal{P}_A of \mathbb{S} refining $\{B, \overline{B}\}$ and having the property stated in the definition of weak incorporability (note that $\mathcal{P}_A \subseteq \mathcal{I}$). 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 $\tau_1 \in T$ such that $S_{\tau_1} = S_t \vee \mathcal{P}_{B_1}$ and \succsim_{τ_1} is faithful to \succsim_t ; second, a context $\tau_2 \in T$ such that $S_{\tau_2} = S_{\tau_1} \vee \mathcal{P}_{B_2}$ and \succsim_{τ_2} is faithful to \succsim_{τ_1} ; and so on for contexts τ_3, \dots, τ_n . Let $\tau := \tau_n$. Property (i) holds because each B_i is representable in context τ_i , hence in context τ . Property (ii) holds as $S_\tau = S_t \vee \mathcal{P}$ with $\mathcal{P} := \mathcal{P}_{B_1} \vee \dots \vee \mathcal{P}_{B_n}$. Property (iii) holds by transitivity of faithfulness. ■

Proof of Lem. 9. This claim follows from Lem. 8 applied to the (finite) set $\mathcal{B} = \{f^{-1}(x) : f \in \mathcal{G}, x \in X_t\}$, by noting that for any $\tau \in T$ F_τ^* is characterizable as the set of $(2^{S_\tau})^*$ -measurable function from \mathbb{S} to X_t . ■

Proof of Lem. 10. Assume Ax. 2 and 5, let $t \in T$ and consider Lem. 5's generalized Savage structure, with set of acts denoted F . Let f, g, A be as specified. First assume $f \succsim_{t,A}^+ g$. Then, by definition, $f' \succsim_t^+ g'$ for some $f', g' \in F$ agreeing with respectively f and 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 $|X_t| \geq 2$ by Ax. 5). Using Lem. 9, we pick a $\tau \in T$ such that $f, g, f', g', h \in F_\tau^*$ and S_τ harmlessly refines S_t (and \succsim_τ is faithful to \succsim_t , which is only needed if the modified equivalence is to be proved). As $h \in F_\tau^*$, A is representable. As $f' \succsim_t^+ g'$, we have $f'_\tau \succsim_\tau g'_\tau$ by Lem. 7. Noting that f'_τ and g'_τ agree respectively with f_τ and g_τ on A_τ and with each other outside A_τ (because of inheriting these properties from analogous ones of f' and g'), we deduce $f_\tau \succsim_{\tau,A} g_\tau$.

Conversely, assume that $f_\tau \succsim_{\tau,A_\tau} g_\tau$ for some $\tau \in T$ satisfying the specified properties. Then there are two functions in F_τ – we may write them as f'_τ and g'_τ for certain $f, g \in F_\tau^*$ – such that $f'_\tau \succsim_\tau g'_\tau$ and such that f'_τ and g'_τ agree with respectively f_τ and g_τ on A_τ and with each other outside A_τ . From $f'_\tau \succsim_\tau g'_\tau$ (and the properties of τ) it follows that $f' \succsim_t^+ g'$, which in turn implies that $f \succsim_{t,A}^+ g$ since f' and g' agree with respectively f and g on A and with each other outside A (they inherit this behaviour from f_τ and g_τ because $f = (f_\tau)^*$, $g = (g_\tau)^*$ and $A = (A_\tau)^*$). ■

Proof of Lem. 11. Assume Ax. 2 and let $t \in T$. Consider Lem. 5's generalized Savage structure and an event $A \in \mathcal{A}$.

First assume A is non-null. Then there are $f, g \in F$ such that $f_{\overline{A}} = g_{\overline{A}}$ and $f \not\prec_t^+ g$. Pick any $h \in F$ taking one value on A and another on \overline{A} (h exists as $|X_t| \geq 2$ by the fact that F contains distinct functions f, g). By Lem. 9, we may choose a context $\tau \in T$ such that $f, g, h \in F_\tau^*$ and S_τ faithfully refines S_t (and such that \succsim_τ is faithful to \succsim_t , something we need to add when proving the equivalence in its modified version). As $h \in F_\tau^*$, A is representable in context τ , i.e., A_τ is

defined. As $f \not\sim_t^+ g$, we have $f_\tau \not\sim_\tau g_\tau$, which (since f_τ and g_τ agree outside A_τ) shows that A_τ is non-null.

Conversely, assume A_τ is non-null (under \succsim_τ) for some $\tau \in T$ with the specified properties. Then we may pick two non-indifferent acts in F_τ which agree outside A ; we may write them as f_τ and g_τ for some $f, g \in F_\tau^*$. Since $f_\tau \not\sim_\tau g_\tau$, we have $f \not\sim_t^+ g$ by Lem. 7. So, as f and g agree outside A , A is non-null. ■

Proof of Lem. 18. Assume Ax. 2, 4 and 5. Let $t \in T$ and $A, B \subseteq \mathbb{S}$. By Ax. 5 there are $x \succ_t y$ in X_t .

