Categorical decision theory and global subjective expected utility representations

Marcus Pivato[*]

June 16, 2024

Abstract

One agent may encounter many sources of uncertainty and many menus of outcomes, which can be combined together into many different decision problems. There may be *analogies* between different uncertainty sources (or different outcome menus). Some uncertainty sources (or outcome menus) may exhibit internal *symmetries*. The agent may also have different levels of *awareness*. In some situations, the state spaces and outcome spaces have additional mathematical structure (e.g. a topology, metric, or differentiable structure), and feasible acts must respect this structure (i.e. they must be continuous, short, or smooth functions). In other situations, the agent might only be aware of a set of abstract "acts", and be unable to specify explicit state spaces and outcome spaces. We introduce a modelling framework that addresses all of these issues. We then define and axiomatically characterize a subjective expected utility representation that is "global" in two senses. First: it posits probabilistic beliefs for all uncertainty sources and utility functions over all outcome menus, which simultaneously rationalize the agents' preferences across all possible decision problems, and which are consistent with the aforementioned analogies, symmetries, and awareness levels. Second: it applies in many mathematical environments (i.e. categories), making it unnecessary to develop a new theory for each one.

Keywords: uncertainty; unawareness; analogy; category theory; Anscombe-Aumann.
JEL classification: D81.

L'algèbre est généreuse, elle donne souvent plus qu'on ne lui demande.
—Jean le Rond d'Alembert

1 Introduction

The standard model of rationality under uncertainty is *subjective expected utility* (SEU). This was axiomatically characterized by landmark results of Savage (1954) and Anscombe and Aumann (1963). The Savage framework is the basis for almost all contemporary

[*]Centre d'Economie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne

research in decision theory. But it also has several shortcomings. It lacks a natural way to simultaneously model the same agent facing multiple decision problems, or facing the same decision problem with different levels of awareness, or access to different information sources. It also lacks a natural way to incorporate analogies between different decision problems, or symmetries within each decision problem. It allows acts to be arbitrary functions between arbitrary sets, so it does not easily adapt to environments where state spaces and outcome spaces have additional structure (e.g. a topology) and acts must respect this structure. Finally, it assumes that the agent already has a complete mental model of the space of possible states of nature and the menu of possible outcomes.

In the literature, each of these problems has been addressed by augmenting the basic Savage framework in some way. This paper will present a modelling framework which allows us to address *all* of these problems at the same time. First, we shall discuss the aforementioned problems in greater detail. Then we shall discuss our proposed solution.

Multiple decision problems. On different occasions, a single agent may be confronted with different sources of uncertainty (e.g. horse races, financial markets, weather, traffic) and different possible sets of outcomes (e.g. financial gains or losses, social status, physical comfort), in different combinations. A holistic model of her beliefs and preferences should be able to represent her attitudes towards *all* of these decision problems simultaneously.

For example, suppose the agent faced N sources of uncertainty represented by state spaces $\mathcal{S}_1, \ldots, \mathcal{S}_N$, along with M possible menus of outcomes, represented by sets $\mathcal{X}_1, \ldots, \mathcal{X}_M$. For any $n \in [1 \ldots N]$ and $m \in [1 \ldots M]$, the agent might be confronted with a decision problem which involves Savage acts from $\mathcal{S}_n$ into $\mathcal{X}_m$. We could construct $N \times M$ distinct SEU representations to deal with these different decision problems. But it would be strange if the agent could have *different* beliefs about $\mathcal{S}_1$ depending on whether it was combined with $\mathcal{X}_1$ or $\mathcal{X}_2$. Likewise, it would be strange if she had a *different* utility function over $\mathcal{X}_1$ depending on whether it was combined with $\mathcal{S}_1$ or $\mathcal{S}_2$. Furthermore, this would fail to recognize relationships between these different spaces —for example, that $\mathcal{X}_1$ and $\mathcal{X}_2$ overlap, or that $\mathcal{S}_1$ is a projection of $\mathcal{S}_2$. It would also be unwieldy to carry around M different beliefs for each of the N state spaces, and N different utility functions for each of the M outcome spaces,

The obvious solution would be to construct a single "grand" state space $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_N$ and a single "grand" outcome space $\mathcal{X} = \mathcal{X}_1 \cup \cdots \cup \mathcal{X}_N$. Equipped with a single probabilistic belief over the $\mathcal{S}$, and a single utility function over $\mathcal{X}$, the agent could then construct SEU preferences for any of the $N \times M$ decision problems she might face.

But such a model is needlessly profligate. It requires the agent to form beliefs about combinations of events that she will never observe: for any events $\mathcal{A}_1 \subseteq \mathcal{S}_1$ and $\mathcal{A}_2 \subseteq \mathcal{S}_2$, she must have belief about $\mathcal{A}_1 \times \mathcal{A}_2$, even though (by hypothesis), she will never encounter a situation where she can simultaneously verify whether $\mathcal{A}_1$ and $\mathcal{A}_2$ obtain. Likewise, she must form meaningful utility comparisons between outcomes which will never appear on the same menu: for any outcome in $\mathcal{X}_1$ and outcome in $\mathcal{X}_2$ (or more generally, any lotteries over outcomes), she must be able to say whether she prefers the first to the second, even though (by hypothesis), this is a choice that she will never be required to make.

Against this, one might argue: even if the agent does not "need" a single belief about $\mathcal{S}$, surely such beliefs should be available "in principle", e.g. for an external analyst modelling the agent's behaviour. Likewise, even if the agent does not "need" a single utility function over all of $\mathcal{X}$, such a utility function should be available "in principle". But there are two problems with this response, one normative and one descriptive. From a *descriptive* perspective, it presents the external analyst with an observability problem: if we can never observe the agent making choices between $\mathcal{X}_1$-lotteries and $\mathcal{X}_2$-lotteries, then how can we impute a single utility function over all $\mathcal{X}$ to this agent? If we never observe her betting on arbitrary subsets of $\mathcal{S}_1 \times \mathcal{S}_2$, then how can we impute a single belief about all of $\mathcal{S}$?

More fundamentally, it is not clear whether such unified beliefs and utility functions are even well-defined. And from a normative perspective, it is not even clear that they *should* be well-defined. This is true even when there is some degree of overlap between different menus, or correlation between different sources of uncertainty. Although the agent's beliefs about particular sources of uncertainty may all be compatible with one another whenever two sources can be compared against each other, there might be no coherent way to combine them into a single, coherent probabilistic belief system about the entire universe. Likewise, although the agent's preferences over different menus may agree whenever these menus overlap, there might be no way to combines these local preferences into a single, coherent global preference order. We shall illustrate with two examples.

Example 1.1. An employee in a large firm knows that during the next year, she will transferred to one of three divisions. Every worker in the firm must undergo an annual performance review. But due to limited resources, the three divisions conduct these reviews during three different but overlapping periods of the calendar year. Division 1 conducts its reviews from January to June. Division 2 conducts them from May until October, and Division 3 conducts them from September until February. Workers in each division can choose the month of their performance review. The employee does not yet know *when* or *where* she will be transferred, but in advance of the transfer, she is asked, "*Hypothetically*, if you moved to Division 1 (or 2 or 3), when would you want your performance review?"

The employee always wants her performance review to be as late as possible. So she reports the following preferences. For Division 1: Jan $\prec$ Feb $\prec$ Mar $\prec$ Apr $\prec$ May $\prec$ Jun. Likewise, for Division 2, May $\prec$ Jun $\prec \cdots \prec$ Oct, and for Division 3, Sep $\prec$ Oct $\prec \cdots \prec$ Feb. Formally, there are three possible menus of outcomes: $\mathcal{X}_1 = \{\text{Jan}, \ldots, \text{Jun}\}$, $\mathcal{X}_2 = \{\text{May}, \ldots, \text{Oct}\}$, and $\mathcal{X}_3 = \{\text{Sep}, \ldots, \text{Feb}\}$. The employee has transitive preferences over each menu, which even agree on the overlap between menus. But there is no "global" preference order over all twelve months that is consistent with these local preferences. $\diamondsuit$

Example 1.2. A doctor encounters patients who may have one of several diseases. There are three main organ systems that may be implicated (Systems 1, 2, and 3). Each of these systems can be have one of two (mutually exclusive) diseases. System 1 can have either Disease a or Disease b, System 2 can have either c or d, and System 3 can have either e or f, Some patients have disease in both Systems 1 and 2, or both Systems 2 and 3, or both Systems 1 and 3. But the doctor never encounters patients who have diseases in all three Systems, because such a combination is always immediately fatal.

From experience, the doctor has formed beliefs about the probability of each disease and each (non-fatal) combination of diseases. We can represent this by introducing six state spaces: $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_3$ $\mathcal{S}_{12}$, $\mathcal{S}_{23}$, and $\mathcal{S}_{13}$, where $\mathcal{S}_j$ represents diseases involving only System j, and $\mathcal{S}_{jk}$ represents diseases involving both System j and System k. Thus, $\mathcal{S}_1 := \{a, b\}$, $\mathcal{S}_2 := \{c, d\}$, $\mathcal{S}_3 := \{e, f\}$, $\mathcal{S}_{12} := \mathcal{S}_1 \times \mathcal{S}_2$, $\mathcal{S}_{23} := \mathcal{S}_2 \times \mathcal{S}_3$, and $\mathcal{S}_{13} := \mathcal{S}_1 \times \mathcal{S}_3$. Probabilistic beliefs are given as follows. The beliefs on the two-element spaces $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_3$ are given by $\mu_1 = \mu_2 = \mu_3 = (0.5, 0.5)$, where these vectors have the obvious meaning. The beliefs on the four-element spaces $\mathcal{S}_{12}$, $\mathcal{S}_{23}$ and $\mathcal{S}_{13}$ are given by the following tables:

μ_{12}	c	d	
a	0.1	0.4	0.5
b	0.4	0.1	0.5
	0.5	0.5	

μ_{23}	e	f	
c	0.1	0.4	0.5
d	0.4	0.1	0.5
	0.5	0.5	

μ_{13}	e	f	
a	0.1	0.4	0.5
b	0.4	0.1	0.5
	0.5	0.5	

Observe that these beliefs are consistent; for example, both μ_{12} and μ_{23} generate the same marginal beliefs $(0.5, 0.5)$ concerning $\{c, d\}$. But there is no "global" probabilistic belief on $\mathcal{S}_1 \times \mathcal{S}_2 \times \mathcal{S}_3$ that is compatible with these beliefs. (See Appendix for proof.)[1] $\diamondsuit$

Information, awareness and model underspecification. Even if we focus on a single decision problem, the same agent may perceive this decision problem in different ways, depending on her level of awareness or the sources of information available to her. One could develop a distinct SEU representation for each awareness level or information source, but this would be unparsimonious. Furthermore, it would not ensure that these different SEU representations were "compatible" with one another —i.e. that they have the same utility function over outcomes, and that they have the "same" beliefs (to the extent that this is meaningful across different awareness levels or information sources). It would be better to develop a single SEU representation which encompasses all awareness/information levels at the same time.

A more fundamental problem with Savage framework is its assumption that the agent can explicitly imagine all possible "states of nature" and all possible "outcomes", and can then conceptualize each "act" as a function mapping states to outcomes. As a descriptive model of how actual humans make decisions, this might be unrealistic (especially if these spaces are large). Even as a normative theory of *ideal* decision-making, it might be too demanding. It might not be an accurate or parsimonious description of how the agent perceives her world. She is certainly aware that she confronts different sources of uncertainty, and that there are various actions that she might take. But she might not be able to explicitly list every possible outcome of these actions, nor every possible contingency which could determine the outcome of an action. Ideally, a decision theory should be neutral about how the agent mentally represents the decision problems she faces. (See Machina 2003 for an excellent discussion of these issues.)

Analogies and symmetries. Agents often exploit *analogies* between novel decision problems and familiar problems. This motivates the use of case studies in business schools,

[1]Similar examples appear in Vorob'ev (1962) and Hammond (2022).

medical schools, and officer training academies. An analogy reveals that two decision problems share important structural features, despite superficial dissimilarities; in other words, it establishes a *homomorphism* between them (Amarante, 2015). So an SEU representation should incorporate any available analogies when specifying beliefs and utility functions.

Agents also often utilize *symmetries* in decision problems —especially in the absence of other information. For example, suppose an agent is confronted with an urn containing an unknown number of red and black balls, and she must bet on the colour of the next ball drawn from the urn. If this is the *only* information available to the agent, then she will likely be indifferent between betting on red and betting on black. For similar reasons, agents normally assign probability 1/2 to each face of an unfamiliar coin, and probability 1/38 to each slot in an unfamiliar roulette wheel. Such reasoning is formalized by the well-known *Principle of Insufficient Reason*, which was implicit in early work on probability theory (e.g. Laplace, 1820), explicitly articulated by Keynes (1921, Ch. IV), and generalized to the *Principle of Transformation Groups* by Jaynes (1968, §VII; 2003, Ch.12). An SEU representation theorem should incorporate such symmetries into the agent's beliefs.

Diversity of mathematical modelling environments. Different mathematical environments may arise in response to different modelling requirements. In some models, the state space and outcome space are measurable spaces, and acts are measurable functions. But in other models, they might be topological spaces and continuous functions, or differentiable manifolds and smooth functions. Instead of functions, it might be appropriate to represent actions using *correspondences* (Ghirardato, 2001), or some other kind of "generalized mapping" from states to outcomes. As noted above, it may sometimes be appropriate to consider a more abstract representation of the decision problem, without *any* explicit specification of "states" or "outcomes". Of course, one could develop separate SEU representation theories to handle each of these situations. For instance, in the setting of topological spaces and continuous functions, Zhou (1999) and Pivato and Vergopoulos (2020) have obtained versions of the Anscombe-Aumann and Savage representation theorems, respectively. But such duplication of effort seems inefficient. It would be better to develop a single, "general purpose" theory, which can be applied in a variety of mathematical environments.

A categorical approach. This paper uses methods from category theory to develop a general SEU representation theory that addresses all of the problems raised above. Category theory is a branch of abstract algebra which provides a powerful and versatile analytical framework in a wide variety of mathematical domains. Roughly speaking, a *category* is a family of mathematical "objects" (e.g. sets, topological spaces, algebraic structures, etc.), connected by a network of relationships (called "morphisms"). The basic philosophy of category theory is that all the relevant mathematical properties of an object should be describable in terms of its *relationships* to other objects in the same family. This allows one to study mathematical objects without making any claims about their "internal structure". This is reminiscent of the methodology of modern economic theory, which seeks models of decision-making that make the least possible ontological commitments about the "internal structure" of agents. This suggests that categorical methods could be useful in decision theory. This paper is a preliminary exploration of this possibility.

The paper introduces *global SEU representations*. These are "global" in two senses. First, they provide a single, holistic representation for an agent's preferences across multiple state spaces, multiple outcome spaces, and/or multiple levels of awareness or information —a representation which incorporates any analogies or symmetries that the agent perceives. Second, the framework and results can be applied in a variety of categories: measurable spaces and measurable maps, topological spaces and continuous maps, differentiable manifolds and smooth maps, etc. State spaces and outcome spaces are represented as *objects* in the category, while acts, analogies, symmetries, and differential awareness/information are represented as *morphisms* between these objects. We make no assumption about the internal structure of these objects and morphisms. All the ingredients of the SEU representation are obtained using the structure of the category itself.

Prior literature. There is already a considerable literature addressing some of the issues mentioned earlier. For example, some models of unawareness involve not a single state space, but a *lattice* of state spaces, representing greater or lesser degrees of awareness of the underlying uncertainty (Heifetz et al., 2006; Ahn and Ergin, 2010; Hayashi, 2012; Schipper, 2013; Dietrich, 2018). In other models, the state space evolves dynamically over time, as the agent learns of new technological possibilities or new contingencies (Karni and Vierø, 2013, 2015, 2017; Dominiak and Tserenjigmid, 2018). Rather than taking the state space as a primitive in the model, some approaches to decision theory treat the *acts* as primitives, and define the state space "endogenously" as a set of possible preferences over these acts (Kreps, 1979; Dekel et al., 2001), or as a set of possible mappings from acts into "outcomes" (Fishburn 1970; Schmeidler and Wakker 1990; Karni and Schmeidler 1991; Karni and Vierø 2013, 2015, 2017; Dominiak and Tserenjigmid 2018; see Karni 2017 for a good review of this approach). Some approaches go further, and altogether dispense with state space (Gilboa and Schmeidler, 2004; Ahn, 2008; Karni, 2006, 2007, 2011, 2013), outcome space (Skiadas, 1997a,b), or both (Blume et al., 2021).

This paper differs from the aforementioned literature in several ways. First, our main goal is to provide a *normative* analysis of "ideal rationality", rather than a *descriptive* account of actual human behaviour. Second and relatedly, the paper focuses on expected utility representations, rather than more general models of ambiguity. Third and most important, the paper it does not develop a single model, but rather, a *modelling framework*, which can be applied to construct and analyse models in a range of environments.

This modelling framework is formulated using tools from category theory. Category theory has already been applied in several parts of theoretical economics. These include: social choice theory (Keiding, 1981; Kijima and Takahara, 1987; Eklund et al., 2010, 2013; Abramsky, 2015), coalitional games (Machover and Terrington, 2014), normal form games (Vassilakis, 1992; Tohmé and Viglizzo, 2023), extensive-form games (Lescanne and Perrinel, 2012; Abramsky and Winschel, 2017; Streufert, 2018, 2020), combinatorial games (i.e. two-player, extensive-form games of perfect information) and the closely related topic of game semantics in formal logic (Joyal, 1977; Hirschowitz et al., 2007; Cockett and Seely, 2007; Honsell and Lenisa, 2011; Honsell et al., 2012a,b, 2014; Lenisa, 2015; Eberhart and Hirschowitz, 2018), and the construction of universal type spaces in games of incomplete information (Moss and Viglizzo, 2004, 2006; Pintér, 2010; Heinsalu, 2014; Fukuda, 2021; Guarino, 2022; Galeazzi and Marti, 2023). The previous literature most closely related to the present paper involves applications of category theory to decision theory (Rozen and Zhitomirski, 2006; Bosi and Herden, 2008; Adachi, 2014; De Oliveira, 2016; Tohmé et al., 2017). Bosi and Herden (2008) consider the construction of (ordinal) utility representations for preferences *without* uncertainty. Rozen and Zhitomirski (2006) and Adachi (2014) consider decisions with uncertainty. De Oliveira (2016) provides an elegant categorical proof of Blackwell's Theorem. Tohmé et al. (2017) have used sheaf theory to model the problem of building consistent preferences over a large menu from preferences over smaller menus. But none of these papers develops a model similar to the one in the present paper.

Organization. Sections 2 and 5 review basic category theory. Sections 3 and 4 introduce *decision environments* and *ex ante preference structures*, which together provide the modelling framework that accommodates all the issues raised earlier in this section. Section 6 defines *global SEU representations*, Section 7 contains an Anscombe-Aumann-type existence theorem for such representations for concrete categories, and Section 8 extends this existence theorem to abstract categories. Appendix A contains the proofs of all results in the text, Appendix B has the proofs of some supplementary results, and Appendix C gives some examples of abstract categories.

Notation. Sets and subsets are denoted by upper-case calligraphic letters $(\mathcal{A}, \mathcal{B}, \ldots)$, and their elements are normally denoted by the corresponding italic letters in either lower- or upper-case $(a, b, \ldots, A, B, \ldots)$. Objects in categories are also denoted by upper-case calligraphic letters. "Standard" categories are denoted in sans serif font (e.g. Set, Meas, Top, etc.), while generic categories, subcategories, or collections of objects in a category are denoted by bold calligraphic font ($\boldsymbol{\mathcal{C}}$, $\boldsymbol{\mathcal{S}}$, $\boldsymbol{\mathcal{X}}$). Morphisms, functions and measures are usually denoted by Greek letters $(\alpha, \beta, , \ldots)$. Vectors are indicated in bold face $(\mathbf{p}, \mathbf{q}, \ldots)$ and their components by the corresponding italics $(p_1, p_2, \ldots)$.

2 Categories[2]

A *category* is a mathematical structure $\boldsymbol{\mathcal{C}}$ with three parts:

- A collection $\boldsymbol{\mathcal{C}}^\circ$ of entities, called the *objects* of $\boldsymbol{\mathcal{C}}$.
- For any pair of objects $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, a collection $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ of entities, called *morphisms* from $\mathcal{A}$ to $\mathcal{B}$.
- For any $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, a *composition* operation $\circ$, such that, for any morphisms $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, we have $\psi \circ \phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$.

[2]Readers familiar with basic category theory can skip ahead to Section 3.

The composition operation has two key algebraic properties:

- *Associativity.* For all objects $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D} \in \boldsymbol{\mathcal{C}}^\circ$ and morphisms $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$, $\beta \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, and $\gamma \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$, we have $\gamma \circ (\beta \circ \alpha) = (\gamma \circ \beta) \circ \alpha$.
- *Identity.* For every object $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$, there is an *identity* morphism $I_\mathcal{A} \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{A})$ such that, for any object $\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, we have $I_\mathcal{A} \circ \phi = \phi$ for all $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{A})$, while $\phi \circ I_\mathcal{A} = \phi$ for all $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$.

$\boldsymbol{\mathcal{C}}$ is called a *concrete category* if the objects in $\boldsymbol{\mathcal{C}}^\circ$ are sets (typically with some additional "structure"), the morphisms in $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ are functions from the set $\mathcal{A}$ to the set $\mathcal{B}$ (which "preserve" this structure), and the composition operation $\circ$ is just ordinary function composition. This paper will work mainly with six categories. In the category Set, objects are *sets*, and morphisms are functions between them. In the category Meas, objects are *measurable spaces* (i.e. sets equipped with sigma algebras), and morphisms are measurable functions between them. In the category Top, objects are *topological spaces*, and morphisms are continuous functions. In the category Metr, objects are *metric spaces*, and morphisms are *short maps* —that is, functions such that the distance between two points is never less than the distance between their images. In the category Diff, objects are *differentiable manifolds*, and morphisms are *smooth functions* —that is, functions that are everywhere infinitely differentiable. In the category UPOVS, objects are *unitary partially ordered vector spaces*, and morphisms are *uniferent order-preserving linear functions* (these will be defined later). But not all categories are concrete. A key feature of the theory in this paper is that it does not require a concrete category. We use the term *abstract category* to refer to a category which may or may not be concrete; for some examples of *non*-concrete categories, see Appendix C. Good introductions to category theory include Spivak (2014), Simmons (2011), Riehl (2017), Leinster (2014), Awodey (2010), Adámek et al. (2009), and Mac Lane (1998) (in roughly increasing order of difficulty).

Isomorphisms, automorphisms, and groups. Recall that a *group* is a structure $(\mathcal{G}, *, e)$ where $\mathcal{G}$ is a set, $*$ is a binary operation on $\mathcal{G}$, and e is an element of $\mathcal{G}$, such that: (i) $f * (g * h) = (f * g) * h$ for all $f, g, h \in \mathcal{G}$; (ii) $e * g = g * e = g$ for all $g \in \mathcal{G}$; and (iii) for every $g \in \mathcal{G}$, there is an element $h \in \mathcal{G}$ (the *inverse* of g) such that $g * h = h * g = e$. An example is the group of permutations of a set, where $*$ is function composition.

Now let $\boldsymbol{\mathcal{C}}$ be a category, and let $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$. A morphism $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ is an *isomorphism* if there is a morphism $\psi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{A})$ (the *inverse* of ϕ) such that $\psi \circ \phi = I_\mathcal{A}$ and $\phi \circ \psi = I_\mathcal{B}$. The composition of isomorphisms is also an isomorphism. An isomorphism from $\mathcal{A}$ to itself is called an *automorphism*. The set of all automorphisms of $\mathcal{A}$ forms a group under composition.

This observation leads to a natural way to represent any group as an abstract, single-object category. Given a group $(\mathcal{G}, *, e)$, we create a category which has only a single object: $\boldsymbol{\mathcal{C}}^\circ = \{\mathcal{C}\}$. All elements of $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{C})$ are automorphisms, and there is a bijection $\phi : \mathcal{G} \longrightarrow \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{C})$ such that, for any $g, h \in \mathcal{G}$, $\phi(g * h) = \phi(g) \circ \phi(h)$ (in other words, ϕ is a *group isomorphism*). In this way, group theory can be seen as a branch of category theory.

Figure 1: *Left.* The outcome place category $\boldsymbol{\mathcal{X}}$ of Example 3.1(b), in the special case of Example 1.1. *Right.* The state place category $\boldsymbol{\mathcal{S}}$ of Example 3.1(c), in the special case of Example 1.2.

3 Decision Environments

Let $\boldsymbol{\mathcal{C}}$ be a category. A *decision environment* on $\boldsymbol{\mathcal{C}}$ is an ordered pair $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$, where $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{\mathcal{X}}$ are *subcategories* of $\boldsymbol{\mathcal{C}}$. In other words, $\boldsymbol{\mathcal{S}}^\circ \subseteq \boldsymbol{\mathcal{C}}^\circ$ and $\boldsymbol{\mathcal{X}}^\circ \subseteq \boldsymbol{\mathcal{C}}^\circ$. Furthermore, for any $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, we have $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2) \subseteq \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_1, \mathcal{S}_2)$; likewise, for any $\mathcal{X}_1, \mathcal{X}_2 \in \boldsymbol{\mathcal{X}}^\circ$, we have $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2) \subseteq \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}_1, \mathcal{X}_2)$. We shall interpret the objects of $\boldsymbol{\mathcal{S}}^\circ$ as "abstract state spaces", and interpret objects of $\boldsymbol{\mathcal{X}}^\circ$ as "abstract outcome spaces". However, if $\boldsymbol{\mathcal{C}}$ is an abstract category, then these might not actually be spaces of any kind. For this reason, we shall refer to the objects of $\boldsymbol{\mathcal{S}}^\circ$ as *state places* and the objects of $\boldsymbol{\mathcal{X}}^\circ$ as *outcome places*.

For any state place $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}^\circ$ and outcome place $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}^\circ$, the morphisms in $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$ represent "abstract acts" —these are devices that somehow transform the abstract "states" in $\mathcal{S}$ into abstract "outcomes" in $\mathcal{X}$. For simplicity, we shall call them *acts*. If $\boldsymbol{\mathcal{C}}$ was a concrete category, then $\mathcal{S}$ and $\mathcal{X}$ would be sets, and the acts in $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$ would be functions. But we shall not assume this.

Heuristically, each state place in $\boldsymbol{\mathcal{S}}^\circ$ represents a source of uncertainty. Suppose that the agent has "beliefs" about these sources of uncertainty. For any state places $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, we interpret each element of $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$ as a $\boldsymbol{\mathcal{C}}$-morphism from $\mathcal{S}_1$ to $\mathcal{S}_2$ that is somehow "compatible" with her beliefs about $\mathcal{S}_1$ and $\mathcal{S}_2$. (We shall make this precise shortly.) For example, if the agent's beliefs took the form of probability measures, then $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$ could be the set of *measure-preserving* functions from $\mathcal{S}_1$ into $\mathcal{S}_2$. However, we shall not (yet) commit to any formal model of the agent's beliefs (e.g. as probabilities), so we shall not (yet) impose any restrictions on the sort of morphisms that can appear in $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$.

Meanwhile, each outcome place in $\boldsymbol{\mathcal{X}}^\circ$ represents an abstract "menu" of possible outcomes. Suppose that the agent has "tastes" over these menus. For any $\mathcal{X}_1, \mathcal{X}_2 \in \boldsymbol{\mathcal{X}}^\circ$, we interpret each element of $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2)$ as a $\boldsymbol{\mathcal{C}}$-morphism from $\mathcal{X}_1$ to $\mathcal{X}_2$ that is somehow "compatible" with her tastes over $\mathcal{X}_1$ and $\mathcal{X}_2$. (We shall make this precise shortly.) For example, if her tastes took the form of preference orders on $\mathcal{X}_1$ and $\mathcal{X}_2$, then $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2)$ could be the set of *order-preserving* $\boldsymbol{\mathcal{C}}$-morphisms from $\mathcal{X}_1$ into $\mathcal{X}_2$. However, we shall not (yet) commit to any formal model of the agent's tastes (e.g. in terms of preference orders or utility functions), so we shall not (yet) impose any restrictions on the sort of morphisms that can appear in $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2)$.

Example 3.1. (a) (*Classic environment*) Let $\mathcal{S}$ and $\mathcal{X}$ be two objects in $\boldsymbol{\mathcal{C}}^\circ$. Let $\boldsymbol{\mathcal{S}}^\circ := \{\mathcal{S}\}$ and $\boldsymbol{\mathcal{X}}^\circ := \{\mathcal{X}\}$. Let $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}, \mathcal{S}) := \{I_\mathcal{S}\}$ and $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{X}) := \{I_\mathcal{X}\}$, where $I_\mathcal{S}$ and $I_\mathcal{X}$ are the identity morphisms on $\mathcal{S}$ and $\mathcal{X}$. Then $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ is the decision environment of an agent who faces a single source of uncertainty (namely, $\mathcal{S}$), and a single set of outcomes (namely, $\mathcal{X}$). When $\boldsymbol{\mathcal{C}} = \mathsf{Set}$, this is the setting of Savage (1954) and Anscombe and Aumann (1963).