First assume $A \succsim_t^+ B$. Then there exist $x, y \in X_t$ such that $x \succ_t^+ y$ and $x_A y_{\bar{A}} \succsim_t^+ x_B y_{\bar{B}}$. So by Lem. 9 there is a context $\tau \in T$ such that $x_A y_{\bar{A}}, x_B y_{\bar{B}} \in F_\tau^*$ (hence, A and B are representable), S_τ harmlessly refines S_t , and \succsim_τ is faithful to \succsim_t (the latter is needed when proving the modified equivalence). By Lem. 7, it follows that $x_{S_\tau} \succ_\tau y_{S_\tau}$ and $(x_A y_{\bar{A}})_\tau \succsim_\tau (x_B y_{\bar{B}})_\tau$. In other words $x \succ_\tau y$ and $x_{A_\tau} y_{S_\tau \setminus A_\tau} \succsim_\tau x_{B_\tau} y_{S_\tau \setminus B_\tau}$. So, $A_\tau \succsim_\tau B_\tau$.

Conversely, assume $A_\tau \succsim_\tau B_\tau$ for a $\tau \in T$ such that A and B are representable and S_τ harmlessly refines S_t . It follows that $A, B \in \mathcal{E}_t$. So by Lem. 8 we may pick a context $\tau' \in T$ such that A and B are representable, $S_{\tau'}$ harmlessly refines S_t , and $\succsim_{\tau'}$ is faithful to \succsim_t . In particular, $X_{\tau'} = X_t$. As $A_\tau \succsim_\tau B_\tau$ we have $A_{\tau'} \succsim_{\tau'} B_{\tau'}$ by belief stability (see Prop. 3, which uses Ax. 2, 4 and 5). Hence there are $x', y' \in X_{\tau'} (= X_t)$ such that $x' \succ_{\tau'} y'$ and $x'_{A_{\tau'}} y'_{S_{\tau'} \setminus A_{\tau'}} \succsim_{\tau'} x'_{B_{\tau'}} y'_{S_{\tau'} \setminus B_{\tau'}}$. In other words, $(x'_\mathbb{S})_{\tau'} \succ_{\tau'} (y'_\mathbb{S})_{\tau'}$ and $(x'_A y'_{\bar{A}})_{\tau'} \succsim_{\tau'} (x'_B y'_{\bar{B}})_{\tau'}$. By Lem. 7 it follows that $x'_\mathbb{S} \succ_t^+ y'_\mathbb{S}$ (i.e., $x' \succ_t^+ y'$) and $x'_A y'_{\bar{A}} \succsim_t^+ x'_B y'_{\bar{B}}$. So $A \succsim_t^+ B$. ■

Proof of Lem. 19. Assume Ax. 2, 4, 5 and 6. Let $A \succsim_t^+ B$, where $t \in T$ and $A, B \subseteq \mathbb{S}$. Let $\tau \in T$ satisfy the conditions stated. I show that $A_\tau \succsim_\tau B_\tau$. As $A \succsim_t^+ B$, we have $A_{\tau'} \succsim_{\tau'} B_{\tau'}$ for *some* $\tau' \in T$ satisfying the analogous conditions, by Lem. 18 (which uses Ax. 2, 4 and 5). As S_τ and $S_{\tau'}$ harmlessly refine S_t , we may pick finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{I}$ of \mathbb{S} such that $S_\tau = S_t \vee \mathcal{P}$ and $S_{\tau'} = S_t \vee \mathcal{P}'$. Using Lem. 40, there is a partition $\mathcal{Q} \subseteq \mathcal{I}$ of \mathbb{S} which refines \mathcal{P} and \mathcal{P}' and contexts $\sigma, \sigma' \in T$ such that $S_\sigma = S_\tau \vee \mathcal{Q}$, $S_{\sigma'} = S_{\tau'} \vee \mathcal{Q}$, \succsim_σ is faithful to \succsim_τ , and $\succsim_{\sigma'}$ to $\succsim_{\tau'}$. Note that $S_\sigma = S_{\sigma'} = S_t \vee \mathcal{Q}$, so that $A_\sigma = A_{\sigma'}$ and $B_\sigma = B_{\sigma'}$. So, by comparative-belief stability (Prop. 3, which uses Ax. 2, 4, 5 and 6), $A_\sigma \succsim_\sigma B_\sigma \Leftrightarrow A_{\sigma'} \succsim_{\sigma'} B_{\sigma'}$. This equivalence reduces to $A_\tau \succsim_\tau B_\tau \Leftrightarrow A_{\tau'} \succsim_{\tau'} B_{\tau'}$ by faithfulness of \succsim_σ to \succsim_τ and of $\succsim_{\sigma'}$ to $\succsim_{\tau'}$. As $A_{\tau'} \succsim_{\tau'} B_{\tau'}$, it follows that $A_\tau \succsim_\tau B_\tau$. ■