(b) (*Variable menus*) Let $\boldsymbol{\mathcal{C}} := \mathsf{Set}$. Let $\mathcal{S} \in \boldsymbol{\mathcal{C}}^\circ$, let $\boldsymbol{\mathcal{S}}^\circ := \{\mathcal{S}\}$, and let $\vec{\boldsymbol{\mathcal{S}}}(\mathcal{S}, \mathcal{S}) := \{I_\mathcal{S}\}$, as in Example (a). Fix some other $\mathcal{W} \in \boldsymbol{\mathcal{C}}^\circ$; it will play the role of a "universal" space of outcomes. Let $\boldsymbol{\mathcal{X}}^\circ$ be a collection of subsets of $\mathcal{W}$. For any $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, if $\mathcal{X} \subseteq \mathcal{Y}$, then let $\vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) := \{\iota^\mathcal{X}_\mathcal{Y}\}$, where $\iota^\mathcal{X}_\mathcal{Y} : \mathcal{X} \hookrightarrow \mathcal{Y}$ is the inclusion map.[3] Meanwhile, if $\mathcal{X} \not\subseteq \mathcal{Y}$, then let $\vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) := \emptyset$. This decision environment describes an agent who faces a single source of uncertainty (namely, $\mathcal{S}$), but confronts a variety of outcome menus (the objects in $\boldsymbol{\mathcal{X}}^\circ$.) We do *not* assume that $\mathcal{W}$ itself is an object in $\boldsymbol{\mathcal{X}}^\circ$. For instance, in Example 1.1, $\mathcal{W} = \{\text{Jan}, \text{Feb}, \cdots, \text{Dec}\}$, but $\boldsymbol{\mathcal{X}}^\circ$ only contains the menus $\mathcal{X}_1 = \{\text{Jan}, \ldots, \text{Jun}\}$, $\mathcal{X}_2 = \{\text{May}, \ldots, \text{Oct}\}$ and $\mathcal{X}_3 = \{\text{Sep}, \ldots, \text{Feb}\}$, along with the three nonempty intersections between these menus, namely $\mathcal{X}_{12} = \{\text{May}, \text{Jun}\}$, $\mathcal{X}_{23} = \{\text{Sep}, \text{Oct}\}$ and $\mathcal{X}_{31} = \{\text{Jan}, \text{Feb}\}$. The morphisms of $\boldsymbol{\mathcal{X}}$ are the inclusion maps, as shown in Figure 1 (left).

(c) (*Variable information or awareness*) Let $\boldsymbol{\mathcal{C}} := \mathsf{Set}$. Let $\mathcal{X}$ be a set, let $\boldsymbol{\mathcal{X}}^\circ := \{\mathcal{X}\}$, and let $\vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{X}) := \{I_\mathcal{X}\}$, as in Example (a). Let $\mathcal{S}$ be another set, and let $\mathcal{E}$ be a collection of equivalence relations on $\mathcal{S}$. For every equivalence relation $\sim$ in $\mathcal{E}$, let $(\mathcal{S}/_\sim)$ be the set of equivalence classes. For any $s \in \mathcal{S}$ and $\sim$ in $\mathcal{E}$, let $[s]_\sim \in (\mathcal{S}/_\sim)$ denote the corresponding equivalence class. For any two equivalence relations $\sim_1$ and $\sim_2$, we say that $\sim_1$ *refines* $\sim_2$ if, for all $s, t \in \mathcal{S}$, we have $(s \sim_1 t) \implies (s \sim_2 t)$. This implies that every $\sim_2$-equivalence class is a union of $\sim_1$ equivalence classes. So there is a unique surjective *quotient map* $\pi : (\mathcal{S}/_{\sim_1}) \longrightarrow (\mathcal{S}/_{\sim_2})$ defined by setting $\pi([s]_{\sim_1}) := [s]_{\sim_2}$ for all $s \in \mathcal{S}$.

This yields a category $\boldsymbol{\mathcal{S}}$, where $\boldsymbol{\mathcal{S}}^\circ := \{(\mathcal{S}/_\sim);\ \sim \in \mathcal{E}\}$, and where for any $\sim_1$ and $\sim_2$ in $\mathcal{E}$, if $\sim_1$ refines $\sim_2$, then $\vec{\boldsymbol{\mathcal{S}}}\left((\mathcal{S}/_{\sim_1}), (\mathcal{S}/_{\sim_2})\right) = \{\pi\}$, where π is the unique quotient map, whereas $\vec{\boldsymbol{\mathcal{S}}}\left((\mathcal{S}/_{\sim_1}), (\mathcal{S}/_{\sim_2})\right) = \emptyset$ otherwise.

There are two ways we can interpret this decision environment. According to the first interpretation, it describes an agent who faces a single source of uncertainty (namely, $\mathcal{S}$), and a single set of outcomes (namely, $\mathcal{X}$), but has access to one of a variety of *information sources* about $\mathcal{S}$. Each equivalence relation $\sim$ in $\mathcal{E}$ represents one possible source of information: if $s \sim t$, then the agent simply cannot distinguish s from t given the information represented by $\sim$. Thus, she cannot perform an act that depends upon the distinction between s and t. Hence, the only feasible acts are functions from $(\mathcal{S}/_\sim)$ to $\mathcal{X}$, so the agent only forms preferences over such acts.

According to the second interpretation, the elements of $\mathcal{E}$ do not represent *informational* constraints, but rather, *cognitive* constraints. Each equivalence relation $\sim$ represent some level of *awareness* the agent might have. According to this interpretation, if $s \sim t$, then the agent is not even *aware* of the distinction between s and t; thus, she cannot even *conceive* of an act that depends upon this distinction. Hence, only functions from $(\mathcal{S}/_\sim)$ to $\mathcal{X}$ are conceivable, so the agent is only *able* to form preferences over such acts.

[3]That is, for all $x \in \mathcal{X}$, $\iota^\mathcal{X}_\mathcal{Y}(x) = x$, seen as an element of $\mathcal{Y}$.

Note that $\mathcal{S}$ itself is not necessarily an object in $\boldsymbol{\mathcal{S}}$, unless the "discrete" equivalence relation (where $s \not\sim t$ whenever $s \neq t$) is an element of $\mathcal{E}$. Recall Example 1.2. Let $\mathcal{S} := \mathcal{S}_1 \times \mathcal{S}_2 \times \mathcal{S}_3$. Define equivalence relations $\sim_1, \sim_2, \sim_3 \sim_{12}, \sim_{23}, \sim_{13}$ as follows: for any $\mathbf{s} = (s_1, s_2, s_3)$ and $\mathbf{t} = (t_1, t_2, t_3)$ in $\mathcal{S}$ and $j \in \{1, 2, 3\}$, stipulate that $\mathbf{s} \sim_j \mathbf{t}$ if and only if $s_j = t_j$. Likewise, for any $j, k \in \{1, 2, 3\}$, stipulate that $\mathbf{s} \sim_{jk} \mathbf{t}$ if and only if $s_j = t_j$ and $s_k = t_k$. In the notation of Example 1.2, the quotient space $(\mathcal{S}/_{\sim_j})$ can be identified with $\mathcal{S}_j$ and the quotient space $(\mathcal{S}/_{\sim_{jk}})$ can be identified with $\mathcal{S}_{jk}$, for all $j, k \in \{1, 2, 3\}$, and the quotient maps correspond to the coordinate projection maps. In this case, $\boldsymbol{\mathcal{S}}^\circ = \{\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_{12}, \mathcal{S}_{23}, \mathcal{S}_{13}\}$ and the only morphisms are whatever coordinate projections exist between these spaces, as shown in Figure 1 (right). $\diamondsuit$

It is also possible to combine the subcategory $\boldsymbol{\mathcal{X}}$ from Example (b) with the subcategory $\boldsymbol{\mathcal{S}}$ from Example (c), to obtain a decision environment with both varying levels of awareness and varying menus of outcomes; this is similar to the model of Dietrich (2018). As is suggested by these examples, for each $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}$ and $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}$, the pair $(\mathcal{S}, \mathcal{X})$ can be interpreted as what Savage called a "small world": a simplified model of reality that contains only the information relevant to a particular decision problem. An agent might not have the time, information, or cognitive capacity necessary to combine all these small worlds into a single "grand world". As in Example 1.1, she might be able to formulate preferences over *small* menus of outcomes, but not be able to formulate them over the set of *all* possible outcomes. As in Example 1.2, she might have well-defined probabilistic beliefs about particular sources of uncertainty, but be unable to formulate a probabilistic belief system that simultaneously encompasses *every* source of uncertainty in the universe. Nevertheless, the decision environment $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ enables her to represent every decision problem that she might encounter with a small world.

4 Ex ante preference structures

Let $\boldsymbol{\mathcal{C}}$ be a category, and let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be a decision environment on $\boldsymbol{\mathcal{C}}$. For every $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ be a weak order[4] on $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$. The collection $\trianglerighteq^{\text{xa}} = \{\succcurlyeq^{\mathcal{S}}_{\mathcal{X}};\ \mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ\}$ is an *ex ante preference structure* on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ if it is "compatible" with the morphisms of $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{\mathcal{X}}$ in the following two senses.

(BP) For any state places $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, any morphism $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, any outcome place $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, and any acts $\alpha, \beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_2, \mathcal{X})$, we have $\alpha \succcurlyeq^{\mathcal{S}_2}_{\mathcal{X}} \beta \iff \alpha \circ \phi \succcurlyeq^{\mathcal{S}_1}_{\mathcal{X}} \beta \circ \phi$.

(TP) For any outcome places $\mathcal{X}_1, \mathcal{X}_2 \in \boldsymbol{\mathcal{X}}^\circ$, any morphism $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2)$, any state place $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, and any acts $\alpha, \beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X}_1)$, we have $\alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}_1} \beta \iff \phi \circ \alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}_2} \phi \circ \beta$.

Compatibility condition (BP) formalizes the earlier informal assertion that the morphisms of the subcategory $\boldsymbol{\mathcal{S}}$ are "belief-preserving". Likewise, condition (TP) formalizes the earlier assertion that the morphisms of the subcategory $\boldsymbol{\mathcal{X}}$ are "taste-preserving".

[4]That is: a complete, transitive, reflexive binary relation.

Example 4.1. (a) (*Classic environment*) Let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be as in Example 3.1(a), with $\boldsymbol{\mathcal{S}}^\circ = \{\mathcal{S}\}$ and $\boldsymbol{\mathcal{X}}^\circ = \{\mathcal{X}\}$, etc. Let $\succcurlyeq$ be a preference order on $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{X})$; then $\{\succcurlyeq\}$ is (trivially) an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$.

(b) (*Variable menus*) Let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be as in Example 3.1(b), where $\boldsymbol{\mathcal{S}}^\circ = \{\mathcal{S}\}$ and $\boldsymbol{\mathcal{X}}^\circ$ is a collection of subsets of some "universal" outcome space $\mathcal{W}$. Let $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, with $\mathcal{X} \subseteq \mathcal{Y}$, and let $\iota : \mathcal{X} \hookrightarrow \mathcal{Y}$ be the inclusion map. For any $\mathcal{X}$-valued acts $\alpha, \beta \in \overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{X})$, the compositions $\iota \circ \alpha$ and $\iota \circ \beta$ are elements of $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{Y})$; heuristically, these are *the same acts* as α and β, but "reframed" as $\mathcal{Y}$-valued acts. Condition (TP) requires that $\alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}} \beta$ if and only if $\iota \circ \alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{Y}} \iota \circ \beta$. This is a version *Independence of Irrelevant Alternatives*: it says the comparison between α and β should be determined only by the outcomes that are actually in the *range* of α and β (i.e. $\mathcal{X}$), and should not be affected by introducing "irrelevant" alternatives (i.e. the other elements of $\mathcal{Y}$).

For example, let $\succcurlyeq^*$ be a preference order on the set $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{W})$. For any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, every function from $\mathcal{S}$ into $\mathcal{X}$ can be seen as a function from $\mathcal{S}$ into $\mathcal{W}$ (because $\mathcal{X} \subseteq \mathcal{W}$); hence $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{X})$ can be treated as a subset of $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{W})$. So, for all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ be the restriction of $\succcurlyeq^*$ to a preference order on $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{X})$. Then $\{\succcurlyeq^{\mathcal{S}}_{\mathcal{X}};\, \mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ\}$ trivially satisfies (TP), so it is an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$.

However, not all ex ante preference structures in this decision environment arise in this way. For instance, consider the formalization of Example 1.1 given at the end of Example 3.1(b). Let μ be some probability distribution $\mathcal{S}$, and suppose the employee has SEU preferences over acts mapping $\mathcal{S}$ into $\mathcal{X}_1$, $\mathcal{X}_2$, $\mathcal{X}_3$, $\mathcal{X}_{12}$, $\mathcal{X}_{23}$, or $\mathcal{X}_{31}$ based on the utility functions shown by the table below:

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
$\mathcal{X}_1$	1	2	3	4	5	6						
$\mathcal{X}_2$					1	2	3	4	5	6		
$\mathcal{X}_3$	5	6							1	2	3	4
$\mathcal{X}_{12}$					1	2						
$\mathcal{X}_{23}$									1	2		
$\mathcal{X}_{31}$	1	2										

It is easily verified that these preferences are consistent with (TP). For example, let $\alpha, \beta : \mathcal{S} \longrightarrow \mathcal{X}_{12}$ be acts. Let $\iota^1_{12} : \mathcal{X}_{12} \hookrightarrow \mathcal{X}_1$ and $\iota^2_{12} : \mathcal{X}_{12} \hookrightarrow \mathcal{X}_2$ be the inclusion maps, as shown in Figure 1. Clearly, the μ-expected utility of α and the μ-expected utility of $\iota^2_{12} \circ \alpha$ are both equal to μ-expected utility of $\iota^1_{12} \circ \alpha$, minus 4. Likewise for β. Thus, $\alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}_{12}} \beta$ if and only if $\iota^1_{12} \circ \alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}_1} \iota^1_{12} \circ \beta$ if and only if $\iota^2_{12} \circ \alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}_2} \iota^2_{12} \circ \beta$, in accord with (TP).

(c) (*Variable information and awareness*) Let $\boldsymbol{\mathcal{C}} = \mathsf{Set}$, and let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be as in Example 3.1(c), where $\boldsymbol{\mathcal{X}}^\circ = \{\mathcal{X}\}$ for some set $\mathcal{X}$, while $\boldsymbol{\mathcal{S}}^\circ$ is the set of quotient spaces obtained from a collection $\mathcal{E}$ of equivalence relations of some set $\mathcal{S}$.

For any $\sim$ in $\mathcal{E}$, every element of $\overrightarrow{\mathsf{Set}}((\mathcal{S}/_\sim), \mathcal{X})$ can be be represented by a function from $\mathcal{S}$ into $\mathcal{X}$ that is constant on each $\sim$-equivalence class. Let $\sim_1$ and $\sim_2$ be in $\mathcal{E}$. Suppose $\sim_1$ refines $\sim_2$, so that $\overrightarrow{\boldsymbol{\mathcal{S}}}((\mathcal{S}/_{\sim_1}), (\mathcal{S}/_{\sim_2})) = \{\pi\}$, where π is the quotient map. If

Figure 2: *Left.* The outcome place category $\boldsymbol{\mathcal{X}}$ of Example 4.2(a). (Bidirectional arrows indicate isomorphisms.) *Right.* The state place category $\boldsymbol{\mathcal{S}}$ of Example 4.2(b).

$\alpha, \beta \in \overrightarrow{\mathsf{Set}}((\mathcal{S}/_{\sim_2}), \mathcal{X})$, then α and $\alpha \circ \pi$ both correspond to the *same* function from $\mathcal{S}$ to $\mathcal{X}$. Likewise, β and $\beta \circ \pi$ both correspond to the same function from $\mathcal{S}$ to $\mathcal{X}$. Condition (BP) requires that $\alpha \succcurlyeq_{\mathcal{X}}^{(\mathcal{S}/\sim_2)} \beta$ if and only if $\alpha \circ \pi \succcurlyeq_{\mathcal{X}}^{(\mathcal{S}/\sim_2)} \beta \circ \pi$, To understand this, note that α and β only depend on the coarser information represented by $\sim_2$, and do not use the finer information represented by $\sim_1$. So condition (BP) is a sort of *Independence of Irrelevant Information*: preferences between two acts should be determined only by the information (or awareness) needed to describe these acts.

For example, let $\succcurlyeq^*$ be a preference order on the set of *all* functions from $\mathcal{S}$ to $\mathcal{X}$. For any $\sim$ in $\mathcal{E}$, let $\succcurlyeq_{\mathcal{X}}^{(\mathcal{S}/\sim)}$ be the restriction of $\succcurlyeq^*$ to $\overrightarrow{\mathsf{Set}}((\mathcal{S}/_{\sim}), \mathcal{X})$ (where we again interpret each element of $\overrightarrow{\mathsf{Set}}((\mathcal{S}/_{\sim}), \mathcal{X})$ as a function from $\mathcal{S}$ into $\mathcal{X}$ that is constant on each $\sim$-equivalence class). Then $\{\succcurlyeq_{\mathcal{X}}^{(\mathcal{S}/\sim)};\ \sim \text{ in } \mathcal{E}\}$ is an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$.

But not all ex ante preference structures in this decision environment arise in this way. Consider the formalization of Example 1.2 given at the end of Example 3.1(c). Let $\alpha, \beta : \mathcal{S}_1 \longrightarrow \mathcal{X}$ be acts. Let $\pi_{12} : \mathcal{S}_{12} \longrightarrow \mathcal{S}_1$ and $\pi_{13} : \mathcal{S}_{13} \longrightarrow \mathcal{S}_1$ be the coordinate projection maps shown in Figure 1. Although they are defined on larger state spaces, the acts $\alpha \circ \pi_{12}$, $\beta \circ \pi_{12}$, $\alpha \circ \pi_{13}$ and $\beta \circ \pi_{13}$ only depend on information about $\mathcal{S}_1$ —i.e. whether the patient has condition a or b. For any utility function u on $\mathcal{X}$, it is easily verified that the μ_{12}-expected utility of $\alpha \circ \pi_{12}$ and the μ_{13}-expected utility of $\alpha \circ \pi_{13}$ both equal $0.5\, u \circ \alpha(a) + 0.5\, u \circ \alpha(b)$, which is the μ_1-expected utility of α. Likewise for β. Thus, $\alpha \succcurlyeq_{\mathcal{X}}^{\mathcal{S}_1} \beta$ if and only if $\alpha \circ \pi_{12} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}_{12}} \beta \circ \pi_{12}$ if and only if $\alpha \circ \pi_{13} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}_{13}} \beta \circ \pi_{13}$, in accord with (BP). $\diamondsuit$

Models similar to Example 4.1(c) have appeared in the literature on "framing effects", "unawareness" and "salience effects" (Heifetz et al., 2006; Ahn and Ergin, 2010; Hayashi, 2012; Schipper, 2013; Dietrich, 2018). However, these papers aim for *descriptive* models of "bounded rationality" in real humans, so they specifically do *not* assume (BP). In contrast, I assume (BP) because I am aiming for in a *normative* model of "ideal rationality". Cohen and Jaffray (1980) have a "variable-information" model of ideal rationality that *does* use a version of (BP) (their Axiom 3). But their other axioms preclude an SEU representation, and instead yield a refinement of *Arrow-Hurwicz* preferences, where the preference between two acts is entirely determined by their minimum and maximum values.

In Example 4.1, $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{\mathcal{X}}$ contained the minimum set of morphisms needed to capture the logical relationships between the different objects. But morphisms can also encode *analogies*, *symmetries* or *invariance properties*, as the next examples show.

Example 4.2. (a) (*Insufficient reason*) Let $\mathcal{C} = \mathsf{Set}$, let $\mathcal{S}$ be a finite set, and let $\mathcal{S}^\circ = \{\mathcal{S}\}$, as in Examples 3.1(a) and 4.1(a). But now, let $\vec{\mathcal{S}}(\mathcal{S},\mathcal{S})$ be the group of *all* permutations of $\mathcal{S}$. Then axiom (BP) encodes a version of Laplace's *Principle of Insufficient Reason*: for any acts $\alpha, \beta \in \overrightarrow{\mathsf{Set}}(\mathcal{S},\mathcal{X})$, if $\alpha \succcurlyeq \beta$, then $\alpha \circ \pi \succcurlyeq \beta \circ \pi$ for all permutations π of $\mathcal{S}$. Heuristically, this describes an agent who regards all elements of $\mathcal{S}$ as indistinguishable, and hence equally likely, ex ante.

(b) (*Analogous preferences*) In Example 1.1, the reader will have noticed that the employee's preferences over the outcome spaces $\mathcal{X}_1$, $\mathcal{X}_2$, and $\mathcal{X}_3$ are structurally analogous: in all cases, she prefers later months to earlier ones. The same is true when comparing their intersections $\mathcal{X}_{12}$, $\mathcal{X}_{23}$, and $\mathcal{X}_{31}$. This is especially clear in the table of utilities in Example 4.1(b). But the decision environment in Example 3.1(b) fails to recognize these analogies. A better model would contain not only the inclusion maps shown in Figure 1, but also a bijection $\mathcal{X}_1 \longrightarrow \mathcal{X}_2$ shifting time forward by four months, and similar bijections $\mathcal{X}_2 \longrightarrow \mathcal{X}_3$, $\mathcal{X}_3 \longrightarrow \mathcal{X}_1$, $\mathcal{X}_{12} \longrightarrow \mathcal{X}_{23}$, $\mathcal{X}_{23} \longrightarrow \mathcal{X}_{31}$, and $\mathcal{X}_{31} \longrightarrow \mathcal{X}_{12}$, along with their inverse maps. The resulting category is shown in Figure 2. The ex ante preference structure in Example 4.1(b) also satisfies (TP) with respect to this larger network of morphisms.

(c)(*Analogous Beliefs*) In Example 1.2, the reader will also have noticed that the doctor's beliefs about the state spaces $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_3$ are analogous: in each case, the doctor assigns equal probability to the two possible diseases. The same is true when comparing their products $\mathcal{S}_{12}$, $\mathcal{S}_{23}$, and $\mathcal{S}_{31}$. The decision environment in Example 3.1(c) fails to recognize these analogies. A better model would include not only the projection maps shown in Figure 1, but also probability-preserving bijections $\phi_2^1 : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, $\phi_3^2 : \mathcal{S}_2 \longrightarrow \mathcal{S}_3$ and $\phi_1^3 : \mathcal{S}_3 \longrightarrow \mathcal{S}_1$ given by

$$\begin{array}{rclcrclcrcl} \phi_2^1(a) & = & c, & \quad & \phi_3^2(c) & = & e, & & \phi_1^3(e) & = & a, \\ \phi_2^1(b) & = & d, & & \phi_3^2(d) & = & f & \text{and} & \phi_1^3(f) & = & b, \end{array}$$

and their inverses, along with the bijections $(\phi_2^1 \times \phi_3^2) : \mathcal{S}_{12} \longrightarrow \mathcal{S}_{23}$, $(\phi_3^2 \times \phi_1^3) : \mathcal{S}_{23} \longrightarrow \mathcal{S}_{31}$ and $(\phi_1^3 \times \phi_2^1) : \mathcal{S}_{31} \longrightarrow \mathcal{S}_{12}$ and their inverses. The resulting category is shown in Figure 2. The ex ante preference structure in Example 4.1(c) also satisfies (BP) with respect to this larger network of morphisms. ◇

5 Functors and natural transformations[5]

Before giving the formal definition of global SEU representations in Section 6, we must first review a few concepts from category theory. Let $\mathcal{C}$ and $\mathcal{D}$ be categories. Let $\vec{\mathcal{C}}$ be the set of all morphisms in $\mathcal{C}$ and let $\vec{\mathcal{D}}$ be the set of all morphisms in $\mathcal{D}$. A (covariant) *functor* from $\mathcal{C}$ to $\mathcal{D}$ consists of a function $F : \mathcal{C}^\circ \longrightarrow \mathcal{D}^\circ$ and a function $F : \vec{\mathcal{C}} \longrightarrow \vec{\mathcal{D}}$, such that for all $\mathcal{A}, \mathcal{B} \in \mathcal{C}^\circ$, the function F maps $\vec{\mathcal{C}}(\mathcal{A},\mathcal{B})$ to $\vec{\mathcal{D}}(F(\mathcal{A}), F(\mathcal{B}))$, and F preserves

[5] Readers familiar with basic category theory can skip ahead to Section 6.

morphism composition. In other words: for all $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^{\circ}$, and all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, $F(\beta \circ \alpha) = F(\beta) \circ F(\alpha)$. We indicate this by writing "$F : \boldsymbol{\mathcal{C}} \Longmapsto \boldsymbol{\mathcal{D}}$".

5A Forgetful functors. One natural class of functors comes from "forgetting" some of the mathematical structure in the objects, reducing them to simpler objects. For example, from any of the categories Meas, Top, Metr, Diff, or UPOVS introduced in Section 2, there is a functor into the category Set, which reduces every object to its underlying set of points, and reduces every morphism to the underlying function. These are called *forgetful functors*. Similarly, there is a functor F : Top $\Longmapsto$ Meas that reduces any topological space $\mathcal{S}$ to the measurable space with the same set of points and the Borel sigma algebra, and reduces every continuous function between topological spaces to the same function, seen as a measurable function between their Borel sigma algebras. There is a functor F : Metr $\Longmapsto$ Top that reduces every metric space into a (metrizable) topological space, and reduces every short map to a continuous map. There is a functor F : Diff $\Longmapsto$ Top that reduces every differentiable manifold into a topological manifold, and reduces every smooth map to a continuous map.

5B Opposite categories and contravariant functors. Let $\boldsymbol{\mathcal{C}}$ be a category. The *opposite* category $\boldsymbol{\mathcal{C}}^{\rm op}$ is the category which has exactly the same objects as $\boldsymbol{\mathcal{C}}$, but where all morphism have their "direction" reversed, and where morphism composition happens in reverse order. For example, the commuting diagram on the left in the category $\boldsymbol{\mathcal{C}}$ becomes the commuting diagram on the right in the category $\boldsymbol{\mathcal{C}}^{\rm op}$

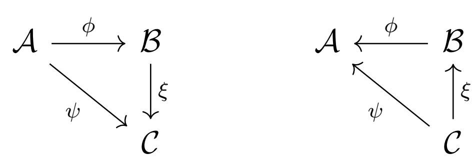

A *contravariant functor* from $\boldsymbol{\mathcal{C}}$ to another category $\boldsymbol{\mathcal{D}}$ is a (covariant) functor from $\boldsymbol{\mathcal{C}}^{\rm op}$ into $\boldsymbol{\mathcal{D}}$. In other words, it is a functor from $\boldsymbol{\mathcal{C}}$ into $\boldsymbol{\mathcal{D}}$ that "reverses the directions of arrows". We indicate this by writing "$F : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \boldsymbol{\mathcal{D}}$".

For example, let Vec be the category of vector spaces and linear functions. For any topological space $\mathcal{X}$, let $\mathfrak{C}_{\rm b}(\mathcal{X})$ be the vector space of all bounded, continuous, real-valued functions on $\mathcal{X}$. This is an object in Vec. For any two spaces $\mathcal{X}$ and $\mathcal{Y}$, and any continuous function $\phi : \mathcal{X} \longrightarrow \mathcal{Y}$, there is an induced function $\overleftarrow{\phi} : \mathfrak{C}_{\rm b}(\mathcal{Y}) \longrightarrow \mathfrak{C}_{\rm b}(\mathcal{X})$ defined by setting $\overleftarrow{\phi}(f) := f \circ \phi$ for any $f \in \mathfrak{C}_{\rm b}(\mathcal{Y})$. It is easily verified that $\overleftarrow{\phi}$ is a linear function —that is, a morphism in Vec. Furthermore, if $\mathcal{Z}$ is a third topological space, and $\psi : \mathcal{Y} \longrightarrow \mathcal{Z}$ is another continuous function, then $\overleftarrow{\psi \circ \phi} = \overleftarrow{\phi} \circ \overleftarrow{\psi}$. Thus, if we define $\mathfrak{C}_{\rm b}(\phi, \mathcal{X}) := \overleftarrow{\phi}$ for any continuous map ϕ, then we get a contravariant functor $\mathfrak{C}_{\rm b}$: Top$^{\rm op} \Longmapsto$ Vec. (This example is typical of a large class of examples. Contravariant functors frequently arise by constructing "spaces of functions" over objects.)