Proof of Lem. 20. 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 $t, t' \in T$. I assume $A \succsim_t^+ B$ and have to prove $A \succsim_{t'}^+ B$. By Lem. 20, as $A \succsim_t^+ B$ we have $A_\tau \succsim_\tau B_\tau$ for a $\tau \in T$ such that A and B are representable and S_τ harmlessly refines S_t . Meanwhile, as $A, B \in \mathcal{R} \subseteq \mathcal{I} \subseteq \mathcal{E}_{t'}$, by Lem. 8 there exists a $\tau' \in T$ such that A and B are

representable in context τ' and $S_{\tau'}$ harmlessly refines $S_{t'}$. As $A_{\tau} \succsim_{\tau} B_{\tau}$ and as A and B belong to a robust algebra (i.e., \mathcal{R}), we have $A_{\tau'} \succsim_{\tau'} B_{\tau'}$ by belief stability on robust algebras (Prop. 3). So $A \succsim_{t'}^+ B$ by Lem. 18. ■

Proof of Lem. 21. Assume Ax. 1–6. Let \mathcal{R} be as in Ax. 6. Let U_t and P_t^+ ($t \in T$) be as defined above.

Claim 1: $P_t^+|_{\mathcal{R}}$ is fine for all $t \in T$. Let $t \in T$. The pair (U_t, P_t^+) represents \succsim_t^+ on $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{E}_t\text{-measurable}\}$ (in Lem. 4's sense). So $(U_t, P_t^+|_{\mathcal{R}})$ represents \succsim_t^+ on $\{f \in X_t^{\mathbb{S}} : f \text{ is } \mathcal{R}\text{-measurable}\}$. By Lem. 5 (applied with $\mathcal{A} = \mathcal{R}$), there is a fine probability measure on \mathcal{R} representing the (belief) relation induced by \succsim_t^+ on \mathcal{R} . This measure represents the same (belief) relation on \mathcal{R} as $P_t^+|_{\mathcal{R}}$, and thus coincides with $P_t^+|_{\mathcal{R}}$ by Lem. 16 and 17. So $P_t^+|_{\mathcal{R}}$ is fine.

Claim 2: $\rho := P_t^+|_{\mathcal{R}}$ is the same for all $t \in T$. Let $t, t' \in T$. By Lem. 20, the functions $P_t^+|_{\mathcal{R}}$ and $P_{t'}^+|_{\mathcal{R}}$ are ordinally equivalent. Since these are fine probability measures by Claim 1, they must coincide by Lem. 16 and 17. ■

Lemma 41 *Under Ax. 1, for any context $t \in T$, two functions $f, g \in X_t^{\mathbb{S}}$ are \succsim_t^+ -comparable (i.e., $f \succsim_t^+ g$ or $g \succsim_t^+ f$) if and only if both are \mathcal{E}_t -measurable.*

Proof. Assume Ax. 1. Let $t \in T$ and $f, g \in X_t^{\mathbb{S}}$. First assume f and g are comparable under \succsim_t^+ . Then f_{τ} and g_{τ} are comparable for some context $\tau \in T$ such that $f, g \in F_{\tau}^*$ and $S_{\tau} = S_t \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{I}$ of \mathbb{S} . Since $f, g \in F_{\tau}^*$, f and g are $(2^{S_{\tau}})^*$ -measurable, which implies \mathcal{E}_t -measurability as $(2^{S_{\tau}})^* = (2^{S_t \vee \mathcal{P}})^* \subseteq \mathcal{E}_t$. Conversely, if f and g are \mathcal{E}_t -measurable, then by Lem. 9 there is a context $\tau \in T$ such that $f, g \in F_{\tau}^*$ and S_{τ} harmlessly refines S_t . By Ax. 1, $f_{\tau} \succsim_{\tau} g_{\tau}$ or $g_{\tau} \succsim_{\tau} f_{\tau}$, which implies that $f \succsim_t^+ g$ or $g \succsim_t^+ f$. ■

Proof of Lem. 22. Assume Ax. 1 and 2. Let $t, \tau \in T$. Assume S_{τ} harmlessly refines S_t . Then $\mathcal{E}_t = \mathcal{E}_{\tau}$ by definition of extrapolated algebras. Now suppose that in addition \succsim_{τ} is faithful to \succsim_t . In view of Lem. 41 it suffices to show that \succsim_t^+ and \succsim_{τ}^+ coincide on the set of \mathcal{E}_{τ} - (or respectively \mathcal{E}_t -)measurable functions in $X_t^{\mathbb{S}}$. Let $f, g \in X_t^{\mathbb{S}}$ be \mathcal{E}_{τ} - (hence, \mathcal{E}_t -)measurable. Then by Lem. 9 there is a context $r \in T$ such that $f, g \in F_r^*$ and S_r harmlessly refines S_{τ} , hence, also S_t . We have $f \succsim_t^+ g \Leftrightarrow f \succsim_{\tau}^+ g$ because each side is equivalent to $f_r \succsim_r g_r$ by Lem. 7. ■

Proof of Lem. 27. Let $(U_t, P_t)_{t \in T}$, ρ , \mathcal{R} and \mathcal{R}_t be as specified. Fix $t \in T$.