5C Hom functors. Let $\boldsymbol{\mathcal{C}}$ be a category, and let $\mathcal{X} \in \boldsymbol{\mathcal{C}}^\circ$. For any other object $\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, consider the set $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{B})$ of morphisms from $\mathcal{X}$ into $\mathcal{B}$. This is an object in Set. For any $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, there is a natural function $\vec{\phi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{B}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{C})$ defined by setting $\vec{\phi}(\alpha) := \phi \circ \alpha$ for all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{B})$. Furthermore, if $\mathcal{D} \in \boldsymbol{\mathcal{C}}^\circ$ is a third object and $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$ is another morphism, then $\overrightarrow{\psi \circ \phi} = \vec{\psi} \circ \vec{\phi}$. Thus, if we define $\overrightarrow{\boldsymbol{\mathcal{C}}}(\phi, \mathcal{X}) := \vec{\phi}$ for every $\boldsymbol{\mathcal{C}}$-morphism ϕ, then we get a covariant functor $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \bullet) : \boldsymbol{\mathcal{C}} \Longmapsto \mathsf{Set}$. This is called a *covariant hom functor*.

On the other hand, for any other object $\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, consider the set $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{X})$ of morphisms from $\mathcal{B}$ into $\mathcal{X}$. This is also an object in Set. For any $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, there is a natural function $\overleftarrow{\phi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{X})$ defined by setting $\overleftarrow{\phi}(\alpha) := \alpha \circ \phi$ for all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$. Furthermore, if $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$ is a third object and $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ is another morphism, then $\overleftarrow{\phi \circ \psi} = \overleftarrow{\psi} \circ \overleftarrow{\phi}$ Thus, if we define $\overrightarrow{\boldsymbol{\mathcal{C}}}(\phi, \mathcal{X}) := \overleftarrow{\phi}$ for every $\boldsymbol{\mathcal{C}}$-morphism ϕ, then we get a contravariant functor $\overrightarrow{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \mathsf{Set}$. This is called a *contravariant hom functor*. (Note the similarity to the definition of $\mathfrak{C}_{\mathrm{b}}(\bullet) : \mathsf{Top}^{\mathrm{op}} \Longmapsto \mathsf{Vec}$ in §5B.)

5D Natural transformations. Let $\boldsymbol{\mathcal{C}}$ and $\boldsymbol{\mathcal{D}}$ be two categories, and let $F : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \boldsymbol{\mathcal{D}}$ and $G : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \boldsymbol{\mathcal{D}}$ be two contravariant functors. A *natural transformation* from F to G is a collection of $\boldsymbol{\mathcal{D}}$-morphisms $\Phi = (\phi_{\mathcal{C}})_{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}$ indexed by the objects of $\boldsymbol{\mathcal{C}}^\circ$, where for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, $\phi_{\mathcal{C}} \in \overrightarrow{\boldsymbol{\mathcal{D}}}(F(\mathcal{C}), G(\mathcal{C}))$, such that for any $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$ and any morphism $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$, the following diagram commutes:

$$\begin{array}{ccc} F(\mathcal{A}) & \xleftarrow{F(\psi)} & F(\mathcal{B}) \\ {\scriptstyle \phi_{\mathcal{A}}}\downarrow & & \downarrow{\scriptstyle \phi_{\mathcal{B}}} \\ G(\mathcal{A}) & \xleftarrow[G(\psi)]{} & G(\mathcal{B}) \end{array}$$

We indicate this by writing: "$\Phi : F \Longrightarrow G$".[6]

For a concrete example, recall the contravariant functor $\mathfrak{C}_{\mathrm{b}} : \mathsf{Top}^{\mathrm{op}} \Longmapsto \mathsf{Vec}$ introduced in §5B. By "forgetting" the vector space structure of the image objects (§5A), we can also interpret this as a functor $\mathfrak{C}_{\mathrm{b}} : \mathsf{Top}^{\mathrm{op}} \Longmapsto \mathsf{Set}$. Now let $\mathcal{X}$ be a topological space, and consider the contravariant hom functor $\overrightarrow{\mathsf{Top}}(\bullet, \mathcal{X}) : \mathsf{Top}^{\mathrm{op}} \Longmapsto \mathsf{Set}$ from §5C. Let $u : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded, continuous, real-valued function. For any topological space $\mathcal{S}$ and any continuous function $\alpha : \mathcal{S} \longrightarrow \mathcal{X}$, we can compose α with u to get a continuous function $u \circ \alpha : \mathcal{S} \longrightarrow \mathbb{R}$. This yields a function $\phi_{\mathcal{S}} : \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X}) \longrightarrow \mathfrak{C}_{\mathrm{b}}(\mathcal{S})$, defined by setting $\phi_{\mathcal{S}}(\alpha) := u \circ \alpha$ for all $\alpha \in \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X})$. If $\mathcal{T}$ is another topological space, and $\psi : \mathcal{S} \longrightarrow \mathcal{T}$ is a continuous function, then it is easily verified that the following diagram commutes:

$$\begin{array}{ccc} \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X}) & \xleftarrow{\overleftarrow{\psi}} & \overrightarrow{\mathsf{Top}}(\mathcal{T}, \mathcal{X}) \\ {\scriptstyle \phi_{\mathcal{S}}}\downarrow & & \downarrow{\scriptstyle \phi_{\mathcal{T}}} \\ \mathfrak{C}_{\mathrm{b}}(\mathcal{S}) & \xleftarrow[\psi^*]{} & \mathfrak{C}_{\mathrm{b}}(\mathcal{T}) \end{array} \qquad \text{where } \overleftarrow{\psi} := \overrightarrow{\mathsf{Top}}(\psi, \mathcal{X}) \text{ and } \psi^* := \mathfrak{C}_{\mathrm{b}}(\psi).^{7}$$

[6] One can likewise define natural transformations between *covariant* functors. But these are never used in this paper.

Thus, the collection $\Phi = (\phi_{\mathcal{S}})_{\mathcal{S}\in\mathsf{Top}^\circ}$ is a natural transformation from the functor $\overrightarrow{\mathsf{Top}}(\bullet, \mathcal{X})$ to the functor $\mathfrak{C}_{\mathrm{b}}$. This example will play a key role in Section 6.

6 Global SEU representations: definition

6A Partially ordered vector spaces. A *partially ordered vector space* (POVS) is a (real) vector space $\mathcal{V}$ equipped with a partial order (a transitive, antisymmetric binary relation) $>$ that is compatible with addition and scalar multiplication in the obvious way.[8] For example: $\mathbb{R}$ is a POVS with the obvious linear order. Here are some other examples.

Example 6.1. For any set $\mathcal{S}$, the vector space $\mathbb{R}^{\mathcal{S}}$ of real-valued functions on $\mathcal{S}$ is a POVS with the pointwise dominance order. If $\mathcal{S}$ is a measurable space, then the space $\mathfrak{L}(\mathcal{S})$ of measurable real-valued functions is a POVS. If $\mathcal{S}$ is a topological space, then the space $\mathfrak{C}(\mathcal{S})$ of continuous real-valued functions is a POVS. If $\mathcal{S}$ is a metric space, then the space $\mathfrak{LL}(\mathcal{S})$ of locally Lipschitz real-valued functions is a POVS. If $\mathcal{S}$ is a differentiable manifold, then the space $\mathfrak{C}^\infty(\mathcal{S})$ of smooth real-valued functions is a POVS. $\diamondsuit$

6B Unitary POVS. Let $\mathcal{V}$ be a POVS. An *order unit* for $\mathcal{V}$ is an element $u \in \mathcal{V}$ such that $u > 0$ and such that, for any $v > 0$ there is some $r \in \mathbb{R}_+$ such that $r\,u \geq v$. A *unitary* partially ordered vector space is a POVS equipped with an order unit. For example, 1 is an order unit for $\mathbb{R}$, making $\mathbb{R}$ a unitary POVS. Here are some other examples.

Example 6.2. We continue the notation of Example 6.1.

(a) Let $\mathcal{S}$ be an abstract set. Let $\ell^\infty(\mathcal{S})$ be the POVS of all *bounded* elements of $\mathbb{R}^{\mathcal{S}}$. (Thus, $\ell^\infty(\mathcal{S}) = \mathbb{R}^{\mathcal{S}}$ only if $\mathcal{S}$ is finite.) This is a unitary POVS: the constant 1-valued function $\mathbf{1}_{\mathcal{S}}$ is an order unit for $\ell^\infty(\mathcal{S})$.

(b) Let $\mathcal{S}$ be a measurable space. Let $\mathfrak{L}^\infty(\mathcal{S})$ be the POVS of all *bounded* elements of $\mathfrak{L}(\mathcal{S})$. (Thus, $\mathfrak{L}^\infty(\mathcal{S}) = \mathfrak{L}(\mathcal{S})$ only if the sigma-algebra of $\mathcal{S}$ is finite.) This is a unitary POVS: $\mathbf{1}_{\mathcal{S}}$ is an order unit for $\mathfrak{L}^\infty(\mathcal{S})$.

(c) Let $\mathcal{S}$ be a topological space. Let $\mathfrak{C}_{\mathrm{b}}(\mathcal{S})$ be the POVS of all *bounded* elements of $\mathfrak{C}(\mathcal{S})$. (Thus, $\mathfrak{C}_{\mathrm{b}}(\mathcal{S}) = \mathfrak{C}(\mathcal{S})$ only if $\mathcal{S}$ is compact.) This is a unitary POVS with order unit $\mathbf{1}_{\mathcal{S}}$.

(d) Let $\mathcal{S}$ be a metric space. Let $\mathfrak{LL}_{\mathrm{b}}(\mathcal{S})$ be the POVS of *bounded*, locally Lipschitz real-valued functions. It is a unitary POVS with order unit $\mathbf{1}_{\mathcal{S}}$.

(e) Let $\mathcal{S}$ be a differentiable manifold. Let $\mathfrak{C}^\infty_{\mathrm{b}}(\mathcal{S})$ be the POVS of *bounded* elements of $\mathfrak{C}^\infty(\mathcal{S})$. (So $\mathfrak{C}^\infty_{\mathrm{b}}(\mathcal{S}) = \mathfrak{C}^\infty(\mathcal{S})$ only if $\mathcal{S}$ is compact.) This is a unitary POVS with order unit $\mathbf{1}_{\mathcal{S}}$. $\diamondsuit$

An order-preserving linear transformation from a unitary POVS $\mathcal{V}_1$ to another unitary POVS $\mathcal{V}_2$ *uniferent* if it sends the order unit of $\mathcal{V}_1$ to the order unit of $\mathcal{V}_2$.

[7] *Proof.* For any $\alpha : \mathcal{T}\longrightarrow\mathcal{X}$, $\phi_{\mathcal{S}} \circ \overleftarrow{\psi}(\alpha) = \phi_{\mathcal{S}}[\overleftarrow{\psi}(\alpha)] = u\circ(\alpha\circ\psi) = (u\circ\alpha)\circ\psi = \psi^*[u\circ\alpha] = \psi^*\circ\phi_{\mathcal{T}}(\alpha)$.

[8] For simplicity, this paper works with partially ordered *real* vector spaces. But many the statements are true for partially ordered vector spaces defined over any linearly ordered field.

6C Utility Frames. Let UPOVS be the category of unitary partially ordered vector spaces and uniferent, order-preserving, linear transformations. A *utility frame* on $\boldsymbol{\mathcal{C}}$ is a contravariant functor $L : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{UPOVS}$.[9]

Example 6.3. (a) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$. For any $\mathcal{S} \in \mathsf{Set}^\circ$, let $L(\mathcal{S}) := \ell^\infty(\mathcal{S})$ with order unit $\mathbf{1}_\mathcal{S}$, as in Example 6.2(a). For any $\mathcal{S}_1, \mathcal{S}_2 \in \mathsf{Set}^\circ$ and $\phi : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, define $L(\phi) : \ell^\infty(\mathcal{S}_2) \longrightarrow \ell^\infty(\mathcal{S}_1)$ by setting $L(\phi)[v] := v \circ \phi$ for all bounded functions $v : \mathcal{S}_2 \longrightarrow \mathbb{R}$.

(b) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Meas}$. For any $\mathcal{S} \in \mathsf{Meas}^\circ$, let $L(\mathcal{S}) := \mathfrak{L}^\infty(\mathcal{S})$ with order unit $\mathbf{1}_\mathcal{S}$, as in Example 6.2(b). For any $\mathcal{S}_1, \mathcal{S}_2 \in \mathsf{Meas}^\circ$ and measurable $\phi : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, define $L(\phi) : \mathfrak{L}^\infty(\mathcal{S}_2) \longrightarrow \mathfrak{L}^\infty(\mathcal{S}_1)$ by setting $L(\phi)[v] := v \circ \phi$ for all bounded measurable $v : \mathcal{S}_2 \longrightarrow \mathbb{R}$.

(c) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Top}$. For any $\mathcal{S} \in \mathsf{Top}^\circ$, let $L(\mathcal{S}) := \mathfrak{C}_{\rm b}(\mathcal{S})$ with order unit $\mathbf{1}_\mathcal{S}$, as in Example 6.2(c). For any $\mathcal{S}_1, \mathcal{S}_2 \in \mathsf{Top}^\circ$ and continuous map $\phi : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, define $L(\phi) : \mathfrak{C}_{\rm b}(\mathcal{S}_2) \longrightarrow \mathfrak{C}_{\rm b}(\mathcal{S}_1)$ by setting $L(\phi)[v] := v \circ \phi$ for all bounded continuous $v : \mathcal{S}_2 \longrightarrow \mathbb{R}$.

(d) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Metr}$. For any $\mathcal{S} \in \mathsf{Metr}^\circ$, let $L(\mathcal{S}) := \mathfrak{LL}_{\rm b}(\mathcal{S})$ with order unit $\mathbf{1}_\mathcal{S}$, as in Example 6.2(d). For any $\mathcal{S}_1, \mathcal{S}_2 \in \mathsf{Metr}^\circ$ and short map $\phi : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, define $L(\phi) : \mathfrak{LL}_{\rm b}(\mathcal{S}_2) \longrightarrow \mathfrak{LL}_{\rm b}(\mathcal{S}_1)$ by setting $L(\phi)[v] := v \circ \phi$ for all bounded locally Lipschitz $v : \mathcal{S}_2 \longrightarrow \mathbb{R}$.