Claim 1: $\mathcal{R}_t = \cup_{\tau \in T: S_{\tau} = S_t \vee \mathcal{P}} \text{ for some finite partition } \mathcal{P} \subseteq \mathcal{R} \text{ of } \mathbb{S}(2^{S_{\tau}})^*$. This claim is provable analogously to the proof of Lem. 3.

Claim 2: For all $\tau \in T$ and finite partitions $\mathcal{P} \subseteq \mathcal{R}$ of \mathbb{S} , there is a $\sigma \in T$ such that $S_{\sigma} = S_{\tau} \vee \mathcal{P}$ and P_{σ}^* extends P_{τ}^* . Consider such τ and \mathcal{P} . Write $\mathcal{P} = \{I_1, \dots, I_n\}$. As each I_i is incorporable and ρ is agreed, we can let $\tau_0 := \tau$ and successively pick $\tau_1, \dots, \tau_n \in T$ such that, for each τ_i , $P_{\tau_i}^*$ extends $P_{\tau_{i-1}}^*$ and

$S_{\tau_i} = S_{\tau_{i-1}} \vee \{I_i, \overline{I_i}\}$. Clearly, $P_{\tau_n}^*$ extends P_τ^* and $S_{\tau_n} = S_\tau \vee \{I_1, \overline{I_1}\} \vee \dots \vee \{I_n, \overline{I_n}\} = S_\tau \vee \mathcal{P}$.

Claim 3: The measures P_τ^* with $S_\tau = S_t \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbb{S} agree pairwise on the domain overlap. Let $\tau, \tau' \in T$ such that $S_\tau = S_t \vee \mathcal{P}$ and $S_{\tau'} = S_t \vee \mathcal{P}'$ for finite partitions $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{R}$ of \mathbb{S} . I show that P_τ^* and $P_{\tau'}^*$ agree on the domain overlap. By Claim 2, there are $\sigma, \sigma' \in T$ such that $S_\sigma = S_\tau \vee \mathcal{P}'$, $S_{\sigma'} = S_{\tau'} \vee \mathcal{P}$, P_σ^* extends P_τ^* , and $P_{\sigma'}^*$ extends $P_{\tau'}^*$. It suffices to show that $P_\sigma^* = P_{\sigma'}^*$. As P_σ and $P_{\sigma'}$ have the same domain $2^{S_\sigma} = 2^{S_{\sigma'}} (= 2^{S_t \vee \mathcal{P} \vee \mathcal{P}'})$, $P_\sigma = P_{\sigma'}$ by R2, whence $P_\sigma^* = P_{\sigma'}^*$.

Claim 4: All desired properties are met by the function ρ_t which to each $A \in \mathcal{R}_t$ assigns $P_\tau^*(A)$ for a (by Claim 1 existing and by Claim 3 arbitrary) $\tau \in T$ such that $S_\tau = S_t \vee \mathcal{P}$ for a finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbb{S} . By definition, ρ_t extends all P_τ^* such that $S_\tau = S_t \vee \mathcal{P}$ for some finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbb{S} . It remains to show that ρ_t is a probability measure. Clearly, $\rho_t(\mathbb{S}) = P_t^*(\mathbb{S}) = 1$. Now consider disjoint $A, B \in \mathcal{R}_t$. By Claim 1 we may pick $\tau, \tau' \in T$ such that $A \in (2^{S_\tau})^*$, $B \in (2^{S_{\tau'}})^*$, $S_\tau = S_t \vee \mathcal{P}$ and $S_{\tau'} = S_t \vee \mathcal{Q}$, for finite partitions $\mathcal{P}, \mathcal{Q} \subseteq \mathcal{R}$ of \mathbb{S} . By Claim 2 we may pick a $\sigma \in T$ such that $S_\sigma = S_t \vee \mathcal{P} \vee \mathcal{Q}$. Now $A, B \in (2^{S_\sigma})^*$ and $\rho_t(A) + \rho_t(B) = P_\sigma^*(A) + P_\sigma^*(B) = P_\sigma^*(A \cup B) = \rho_t(A \cup B)$. ■

Proof of Lem. 28. Let $(U_t, P_t)_{t \in T}$, ρ , \mathcal{R} , ρ_t and \mathcal{R}_t be as specified. Fix $t \in T$. The proof proceeds in two steps.

Claim 1: $\mathbb{E}_{\rho_t}(U_t \circ f) \geq \mathbb{E}_{\rho_t}(U_t \circ g) \Leftrightarrow f \succsim_t^+ g$ for all \mathcal{R}_t -measurable $f, g \in X_t^\mathbb{S}$. Let $f, g \in X_t^\mathbb{S}$ be \mathcal{R}_t -measurable. We may pick a finite partition $\mathcal{P} \subseteq \mathcal{R}$ of \mathbb{S} such that f and g are $(2^{S_t \vee \mathcal{P}})^*$ -measurable, and then pick a $\sigma \in T$ such that $S_\sigma = S_t \vee \mathcal{P}$ (for details see Claims 1 and 2 in Lem. 27's proof). W.l.o.g. $X_\sigma = X_t$ by independence of outcome and state awareness. The desired equivalence holds as $\mathbb{E}_{\rho_t}(U_t \circ f) \geq \mathbb{E}_{\rho_t}(U_t \circ g) \Leftrightarrow \mathbb{E}_{P_\sigma}(U_\sigma \circ f_\sigma) \geq \mathbb{E}_{P_\sigma}(U_\sigma \circ g_\sigma) \Leftrightarrow f_\sigma \succsim_\sigma g_\sigma \Leftrightarrow f \succsim_t^+ g$, where the last ' \Leftrightarrow ' holds by Lem. 7 and the first ' \Leftrightarrow ' holds as ρ_t extends P_σ^* and U_σ is an increasing affine transformation of U_t (by R1 and the fact that $X_\sigma = X_t$).

Claim 2: ρ_t is fine and U_t is non-constant. Non-constancy of U_t holds as U_t is part of representation in Thm. 1's sense. Further, as $\mathcal{R} \subseteq \mathcal{R}_t \subseteq \mathcal{E}_t$ where by Lem. 26 \mathcal{R} is an algebra as in Ax. 6 (and \mathcal{E}_t is the extrapolated algebra), we know by Lem. 4 that the restriction of \succsim_t^+ to $\{f \in X_t^\mathbb{S} : f \text{ is } \mathcal{R}_t\text{-measurable}\}$ has a representation (U'_t, P'_t) in Lem. 4's sense; in particular, P'_t is a fine probability measure on \mathcal{R}_t . By Claim 1, ρ represents the same probability order on \mathcal{R}_t as P'_t . Hence $\rho = P'_t$ by Lem. 16 and 17. So ρ is itself fine. ■

From now on the restriction to exhaustive states is lifted.

Lemma 42 *If an algebra \mathcal{R} on \mathbb{S} is robust, then w.r.t. any Δ -substructure ($\Delta \in \Pi$) the (trace) algebra $\mathcal{R}|_{\mathbb{S}_\Delta}$ on \mathbb{S}_Δ is robust.*

Proof. Consider a robust algebra \mathcal{R} on \mathbb{S} , a $\Delta \in \Pi$, contexts $t, t' \in \Delta$, and $\mathcal{R}|_{\mathbb{S}_\Delta}$ -determined acts $f, g \in F_t$ and $f', g' \in F_{t'}$ such that f is equivalent to f' , and g to g' . We must show that $f \succsim_t g \Leftrightarrow f' \succsim_{t'} g'$. This holds because (i) \mathcal{R} is robust, and (ii) the $\mathcal{R}|_{\mathbb{S}_\Delta}$ -determinedness of the four acts implies (in fact, is equivalent to) their \mathcal{R} -determinedness. ■

Lemma 43 *Assume Ax. 2. If an objective event $A \subseteq \mathbb{S}$ is incorporable, then w.r.t. any Δ -substructure ($\Delta \in \Pi$) $A \cap \mathbb{S}_\Delta$ is incorporable.*

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

- $S_\tau = S_t \vee \{A \cap \mathbb{S}_\Delta, \mathbb{S}_\Delta \setminus (A \cap \mathbb{S}_\Delta)\}$: this holds because

$$S_\tau = S_{t'} = S_t \vee \{A, \bar{A}\} = S_t \vee \{A \cap \mathbb{S}_\Delta, \mathbb{S}_\Delta \setminus (A \cap \mathbb{S}_\Delta)\}.$$

- \succsim_τ is faithful to \succsim_t : As $\succsim_{t'}$ is faithful to \succsim_t , $X_{t'} \supseteq X_t$, i.e., $X_{t'} \supseteq X_\tau$. So, as also $S_{t'} = S_\tau$, the relation \succsim_τ is the restriction of $\succsim_{t'}$ to $F_\tau (\subseteq F_{t'})$ by preference stability (see Prop. 1, which uses Ax. 2). Hence, not only $\succsim_{t'}$, but also \succsim_τ is faithful to \succsim_t . ■

Proof of Lem. 30. Let $\Delta \in \Pi$. The Δ -substructure 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 \mathbb{S} as in Ax. 6 (for the general structure). I show that the substructure satisfies Ax. 6 in virtue of the trace algebra $\mathcal{R}|_{\mathbb{S}_\Delta}$. By Lem. 42 and 43, $\mathcal{R}|_{\mathbb{S}_\Delta}$ is, w.r.t. the substructure, a robust algebra (on \mathbb{S}_Δ) composed of incorporable objective events. Now consider a $t \in \Delta$, acts $f \succ_t g$ in F_t , and an $x \in X_t$. By Ax. 6 for the general structure, we may partition \mathbb{S} into some $R_1, \dots, R_n \in \mathcal{R}$ such that, in some context $t' \in T$ where $S_{t'} = S_t \vee \{R_1, \dots, R_n\}$ (so each R_i is representable by an $A_i \subseteq S_{t'}$) and $X_{t'} \supseteq X_t$ (so $F_{t'}$ contains acts f' equivalent to f and g' equivalent to g), we have $f'_{S_{t'} \setminus A_i} x_{A_i} \succ_{t'} g'$ and $f' \succ_{t'} g'_{S_{t'} \setminus A_i} x_{A_i}$ for all A_i . To complete the proof of Ax. 6 for the substructure, it suffices to note that (i) $t' \in \Delta$ because $\mathbb{S}_{t'} = \mathbb{S}_t$ (as $S_{t'} = S_t \vee \{R_1, \dots, R_n\}$), and (ii) \mathbb{S}_Δ is partitioned into (the non-empty sets among) $R_1 \cap \mathbb{S}_\Delta, \dots, R_n \cap \mathbb{S}_\Delta \in \mathcal{R}|_{\mathbb{S}_\Delta}$, where each such $R_i \cap \mathbb{S}_\Delta$ is represented by A_i . ■