(e) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Diff}$. For any $\mathcal{S} \in \mathsf{Diff}^\circ$, let $L(\mathcal{S}) := \mathfrak{C}^\infty_{\rm b}(\mathcal{S})$ with order unit $\mathbf{1}_\mathcal{S}$, as in Example 6.2(e). For any $\mathcal{S}_1, \mathcal{S}_2 \in \mathsf{Diff}^\circ$ and smooth map $\phi : \mathcal{S}_1 \longrightarrow \mathcal{S}_2$, define $L(\phi) : \mathfrak{C}^\infty_{\rm b}(\mathcal{S}_2) \longrightarrow \mathfrak{C}^\infty_{\rm b}(\mathcal{S}_1)$ by setting $L(\phi)[v] := v \circ \phi$ for all bounded smooth $v : \mathcal{S}_2 \longrightarrow \mathbb{R}$. ♢

As these examples suggest, in general we will think of elements of $L(\mathcal{S})$ as abstract "utility functions" on $\mathcal{S}$. If $\boldsymbol{\mathcal{D}}$ is a subcategory of $\boldsymbol{\mathcal{C}}$, then let $L_{\upharpoonright \boldsymbol{\mathcal{D}}} : \boldsymbol{\mathcal{D}}^{\rm op} \Longmapsto \mathsf{UPOVS}$ be the *restriction* of L to a utility frame on $\boldsymbol{\mathcal{D}}$ (defined in the obvious way). Let $L : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{UPOVS}$ and $L' : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{UPOVS}$ be two utility frames. We say that L is a *utility subframe* of L' if $L(\mathcal{C})$ is a linear subspace of $L'(\mathcal{C})$ (with the same partial order and order unit) for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, and $L(\phi) = L'(\phi)_{\upharpoonright L(\mathcal{C})} : L(\mathcal{C}) \longrightarrow L(\mathcal{B})$ for all $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$. Example 6.3(b) is a utility subframe of Example 6.3(a) (restricted to Meas). Example 6.3(c) is a utility subframe of Example 6.3(b) (restricted to Top). Example 6.3(d) is a utility subframe of Example 6.3(c) (restricted to Metr), and Example 6.3(e) is a utility subframe of Example 6.3(c) (restricted to Diff). In Section 7, we construct a general utility frame that subsumes all the cases in Example 6.3 as utility subframes.

6D Utility functionals. Fix a utility frame $L : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{UPOVS}$. Let $F : \mathsf{UPOVS} \Longmapsto \mathsf{Set}$ be the forgetful functor from §5A, and let $\underline{L} := F \circ L : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{Set}$. (Thus, for any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, $\underline{L}(\mathcal{C})$ is the set of elements in $L(\mathcal{C})$, but regarded as a *set*, rather than an ordered vector space.) Let $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. Recall the contravariant hom functor $\overrightarrow{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) : \boldsymbol{\mathcal{C}}^{\rm op} \Longmapsto \mathsf{Set}$ from §5C. For any $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_1, \mathcal{S}_2)$, let $\overleftarrow{\phi} := \overrightarrow{\boldsymbol{\mathcal{C}}}(\phi, \mathcal{X})$; in other words, $\overleftarrow{\phi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_2, \mathcal{X}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_1, \mathcal{X})$ is the function defined by $\overleftarrow{\phi}(\alpha) := \alpha \circ \phi$ for all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}_2, \mathcal{X})$.

An (L-valued) *utility functional* for $\mathcal{X}$ is a natural transformation $U_\mathcal{X} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ (§5D). In other words, $U_\mathcal{X} = (U^\mathcal{C}_\mathcal{X})_{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}$, where for any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, $U^\mathcal{C}_\mathcal{X} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \underline{L}(\mathcal{C})$

[9] In the terminology of category theory, L is a UPOVS-valued *presheaf*.

is a function such that, for any $\mathcal{C}_1, \mathcal{C}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}_1, \mathcal{C}_2)$, the following diagram commutes:

$$\begin{array}{ccc} \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}_1,\mathcal{X}) & \xleftarrow{\overleftarrow{\phi}} & \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}_2,\mathcal{X}) \\ {\scriptstyle U_{\mathcal{X}}^{\mathcal{C}_1}}\Big\downarrow & & \Big\downarrow{\scriptstyle U_{\mathcal{X}}^{\mathcal{C}_2}} \\ \underline{L}(\mathcal{C}_1) & \xleftarrow[\underline{L}(\phi)]{} & \underline{L}(\mathcal{C}_2) \end{array} \qquad (1)$$

Example 6.4. (a) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$ and define L as in Example 6.3(a). Let $u : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded function. For any $\mathcal{C} \in \mathsf{Set}^\circ$ and any function $\alpha : \mathcal{C} \longrightarrow \mathcal{X}$, define $U_{\mathcal{X}}^{\mathcal{C}}(\alpha) := u \circ \alpha : \mathcal{C} \longrightarrow \mathbb{R}$. Then $U_{\mathcal{X}}^{\mathcal{C}}(\alpha) \in L(\mathcal{C})$. For any $\phi : \mathcal{C}_1 \longrightarrow \mathcal{C}_2$ and $\alpha : \mathcal{C}_2 \longrightarrow \mathcal{X}$, we have $U_{\mathcal{X}}^{\mathcal{C}_1} \circ \overleftarrow{\phi}(\alpha) = U_{\mathcal{X}}^{\mathcal{C}_1}[\overleftarrow{\phi}(\alpha)] = u \circ (\alpha \circ \phi) = (u \circ \alpha) \circ \phi = \underline{L}(\phi)[u \circ \alpha] = \underline{L}(\phi) \circ U_{\mathcal{X}}^{\mathcal{C}_2}(\alpha)$. Thus, the diagram (1) commutes.

(b) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Meas}$ and L is as in Example 6.3(b). Then we use the same construction as part (a), but we require $u : \mathcal{X} \longrightarrow \mathbb{R}$ to be both bounded and measurable.

(c) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Top}$ and L is as in Example 6.3(c). Then we use the same construction as part (a), but we require $u : \mathcal{X} \longrightarrow \mathbb{R}$ to be both bounded and continuous.

(d) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Metr}$ and L is as in Example 6.3(d). Then we use the same construction as part (a), but we require $u : \mathcal{X} \longrightarrow \mathbb{R}$ to be both bounded and locally Lipschitz.

(e) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Diff}$ and L is as in Example 6.3(e). Then we use the same construction as part (a), but we require $u : \mathcal{X} \longrightarrow \mathbb{R}$ to be both bounded and smooth. ♢

In all five cases of Example 6.4, u was an element of $L(\mathcal{X})$, and we defined $U_{\mathcal{X}}^{\mathcal{C}}(\alpha) = L(\alpha)[u]$ for all $\alpha : \mathcal{C} \longrightarrow \mathcal{X}$.[10] The next result shows that this is actually the general case.

Proposition 6.5 *Let $L : \boldsymbol{\mathcal{C}}^{\text{op}} \Longmapsto \mathsf{UPOVS}$ be a utility frame, and let $\mathcal{X} \in \boldsymbol{\mathcal{C}}^\circ$.*

(a) *Let $u \in L(\mathcal{X})$. For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, define $U_{\mathcal{X}}^{\mathcal{C}} : \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \underline{L}(\mathcal{C})$ as follows:*

$$U_{\mathcal{X}}^{\mathcal{C}}(\alpha) \quad := \quad L(\alpha)[u], \quad \text{for all } \alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}). \qquad (2)$$

Then $U_{\mathcal{X}} = (U_{\mathcal{X}}^{\mathcal{C}})_{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}$ is a utility functional on $\mathcal{X}$.

(b) *Conversely, for any utility functional $U_{\mathcal{X}}$ on $\mathcal{X}$, there is a unique $u \in L(\mathcal{X})$ such that $U_{\mathcal{X}}$ arises from u as in equation (2).*

6E Local SEU representations. Let $\boldsymbol{\mathcal{C}}$ be a category, and let $L : \boldsymbol{\mathcal{C}}^{\text{op}} \longrightarrow \mathsf{UPOVS}$ be a utility frame. For any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, a ***belief about*** $\mathcal{C}$ is an order-preserving linear functional $\rho : L(\mathcal{C}) \longrightarrow \mathbb{R}$, such that $\rho(\mathbf{1}) = 1$. In other words: ρ is a UPOVS-morphism from $L(\mathcal{C})$ to $\mathbb{R}$, where we regard $\mathbb{R}$ as a unitary POVS with order unit 1.

[10] Recall: if α is a morphism from $\mathcal{C}$ into $\mathcal{X}$, then $L(\alpha)$ is a function from $L(\mathcal{X})$ to $L(\mathcal{C})$, because L is a contravariant functor. Thus, $L(\alpha)[u] \in L(\mathcal{C})$ because $u \in L(\mathcal{X})$.

Example 6.6. Let $\boldsymbol{\mathcal{C}} = \mathsf{Set}$ and define $L : \mathsf{Set}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ as in Example 6.3(a). Let $\mathcal{S}$ be a set, let $\wp(\mathcal{S})$ be the power set of $\mathcal{S}$, and let μ be a probability measure on $\wp(\mathcal{S})$. Define $\rho_{\mathcal{S}} : L(\mathcal{S}) \longrightarrow \mathbb{R}$ by setting $\rho_{\mathcal{S}}(v) := \int_{\mathcal{S}} v \, \mathrm{d}\mu$ for all $v \in L(\mathcal{S})$. Then $\rho_{\mathcal{S}}$ is an order-preserving linear functional with $\rho_{\mathcal{S}}(\mathbf{1}) = 1$; thus, it is a belief about $\mathcal{S}$. (Similar examples appear in categories such as Meas, Top, and Metr; see Proposition 6.8 below.) ♢

Now let $\mathcal{S}$ and $\mathcal{X}$ be objects in $\boldsymbol{\mathcal{C}}^{\circ}$, and let $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ be a preference order on $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$. An ($L$-valued) *local subjective expected utility* (SEU) *representation* for $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ consists of a belief $\rho_{\mathcal{S}}$ about $\mathcal{S}$ and an L-valued utility functional $U_{\mathcal{X}}$ on $\mathcal{X}$, such that:

$$\text{and all } \alpha, \beta \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X}), \qquad \alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}} \beta \quad \Longleftrightarrow \quad \rho_{\mathcal{S}}\left[U^{\mathcal{S}}_{\mathcal{X}}(\alpha)\right] \geq \rho_{\mathcal{S}}\left[U^{\mathcal{S}}_{\mathcal{X}}(\beta)\right]. \tag{3}$$

Example 6.7. Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$ and define $L : \mathsf{Set}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ as in Example 6.3(a). Let $\mathcal{S}$ and $\mathcal{X}$ be sets, and let ρ be a belief on $\mathcal{S}$ defined by a probability measure μ on $\wp(\mathcal{S})$, as in Example 6.6. Let $u : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded function, and define utility functional $U_{\mathcal{X}} : \vec{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(a). So for any $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{C})$, we have $U^{\mathcal{S}}_{\mathcal{X}}(\alpha) = u \circ \alpha$. Thus, $\rho_{\mathcal{S}}\left[U^{\mathcal{S}}_{\mathcal{X}}(\alpha)\right] = \int_{\mathcal{S}} u \circ \alpha \ \mathrm{d}\mu$. So for all $\alpha, \beta \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$, formula (3) says $\alpha \succcurlyeq^{\mathcal{S}}_{\mathcal{X}} \beta \iff \int_{\mathcal{S}} u \circ \alpha \, \mathrm{d}\mu \geq \int_{\mathcal{S}} u \circ \beta \, \mathrm{d}\mu$. ♢

Suppose that $\trianglerighteq^{\mathrm{xa}}$ is an ex ante preference structure on a decision environment $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$. For each choice of state place $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}^{\circ}$ and outcome place $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}^{\circ}$, we could construct a local SEU representation (3) for the preference order $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$. But such an inchoate collection of local SEU representations would be unsatisfactory, for two reasons.

First, as noted in Section 1, it would be strange if the agent had *different* beliefs about $\mathcal{S}$ for each possible outcome place $\mathcal{X}$. Likewise, it would be strange if she had a *different* utility functional over $\mathcal{X}$ for each possible state place $\mathcal{S}$. For each $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}^{\circ}$, we want a *single* belief $\rho_{\mathcal{S}}$ that yields a local SEU representation (3) for *every* choice of $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}^{\circ}$. Likewise, for each $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}^{\circ}$, we want a *single* utility functional $U_{\mathcal{X}}$ that yields a local SEU representation (3) for *every* choice of $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}^{\circ}$.

Second, the agent's preferences are *congruent* across different choices of $\mathcal{S}$ in $\boldsymbol{\mathcal{S}}^{\circ}$ and different choices of $\mathcal{X}$ in $\boldsymbol{\mathcal{X}}^{\circ}$, as formalized by the properties (BP) and (TP) from Section 4. Thus, for every pair of state places $\mathcal{S}_1$ and $\mathcal{S}_2$ in $\boldsymbol{\mathcal{S}}^{\circ}$ and every morphism $\phi \in \vec{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, we want ϕ to somehow transform the agent's beliefs about $\mathcal{S}_1$ into her beliefs about $\mathcal{S}_2$. Likewise, for every pair of outcome places $\mathcal{X}_1$ and $\mathcal{X}_2$ in $\boldsymbol{\mathcal{X}}^{\circ}$ and every morphism $\phi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}_1, \mathcal{X}_2)$, we want ϕ to somehow transform the agent's utility over $\mathcal{X}_1$ into her utility over $\mathcal{X}_2$. Such a collection of beliefs and utilities would provide a "global" SEU representation for the entire ex ante preference structure $\trianglerighteq^{\mathrm{xa}}$. In the following subsections, we shall develop the components that we need to define such a global SEU representation.

6F Belief systems. Let $\mathbb{R}$ be the unitary POVS of real numbers, with order unit 1. Let $L : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ be a utility frame. Let $\boldsymbol{\mathcal{S}}$ be a subcategory of $\boldsymbol{\mathcal{C}}$ (e.g. the category of state places in a decision environment). Let $L_{\upharpoonright \boldsymbol{\mathcal{S}}} : \boldsymbol{\mathcal{S}}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ be the restriction of L to $\boldsymbol{\mathcal{S}}$. A *belief system* for $\boldsymbol{\mathcal{S}}$ is a collection $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}}$, where, for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, $\rho_{\mathcal{S}}$ is a belief about

$\mathcal{S}$ (i.e. a uniferent, order-preserving linear transformation $\rho_{\mathcal{S}} : L(\mathcal{S}) \longrightarrow \mathbb{R}$) such that, for any $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, and any $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, the following diagram commutes:[11]

$$\begin{array}{ccc} L(\mathcal{S}_1) & \xleftarrow{\;L(\phi)\;} & L(\mathcal{S}_2) \\ & {\scriptstyle \rho_{\mathcal{S}_1}} \searrow \quad \swarrow {\scriptstyle \rho_{\mathcal{S}_2}} & \\ & \mathbb{R} & \end{array} \tag{4}$$

Proposition 6.8 *Let $L : \boldsymbol{\mathcal{C}}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ be a utility frame, let $\boldsymbol{\mathcal{S}}$ be a subcategory of $\boldsymbol{\mathcal{C}}$, and let $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a belief system.*

(a) *Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$, and $L : \mathsf{Set}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ is as in Example 6.3(a). Then for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there is a unique finitely additive probability measure $\mu_{\mathcal{S}}$ on the power set $\wp(\mathcal{S})$, such that $\rho_{\mathcal{S}} : L(\mathcal{S}) \longrightarrow \mathbb{R}$ is defined by*

$$\rho_{\mathcal{S}}(v) \quad = \quad \int_{\mathcal{S}} v \ \ \mathrm{d}\mu_{\mathcal{S}}, \qquad \textit{for all } v \in L(\mathcal{S}). \tag{5}$$

Furthermore, for all $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, we have

$$\phi(\mu_{\mathcal{S}_1}) = \mu_{\mathcal{S}_2}, \qquad \textit{for all } \phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2).^{12} \tag{6}$$

(b) *Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Meas}$, and $L : \mathsf{Meas}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ is as in Example 6.3(b). Then for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there is a unique finitely additive probability measure $\mu_{\mathcal{S}}$ on the sigma-algebra of $\mathcal{S}$, satisfying equations (5) and (6).*

(c) *Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Top}$, and $L : \mathsf{Top}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ is as in Example 6.3(c). Suppose that all objects in $\boldsymbol{\mathcal{S}}^\circ$ are normal Hausdorff spaces. Then for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there is a unique finitely additive normal probability measure $\mu_{\mathcal{S}}$ on the Borel sigma-algebra of $\mathcal{S}$, satisfying equations (5) and (6).*[13]

Furthermore, if $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{S}$ is compact, then $\mu_{\mathcal{S}}$ is countably additive.[14]

(d) *Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Metr}$, and $L : \mathsf{Metr}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ is as in Example 6.3(d). Suppose that all objects in $\boldsymbol{\mathcal{S}}^\circ$ are compact. Then for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there is a unique (countably additive) Borel probability measure $\mu_{\mathcal{S}}$ on $\mathcal{S}$ satisfying equations (5) and (6).*[15]

[11] In the terminology of category theory, $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ is a *co-cone* from the functor $L_{\upharpoonright \boldsymbol{\mathcal{S}}}$ to the object $\mathbb{R}$ in the category UPOVS.

[12] This means: for all $\mathcal{B} \subseteq \mathcal{S}_2$, $\mu_{\mathcal{S}_2}(\mathcal{B}) = \mu_{\mathcal{S}_1}[\phi^{-1}(\mathcal{B})]$.

[13] A Borel measure μ on a space $\mathcal{S}$ is *normal* if it is both *inner-regular* and *outer-regular*. In other words, for any Borel subset $\mathcal{B} \subseteq \mathcal{S}$, we have $\mu(\mathcal{B}) = \sup_{\substack{\mathcal{F} \subseteq \mathcal{B} \\ \mathcal{K} \text{ closed}}} \mu(\mathcal{F}) = \inf_{\substack{\mathcal{B} \subseteq \mathcal{O} \subseteq \mathcal{S} \\ \mathcal{O} \text{ open}}} \mu(\mathcal{O})$.

[14] This statement is also true when $\mathcal{S}$ is a *locally* compact Hausdorff space, if $L : \mathsf{Top}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ is the utility frame in which $L(\mathcal{S})$ is the space of continuous real-valued functions on $\mathcal{S}$ with *compact support*.

[15] Every Borel probability measure on a metric space is normal (Aliprantis and Border, 2006, Thm.12.5).

Remark. Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Diff}$, and $L : \mathsf{Diff}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ is as in Example 6.3(e). If $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ is compact, then $\mathfrak{C}^\infty(\mathcal{S})$ coincides with the *Schwartz space* of $\mathcal{S}$.[16] The dual of Schwartz space is the space of *tempered distributions* (see Folland 1984, §8.5, p.258, or Katznelson 2004, §VI.4, p.146.) Thus, we could state a "part (e)" of Proposition 6.8 stating that, for any compact manifold $\mathcal{S}$, there is a unique tempered distribution $\mu_{\mathcal{S}}$ on $\mathcal{S}$, satisfying analogies to equations (5) and (6). ♢

The situation described in the last sentence of Proposition 6.8(c) arises in a wide variety of cases. A ***Riesz space*** is a POVS $\mathcal{V}$ where the partial order is a ***lattice***—in other words, any $u, v \in \mathcal{V}$ have a supremum $u \vee v$ and an infimum $u \wedge v$ in $\mathcal{V}$. All of the spaces in Example 6.2 are Riesz spaces, where $(u \vee v)(s) := \max\{u(s), v(s)\}$ and $(u \wedge v)(s) := \min\{u(s), v(s)\}$ for all $u, v \in \mathcal{V}$ and $s \in \mathcal{S}$. A Riesz space $\mathcal{V}$ is ***Archimedean*** if, for any $u, v \in \mathcal{V}$, if $0 \leq n\, u \leq v$ for all $n \in \mathbb{N}$, then $u = 0$. All of the spaces in Example 6.2 are Archimedean.

For any $v \in \mathcal{V}$, define $|v| := (v \vee 0) + ((-v) \vee 0)$; then $|v| \geq 0$. If $\mathcal{V}$ has an order unit 1, then for all $v \in \mathcal{V}$ we define $\|v\| := \min\{r \in \mathbb{R}_+;\ \ r\,1 \geq |v|\}$. This is a norm on $\mathcal{V}$. We say that $\mathcal{V}$ is an ***M-space*** if it is a Banach space with this norm.

Let $\mathsf{UARiesz}$ be the category of unitary Archimedean Riesz spaces and uniferent order-preserving linear functions (a subcategory of UPOVS). Let MSpace be the category of M-spaces and *continuous*, uniferent order-preserving linear functions (a subcategory of $\mathsf{UARiesz}$). Let CHS be the category of compact Hausdorff spaces and continuous maps.

Proposition 6.9 *Let $\boldsymbol{\mathcal{C}}$ be a category, and let $L : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \mathsf{UARiesz}$ be a utility frame.*

- **(a)** *There is a functor $H : \boldsymbol{\mathcal{C}} \Longmapsto \mathsf{CHS}$, and for all $\mathcal{S} \in \boldsymbol{\mathcal{C}}^\circ$, there is a uniferent, order-preserving linear function $\text{ß}_{\mathcal{S}} : L(\mathcal{S}) \longrightarrow \mathfrak{C}(H(\mathcal{S}))$, whose image is both order-dense and uniformly dense in $\mathfrak{C}(H(\mathcal{S}))$. In fact, if $L(\mathcal{S})$ is an M-space, then $\text{ß}_{\mathcal{S}}$ is a Riesz isomorphism from $L(\mathcal{S})$ to $\mathfrak{C}(H(\mathcal{S}))$.*

- **(b)** *The collection $(\text{ß}_{\mathcal{S}})_{\mathcal{S} \in \boldsymbol{\mathcal{C}}^\circ}$ determines a natural transformation $\text{ß} : L \Longrightarrow \mathfrak{C} \circ H$. If L is a functor into MSpace, then this is a natural isomorphism.*

- **(c)** *Let $\boldsymbol{\mathcal{S}}$ be a subcategory of $\boldsymbol{\mathcal{C}}$, and let $(\rho_{\mathcal{S}})_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a belief system on $\boldsymbol{\mathcal{S}}$. For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there is a unique Borel probability measure $\mu_{\mathcal{S}}$ on $H(\mathcal{S})$ such that*

$$\rho_{\mathcal{S}}(v) \;=\; \int_{H(\mathcal{S})} \text{ß}_{\mathcal{S}}(v)\ \mathrm{d}\mu_{\mathcal{S}}, \qquad \textit{for all } v \in L(\mathcal{S}).^{17} \tag{7}$$

Furthermore, for any $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$ and $\phi \in \vec{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, we have

$$\phi^*(\mu_{\mathcal{S}_1}) \quad = \quad \mu_{\mathcal{S}_2}, \tag{8}$$

where $\phi^* := H(\phi) : H(\mathcal{S}_1) \longrightarrow H(\mathcal{S}_2)$.

[16] This is the space of smooth real-valued functions on $\mathcal{S}$ whose derivatives all "vanish at infinity" faster than the reciprocal of any polynomial. But if $\mathcal{S}$ is compact, this is true for any smooth function.

[17] To understand this equation, recall that $\text{ß}_{\mathcal{S}}(v)$ is a real-valued function on $H(\mathcal{S})$.

The natural transformation in part (b) means that, for all intents and purposes, *any* UARiesz-valued utility frame is obtained by transforming each object in $\boldsymbol{\mathcal{C}}$ into a compact Hausdorff space (via some functor H), and then considering some Riesz space of continuous real-valued functions on that compact Hausdorff space. For an MSpace-valued utility frame, the natural isomorphism yields an even stronger statement: each object in $\boldsymbol{\mathcal{C}}$ is transformed into the M-space of *all* continuous real-valued functions on its H-associated compact Hausdorff space. In part (c), equations (7) and (8) can be seen as versions of equations (5) and (6) from Proposition 6.8, translated through the functor H.

6G Positive affine transformations. A *positive affine transformation* is an increasing bijection $\phi : \mathbb{R} \longrightarrow \mathbb{R}$ of the form $\phi(r) = a\,r + b$ for all $r \in \mathbb{R}$, where $a > 0$ and $b \in \mathbb{R}$ are constants. If $b \geq 0$ and $0 < a \leq 1$, then we say ϕ is an *affine contraction*. The set of all positive affine transformations forms a group Aff under composition, which we can regard as a single-object category, as explained in Section 2.

Let $\mathcal{V}$ be a unitary POVS with order unit $\mathbf{1}$. A *positive affine transformation* of $\mathcal{V}$ is an order-preserving bijection $\phi : \mathcal{V} \longrightarrow \mathcal{V}$ of the form $\phi(\mathbf{v}) = a\,\mathbf{v} + b\mathbf{1}$ for all $\mathbf{v} \in \mathcal{V}$, where $a > 0$ and $b \in \mathbb{R}$ are constants.[18] The set of all positive affine transformations of $\mathcal{V}$ forms a group $\mathrm{Aff}(\mathcal{V})$ under composition. There is clearly a natural group isomorphism $\mathsf{Aff} \longrightarrow \mathrm{Aff}(\mathcal{V})$. For any $\phi \in \mathsf{Aff}$, let $\phi_{\mathcal{V}}$ denote the corresponding element of $\mathrm{Aff}(\mathcal{V})$.

6H Utility systems. For any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, recall the covariant hom functor $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \bullet) : \boldsymbol{\mathcal{C}} \Longmapsto \mathsf{Set}$ from §5C. For any $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{Y})$, let $\vec{\phi} := \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \phi)$; in other words, $\vec{\phi} : \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{Y})$ is the function defined by $\vec{\phi}(\alpha) := \phi \circ \alpha$ for all $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$.

Let $L : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ be a utility frame. Let $\boldsymbol{\mathcal{X}}$ be a subcategory of $\boldsymbol{\mathcal{C}}$ (for example, the category of outcome places in a decision environment). An (L-valued) *utility system* on $\boldsymbol{\mathcal{X}}$ is an ordered pair (U, A), in which

- $A : \boldsymbol{\mathcal{X}} \Longmapsto \mathsf{Aff}$ is a functor; and
- $U = (U_{\mathcal{X}})_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$, where $U_{\mathcal{X}} : \vec{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ is a utility functional for each $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$;

such that, for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, all $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ and all morphisms $\phi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, the following diagram commutes:

$$\begin{array}{ccc} \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) & \xrightarrow{\ \vec{\phi}\ } & \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{Y}) \\ {\scriptstyle U_{\mathcal{X}}^{\mathcal{C}}}\downarrow & & \downarrow{\scriptstyle U_{\mathcal{Y}}^{\mathcal{C}}} \\ \underline{L}(\mathcal{C}) & \xrightarrow[\widehat{\phi}_{L(\mathcal{C})}]{} & \underline{L}(\mathcal{C}) \end{array} \qquad \text{where } \widehat{\phi} \; := \; A(\phi). \qquad (9)$$

Here, $\widehat{\phi}_{L(\mathcal{C})}$ is the automorphism of the set $\underline{L}(\mathcal{C})$ obtained from the affine transformation $\widehat{\phi}$. We shall call A the *affinity functor* of the utility system.

[18] ϕ is not linear if $b \neq 0$, so it is generally *not* a morphism in the category POVS.

Example 6.10. Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$, Meas, Top, Metr or Diff, and define L as in Example 6.3. For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_\mathcal{X} : \mathcal{X} \longrightarrow \mathbb{R}$ be a function (measurable, continuous, Lipschitz or smooth, as appropriate), and define the utility functional $U_\mathcal{X} = (U^\mathcal{C}_\mathcal{X})_{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4. Let $A : \boldsymbol{\mathcal{X}} \longrightarrow \mathsf{Aff}$ be a functor, and for all $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, suppose that $u_\mathcal{Y} \circ \phi = \widehat{\phi} \circ u_\mathcal{X}$, where $\widehat{\phi} := A(\phi)$ (an affine function from $\mathbb{R}$ to itself). Then the collection $(U^\mathcal{C}_\mathcal{X})^{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ together with A is a utility system on $\boldsymbol{\mathcal{X}}$.