Proof of Lem. 32. Assume Ax. 1–6. Let $\mathcal{R}, (P_t)_{t \in T}$, and $(\rho_\Delta)_{\Delta \in \Pi}$ be as specified. Each ρ_Δ induces a function π_Δ on \mathcal{R} via $\pi_\Delta(A) := \rho_\Delta(A \cap \mathbb{S}_\Delta)$ ($A \in \mathcal{R}$).

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 $\pi_\Delta(\mathbb{S}) = \rho_\Delta(\mathbb{S}_\Delta) = 1$

and as for disjoint $A, B \in \mathcal{R}$ we have $\pi_\Delta(A \cup B) = \rho_\Delta((A \cup B) \cap S_\Delta) = \rho_\Delta((A \cap S_\Delta) \cup (B \cap S_\Delta)) = \rho_\Delta(A \cap S_\Delta) + \rho_\Delta(B \cap S_\Delta) = \pi_\Delta(A) + \pi_\Delta(B)$. Second, I show fineness. Let $\epsilon > 0$. As ρ_Δ is fine, we may partition \mathbb{S}_Δ into $R_1, \dots, R_n \in \mathcal{R}|_{\mathbb{S}_\Delta}$ such that $\rho_\Delta(R_i) < \epsilon$ for all R_i . Each R_i is in $\mathcal{R}|_{\mathbb{S}_\Delta}$; so we write $R_i = B_i \cap \mathbb{S}_\Delta$ for $B_i \in \mathcal{R}$. We may take B_1, \dots, B_n to partition \mathbb{S} , by the argument in fn. 62. Now $\pi_\Delta(B_i) = \rho_\Delta(B_i \cap \mathbb{S}_\Delta) = \rho_\Delta(R_i) < \epsilon$ for all B_i . So π_Δ is fine.

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. 16 and 17, 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 $t \in \Delta$ in which A and B are representable. The events $A_t, B_t (\subseteq S_t)$ representing A and B respectively also represent $A \cap \mathbb{S}_\Delta$ and $B \cap \mathbb{S}_\Delta$ respectively. Now (*) $\pi_\Delta(A) \geq \pi_\Delta(B) \Leftrightarrow A_t \succsim_t B_t$, since $\pi_\Delta(A) \geq \pi_\Delta(B) \Leftrightarrow \rho_\Delta(A \cap \mathbb{S}_\Delta) \geq \rho_\Delta(B \cap \mathbb{S}_\Delta) \Leftrightarrow P_t(A_t) \geq P_t(B_t) \Leftrightarrow A_t \succsim_t B_t$, where the second equivalence holds as ρ_Δ is agreed among $(P_\delta)_{\delta \in \Delta}$ and A_t and B_t represent $A \cap \mathbb{S}_\Delta$ and $B \cap \mathbb{S}_\Delta$ respectively. Analogously, as A and B are incorporable we may pick a context $t' \in \Delta'$ where A and B are representable; as before, (**) $\pi_{\Delta'}(A) \geq \pi_{\Delta'}(B) \Leftrightarrow A_{t'} \succsim_{t'} B_{t'}$. As A and B belong to the robust algebra \mathcal{R} , $A_t \succsim_t B_t \Leftrightarrow A_{t'} \succsim_{t'} B_{t'}$ by Prop. 5, 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 agreed among the P_t ($t \in T$). For any $t \in T$, recall that P_t^* is the function of (representable) objective events $A \subseteq \mathbb{S}$ induced by P_t ; let P_t^{**} be the analogous function induced by P_t w.r.t. the Δ_t -substructure. So P_t^{**} is a function of (representable) $A \subseteq \mathbb{S}_{\Delta_t}$. Now let $A \in \mathcal{R}$, $t \in T$, and $\Delta := \Delta_t$. As ρ_Δ is agreed among $(P_r)_{r \in \Delta}$, there is a $\tau \in \Delta$ such that P_τ^{**} extends P_t^{**} , $S_\tau = S_t \vee \{(A \cap \mathbb{S}_\Delta), \mathbb{S}_\Delta \setminus (A \cap \mathbb{S}_\Delta)\}$ and $P_\tau^{**}(A \cap \mathbb{S}_\Delta) = \rho_\Delta(A \cap \mathbb{S}_\Delta)$. Turning to the general structure, we must show that (i) P_τ^* extends P_t^* , (ii) $S_\tau = S_t \vee \{A, \mathbb{S} \setminus A\}$, and (iii) $P_\tau^*(A) = \pi_\Delta(A)$. Claim (i) holds as for all $B \subseteq \mathbb{S}$ in the domain of P_t^* , hence of P_τ^* , $P_t^*(B) = P_t^{**}(B \cap \mathbb{S}_\Delta) = P_\tau^{**}(B \cap \mathbb{S}_\Delta) = P_\tau^*(B)$, where the second equality holds as P_τ^{**} extends P_t^{**} , while the first and third hold as, in each of the contexts t and τ , B has the same representation as $B \cap \mathbb{S}_\Delta$. Claim (ii) holds as $S_\tau = S_t \vee \{(A \cap \mathbb{S}_\Delta), \mathbb{S}_\Delta \setminus (A \cap \mathbb{S}_\Delta)\} = S_t \vee \{A, \mathbb{S} \setminus A\}$. Claim (iii) holds as $P_\tau^*(A) = P_\tau^{**}(A \cap \mathbb{S}_\Delta) = \rho_\Delta(A \cap \mathbb{S}_\Delta) = \pi_\Delta(A) = \rho(A)$. ■

Proof of Lem. 35. Assume Ax. 1–6. Let t, t', \mathcal{R}, f, g be as given. Note that $\mathcal{R}|_{\mathbb{S}_t}$ is included in the extrapolated algebra \mathcal{E}_t , as by Lem. 43 $\mathcal{R}|_{\mathbb{S}_t}$ consists of (w.r.t. the Δ_t -substructure) incorporable objective events. As f and g are \mathcal{R} -measurable, $f_{\mathbb{S}_t}$ and $g_{\mathbb{S}_t}$ are $\mathcal{R}|_{\mathbb{S}_t}$ -measurable, so (as $\mathcal{R}|_{\mathbb{S}_t} \subseteq \mathcal{E}_t$) \mathcal{E}_t -measurable. Hence by Lem. 9 (applied to the substructure) we may pick a $\tau \in \Delta_t$ such that $f_{\mathbb{S}_t} = \hat{f}^*$ and $g_{\mathbb{S}_t} = \hat{g}^*$ for certain $\hat{f}, \hat{g} \in F_\tau$ and S_τ harmlessly refines S_t ; so, by Lem. 7, $f_{\mathbb{S}_t} \succsim_t^+ g_{\mathbb{S}_t} \Leftrightarrow \hat{f} \succsim_\tau \hat{g}$. By analogous arguments, we may pick a $\tau' \in \Delta_{t'}$ such that $f_{\mathbb{S}_{t'}} = \tilde{f}^*$ and $g_{\mathbb{S}_{t'}} = \tilde{g}^*$ for certain $\tilde{f}, \tilde{g} \in F_{\tau'}$ and $S_{\tau'}$ harmlessly refines $S_{t'}$; so,

$f_{\mathbb{S}_{t'}} \succsim_{t'}^+ g_{\mathbb{S}_{t'}} \Leftrightarrow \hat{f} \succsim_{\tau'} \tilde{g}$. As $f_{\mathbb{S}_t} \succsim_t^+ g_{\mathbb{S}_t} \Leftrightarrow \hat{f} \succsim_{\tau} \hat{g}$ and $f_{\mathbb{S}_{t'}} \succsim_{t'}^+ g_{\mathbb{S}_{t'}} \Leftrightarrow \tilde{f} \succsim_{\tau'} \tilde{g}$, it suffices to show that $\hat{f} \succsim_{\tau} \hat{g} \Leftrightarrow \tilde{f} \succsim_{\tau'} \tilde{g}$. 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 $\mathbb{S}_t = \mathbb{S}_{\tau}$ and \tilde{f}^* on $\mathbb{S}_{t'} = \mathbb{S}_{\tau'}$) 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 $\mathbb{S}_t = \mathbb{S}_{\tau}$ and \tilde{g}^* on $\mathbb{S}_{t'} = \mathbb{S}_{\tau'}$). ■

Proof of Lem. 37. Let $(U_t, P_t)_{t \in T}$, ρ and \mathcal{R} be as assumed. Let $\Delta \in \Pi$. w.r.t. the Δ -substructure, the subsystem $(U_t, P_t)_{t \in \Delta}$ is still an EU rationalization obeying R1 and R2, as all this is inherited from the full system. It suffices to show R3. Now

$$B \cap \mathbb{S}_{\Delta} = C \cap \mathbb{S}_{\Delta} \Rightarrow \rho(B) = \rho(C) \text{ for all } B, C \in \mathcal{R}, \quad (2)$$

because any $B, C \in \mathcal{R}$ are (by ρ 's agreedness) representable in some context $t \in \Delta$, for which $\rho(B) = P_t^*(B) = P_t^*(B \cap \mathbb{S}_{\Delta})$ (the last equality holds as B and $B \cap \mathbb{S}_{\Delta}$ are represented by the same subjective event) and similarly $\rho(C) = P_t^*(C) = P_t^*(C \cap \mathbb{S}_{\Delta})$. Now the function ρ induces a function $\rho_{\Delta} : \mathcal{R}|_{\mathbb{S}_{\Delta}} \rightarrow [0, 1]$ by defining, for any $A \in \mathcal{R}|_{\mathbb{S}_{\Delta}}$, $\rho_{\Delta}(A) := \rho(B)$, where B is some (hence by (2) *any*) member of \mathcal{R} such that $B \cap \mathbb{S}_{\Delta} = A$. By construction, $\rho_{\Delta}(B \cap \mathbb{S}_{\Delta}) = \rho(B)$ 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 $\rho_{\Delta}(\mathbb{S}_{\Delta}) = \rho(\mathbb{S}) = 1$, and since any disjoint $A, A' \in \mathcal{R}|_{\mathbb{S}_{\Delta}}$ may be written as $A = B \cap \mathbb{S}_{\Delta}$ and $A' = B' \cap \mathbb{S}_{\Delta}$ 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 \mathbb{S} into $B_1, \dots, B_n \in \mathcal{R}$ such that $\rho(B_i) < \epsilon$ for all $i = 1, \dots, n$; consequently \mathbb{S}_{Δ} is partitioned⁶³ into $B_1 \cap \mathbb{S}_{\Delta}, \dots, B_n \cap \mathbb{S}_{\Delta} \in \mathcal{R}|_{\mathbb{S}_{\Delta}}$ and $\rho_{\Delta}(B_i \cap \mathbb{S}_{\Delta}) = \rho(B_i) < \epsilon$ for all i .

Claim 2: ρ_{Δ} is agreed (w.r.t. the Δ -substructure). For any $\delta \in \Delta$, let P_{δ}^* be (as usual) the function of representable objective events induced by P_{δ} , and let P_{δ}^{**} be the analogous function defined w.r.t. the Δ -substructure; so P_{δ}^* is a function of (representable) subsets of \mathbb{S} , whereas P_{δ}^{**} is a function of (representable) subsets of \mathbb{S}_{Δ} . Now consider a $t \in \Delta$ and an $A \in \mathcal{R}|_{\mathbb{S}_{\Delta}}$. We need to show that there is a $\tau \in \Delta$ such that (a) P_{τ}^{**} extends P_t^{**} , (b) $S_{\tau} = S_t \vee \{A, \mathbb{S}_{\Delta} \setminus A\}$, and (c) $P_{\tau}^{**}(A) = \rho_{\Delta}(A)$. Write A as $B \cap \mathbb{S}_{\Delta}$ for some $B \in \mathcal{R}$. As ρ is agreed w.r.t. the general structure, there is a $\tau \in T$ such that P_{τ}^* extends P_t^* , $S_{\tau} = S_t \vee \{B, \overline{B}\}$, and $P_{\tau}^*(B) = \rho(B)$. We may assume w.l.o.g. that $\tau \in \Delta$, as one may verify using independence between outcome and state awareness and Rem. 21. Condition (a) holds because, when restricted to subsets of \mathbb{S}_{Δ} , P_{τ}^* coincides with P_{τ}^{**} and P_t^* coincides with P_t^{**} .

⁶³in the broad sense that allows some of $B_1 \cap \mathbb{S}_{\Delta}, \dots, B_n \cap \mathbb{S}_{\Delta}$ to be empty

Condition (b) holds because $S_t \vee \{B, \overline{B}\} = S_t \vee \{A, \mathbb{S}_\Delta \setminus A\}$. Condition (c) holds because, as $A \subseteq \mathbb{S}_\Delta$, we have $P_\tau^{**}(A) = P_\tau^*(A)$ 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
- Bolker, E. (1966) Functions Resembling Quotients of Measures, *Transactions of the American Mathematical Society* 124: 292–312
- Dekel, E., Lipman, B. L., Rustichini, A. (1998) Standard state-space models preclude unawareness, *Econometrica* 66: 159–73
- Dietrich, F., List, C. (2016a) Mentalism versus Behaviourism in Economics: a philosophy-of-science perspective, *Economics and Philosophy* 32(2): 249–281
- Dietrich, F., List, C. (2016b) Reason-based choice and context-dependence: an explanatory framework, *Economics and Philosophy* 32(2): 175–229
- 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
- Heifetz, A., Meier, M., Schipper, B. (2006) Interactive unawareness, *Journal of Economic Theory* 130: 78–94
- Hill, B. (2010) Awareness Dynamics, *Journal of Philosophical Logic* 39: 113–37
- Jeffrey, R. C. (1983) *The Logic of Decision*, 2nd ed., Chicago: University of Chicago Press
- 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
- Kochov, A. (2016) A behavioral definition of unforeseen contingencies, *working paper*, University of Rochester
- 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

- Schipper, B. (2013) Awareness-dependent subjective expected utility, *International Journal of Game Theory* 42: 725-753
- 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
- Tversky, A., Koehler, D.J. (1994) Support theory: a nonextensional representation of subjective probability, *Psychological Rev.* 101: 547-67
- Wakker, P. (1981) Agreeing probability measures for comparative probability structures, *Annals of Statistics* 9: 658-62