For a concrete example, consider the table in Example 4.1(b) that specifies utility functions u_1, u_2, u_3, u_{12}, u_{23}, and u_{31} for the six outcome spaces $\mathcal{X}_1$, $\mathcal{X}_2$, $\mathcal{X}_3$, $\mathcal{X}_{12}$, $\mathcal{X}_{23}$, and $\mathcal{X}_{31}$ from Example 1.1. Recall that the $\boldsymbol{\mathcal{X}}$-morphisms between these objects are the inclusion maps shown in Figure 1. By inspecting the table in Example 4.1(b), one can see that:

$$\begin{array}{rclcrclcrcl} u_1 \circ \iota^1_{12} &=& u_{12} + 4, && u_2 \circ \iota^2_{23} &=& u_{23} + 4, && u_3 \circ \iota^3_{31} &=& u_{31} + 4, \\ u_1 \circ \iota^1_{31} &=& u_{31}, && u_2 \circ \iota^2_{12} &=& u_{12}, & \text{and} & u_3 \circ \iota^3_{23} &=& u_{23}. \end{array}$$

Let $I : \mathbb{R} \longrightarrow \mathbb{R}$ be the identity map, and let $\phi(r) := r + 4$ for all $r \in \mathbb{R}$. Then $I, \phi \in \mathsf{Aff}$. Consider the functor $A : \boldsymbol{\mathcal{X}} \Longmapsto \mathsf{Aff}$ defined by $A(\iota^1_{12}) = A(\iota^2_{23}) = A(\iota^3_{31}) = \phi$ while $A(\iota^2_{12}) = A(\iota^3_{23}) = A(\iota^1_{31}) = I$.[19] If we define $U_\mathcal{X} = (U^\mathcal{C}_\mathcal{X})_{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}$ for all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$ as in Example 6.4, then $(U^\mathcal{C}_\mathcal{X})^{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ together with A is a utility system on $\boldsymbol{\mathcal{X}}$. ♢

6I Cardinal equivalence. $L : \boldsymbol{\mathcal{C}}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$ be a utility frame, and let (U, A) and $(\widehat{U}, \widehat{A})$ be two L-valued utility systems on a subcategory $\boldsymbol{\mathcal{X}}$. For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $\gamma^\mathcal{X} \in \mathsf{Aff}$. The collection $\Gamma = (\gamma^\mathcal{X})_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ is a *cardinal equivalence* from (U, A) to $(\widehat{U}, \widehat{A})$ if:

- For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, and all $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$, $\widehat{U}^\mathcal{C}_\mathcal{X}(\phi) = \gamma^\mathcal{X}_{L(\mathcal{C})} \circ U^\mathcal{C}_\mathcal{X}(\phi)$.
- For all $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, $\widehat{A}(\phi) = \gamma^\mathcal{Y}_{L(\mathcal{C})} \circ A(\phi) \circ (\gamma^\mathcal{X}_{L(\mathcal{C})})^{-1}$.

We then say that (U, A) and $(\widehat{U}, \widehat{A})$ are *cardinally equivalent*. As the name suggests, this means that (U, A) and $(\widehat{U}, \widehat{A})$ encode the same cardinal utility information. It is easily verified that we can replace (U, A) with $(\widehat{U}, \widehat{A})$ in a global SEU representation (3) without changing the ex ante preference structure.

Example 6.11. Again consider the table in Example 4.1(b) that specifies utility functions for the six outcome spaces $\mathcal{X}_1$, $\mathcal{X}_2$, $\mathcal{X}_3$, $\mathcal{X}_{12}$, $\mathcal{X}_{23}$, and $\mathcal{X}_{31}$ from Example 1.1. Let (U, A) be the resulting utility system, as explained at the end of Example 6.10. Now consider the following collection of utility functions:

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
$\mathcal{X}_1$	2	4	6	8	10	12						
$\mathcal{X}_2$					3	6	9	12	15	18		
$\mathcal{X}_3$	25	26							21	22	23	24
$\mathcal{X}_{12}$					15	25						
$\mathcal{X}_{23}$									107	207		
$\mathcal{X}_{31}$	1	2										

[19] Since the category $\boldsymbol{\mathcal{X}}$ in this example has no nontrivial compositions of morphisms, *any* function mapping each morphism to an element of Aff is automatically a functor.

Let $(\widehat{U}, \widehat{A})$ be the resulting utility system.[20] Define

$$\begin{array}{rclrclrcl} \gamma^{\mathcal{X}_1}(r) & := & 2\,r, & \gamma^{\mathcal{X}_2}(r) & := & 3\,r, & \gamma^{\mathcal{X}_3}(r) & := & r+20, \\ \gamma^{\mathcal{X}_{12}}(r) & := & 10\,r+5, & \gamma^{\mathcal{X}_{23}}(r) & := & 100\,r+7, \quad \text{and} & \gamma^{\mathcal{X}_{31}}(r) & := & r, \quad \text{for all } r \in \mathbb{R}. \end{array}$$

Then $\Gamma = (\gamma^{\mathcal{X}})_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ is a cardinal equivalence from (U, A) to $(\widehat{U}, \widehat{A})$. ♢

6J Global SEU representations. Let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be a decision environment in a category $\boldsymbol{\mathcal{C}}$, and let $\underline{\rhd}^{\text{xa}}$ be an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$. A *global subjective expected utility* (SEU) *representation* for $\underline{\rhd}^{\text{xa}}$ consists of:

- A utility frame $L : \boldsymbol{\mathcal{C}}^{\text{op}} \longrightarrow \mathsf{UPOVS}$;
- A belief system $(\rho_{\mathcal{S}})_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$; and
- A utility system given by $(U^{\mathcal{C}}_{\mathcal{X}})^{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ and $A : \boldsymbol{\mathcal{X}} \longrightarrow \mathsf{Aff}$;

such that for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, the pair $(\rho_{\mathcal{S}}, U_{\mathcal{X}})$ is a local SEU representation (3) for $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$.

Example 6.12. (a) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Set}$, and let $L : \mathsf{Set}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ be as in Example 6.3(a). Let $\{\mu_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a collection of finitely additive probability measures defining a belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$, as in Proposition 6.8(a). Let $A : \boldsymbol{\mathcal{X}} \longrightarrow \mathsf{Aff}$ be a functor. For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded function, and define $U_{\mathcal{X}} : \overrightarrow{\mathsf{Set}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(a). For all $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, suppose that $u_{\mathcal{Y}} \circ \phi = \widehat{\phi} \circ u_{\mathcal{X}}$, where $\widehat{\phi} := A(\phi)$. Define the utility system $(U^{\mathcal{C}}_{\mathcal{X}})^{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ as in Example 6.10.

For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, define a weak order $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ on $\overrightarrow{\mathsf{Set}}(\mathcal{S}, \mathcal{X})$ by formula (3). Then the system $\underline{\rhd}^{\text{xa}} = (\succcurlyeq^{\mathcal{S}}_{\mathcal{X}})^{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ is an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$, and the data L, $(\rho_{\mathcal{S}})_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$, $(U^{\mathcal{C}}_{\mathcal{X}})^{\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ}_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ}$ and A together determine a global SEU representation for $\underline{\rhd}^{\text{xa}}$.

(b) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Meas}$, and let $L : \mathsf{Meas}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ be as in Example 6.3(b). Let $\{\mu_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a collection of finitely additive probability measures defining a belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$, as in Proposition 6.8(b). For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded measurable function, and define $U_{\mathcal{X}} : \overrightarrow{\mathsf{Meas}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(b). Now proceed as in case (a) to get an ex ante preference structure with a global SEU representation.

(c) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Top}$, and let $L : \mathsf{Top}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ be as in Example 6.3(c). Let $\{\mu_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a set of finitely additive Borel probability measures defining a belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$, as in Proposition 6.8(c). For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded continuous function, and define $U_{\mathcal{X}} : \overrightarrow{\mathsf{Top}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(c). Now proceed as in case (a).

(d) Suppose $\boldsymbol{\mathcal{C}} = \mathsf{Metr}$, and let $L : \mathsf{Metr}^{\text{op}} \Longrightarrow \mathsf{UPOVS}$ be as in Example 6.3(d). Let $\{\mu_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ be a collection of Borel probability measures defining a belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$, as in Proposition 6.8(d). For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded locally Lipschitz function, and define $U_{\mathcal{X}} : \overrightarrow{\mathsf{Metr}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(d). Now proceed as in (a).

[20] $\widehat{A}$ is much more complicated than A. For example, $\widehat{A}(\iota^1_{12}) = \phi$ where $\phi(r) = r/5 + 7$.

(e) Suppose $\mathcal{C} = \mathsf{Diff}$, and let $L : \mathsf{Diff}^{\rm op} \Longmapsto \mathsf{UPOVS}$ be as in Example 6.3(e). Let $\{\mu_{\mathcal{S}}\}_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}$ be a collection of tempered distributions defining a belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}$, as in the remark following Proposition 6.8. For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ be a bounded smooth function, and define $U_{\mathcal{X}} : \overrightarrow{\mathsf{Diff}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$ as in Example 6.4(e). Now proceed as in case (a). ♢

If $(L, (\rho_{\mathcal{S}})_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}, U, A)$ is a global SEU representation for $\underline{\triangleright}^{\rm xa}$, and $(\widehat{U}, \widehat{A})$ is another utility system on $\boldsymbol{\mathcal{X}}$ that is cardinally equivalent to (U, A), then it is easily verified that $(L, (\rho_{\mathcal{S}})_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}, \widehat{U}, \widehat{A})$ is also global SEU representation for $\underline{\triangleright}^{\rm xa}$. Thus, global SEU representations can only be unique up to cardinal equivalence of their utility systems.

Our goal now is to reverse the logic of Example 6.12. Instead of *stipulating* all the pieces of a global SEU representation and using these to construct an ex ante preference structure, suppose that we *begin* with an ex ante preference structure $\underline{\triangleright}^{\rm xa}$ on a decision environment $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$. Under what conditions does $\underline{\triangleright}^{\rm xa}$ admit a global SEU representation?

7 Existence of global SEU in concrete categories

The main result of this paper yields a global SEU representation under broad conditions. In this section, we formulate this result for the special case of concrete categories, while in Section 8, we shall state a more general version that also applies to abstract categories. But before stating the theorem, we need a bit more machinery.

7A Concrete categories. Throughout this section, $\mathcal{C}$ denotes a concrete category; in other words, the objects in $\mathcal{C}$ are sets, and the morphisms are functions between these sets. For consistency with the framework developed in Section 8, we shall use the following notation. For any object $\mathcal{C}$ in $\mathcal{C}^\circ$, let $\underline{\mathcal{C}}$ denote the underlying set, with $\underline{a}, \underline{b}, \underline{c}$ etc. denoting generic elements. For any objects $\mathcal{B}, \mathcal{C} \in \mathcal{C}^\circ$ and any morphism $\phi \in \overrightarrow{\mathcal{C}}(\mathcal{B}, \mathcal{C})$, let $\underline{\phi} : \underline{\mathcal{B}} \longrightarrow \underline{\mathcal{C}}$ denote the underlying function. We shall also make the following assumption:

(CM) For all $\mathcal{B}, \mathcal{C} \in \mathcal{C}^\circ$ and all $\underline{c} \in \underline{\mathcal{C}}$, there is a *constant* morphism $\phi \in \overrightarrow{\mathcal{C}}(\mathcal{B}, \mathcal{C})$ such that $\underline{\phi}(\underline{b}) = \underline{c}$ for all $\underline{b} \in \underline{\mathcal{B}}$.

This assumption is satisfied by Set, Meas, Top, Metr, Diff, and many concrete categories. (But it is not satisfied by Vec, UPOVS, and other categories of "algebraic" objects.)

7B The canonical frame. For all $\mathcal{C} \in \mathcal{C}^\circ$, let $L(\mathcal{C}) := \ell^\infty(\underline{\mathcal{C}})$. This is a unitary, partially ordered vector space (in fact, it is a unitary, Archimedean Riesz space). For all $\mathcal{A}, \mathcal{B} \in \mathcal{C}^\circ$ and $\phi \in \overrightarrow{\mathcal{C}}(\mathcal{A}, \mathcal{B})$, we define $L(\phi) := \phi^* : L(\mathcal{B}) \longrightarrow L(\mathcal{A})$ by setting $\phi^*(v) := v \circ \underline{\phi}$ for all $v \in \ell^\infty(\underline{\mathcal{B}})$. (Recall that $\underline{\phi} : \underline{\mathcal{A}} \longrightarrow \underline{\mathcal{B}}$, and $v : \underline{\mathcal{B}} \longrightarrow \mathbb{R}$, so $v \circ \underline{\phi} : \underline{\mathcal{A}} \longrightarrow \mathbb{R}$.) It is easily verified that L is a utility frame on $\mathcal{C}$. We shall call this the *canonical frame* on $\mathcal{C}$.

7C Connected categories. Let $\mathcal{C}$ be a category. We define a binary relation $\sim$ on $\mathcal{C}^\circ$ as follows: for any $\mathcal{X}, \mathcal{Y} \in \mathcal{C}^\circ$, write $\mathcal{X} \sim \mathcal{Y}$ if either $\overrightarrow{\mathcal{C}}(\mathcal{X}, \mathcal{Y}) \neq \emptyset$ or $\overrightarrow{\mathcal{C}}(\mathcal{Y}, \mathcal{X}) \neq \emptyset$. The structure $(\mathcal{C}^\circ, \sim)$ is an undirected graph. The category $\mathcal{C}$ is *connected* if this graph is connected. For example, the two categories shown in Figure 1 are connected.

7D Ex post preferences. Let $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ be a decision environment in $\boldsymbol{\mathcal{C}}$, and let $\trianglerighteq^{\mathrm{xa}}$ be an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$. Fix $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. For all $\underline{x} \in \underline{\mathcal{X}}$, (CM) yields a (unique) constant morphism $\kappa^{\mathcal{S}}_{\underline{x}} \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X})$ with the value $\underline{x}$. We define a weak order $\succsim^{\mathrm{xp}}_{\mathcal{S},\mathcal{X}}$ on $\underline{\mathcal{X}}$ by stipulating, for all $\underline{x}, \underline{y} \in \underline{\mathcal{X}}$, that $\underline{x} \succsim^{\mathrm{xp}}_{\mathcal{S},\mathcal{X}} \underline{y}$ if and only if $\kappa^{\mathcal{S}}_{\underline{x}} \succsim^{\mathcal{S}}_{\mathcal{X}} \kappa^{\mathcal{S}}_{\underline{y}}$.

Lemma 7.1 *If the category $\boldsymbol{\mathcal{S}}$ is connected, then $\succsim^{\mathrm{xp}}_{\mathcal{S},\mathcal{X}}$ is independent of $\mathcal{S}$.*

In light of Lemma 7.1, there is a weak order $\succsim^{\mathrm{xp}}_{\mathcal{X}}$ on $\underline{\mathcal{X}}$, such that for any $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and any $\underline{x}, \underline{y} \in \underline{\mathcal{X}}$, we have $\kappa^{\mathcal{S}}_{\underline{x}} \succsim^{\mathcal{S}}_{\mathcal{X}} \kappa^{\mathcal{S}}_{\underline{y}}$ if and only if $\underline{x} \succsim^{\mathrm{xp}}_{\mathcal{X}} \underline{y}$. Following the standard terminology in decision theory, we shall refer to $\succsim^{\mathrm{xp}}_{\mathcal{X}}$ as the *ex post preference order* on $\mathcal{X}$.

7E Split epimorphisms. Let $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$. A morphism $\pi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ is called a *split epimorphism* (or a *retraction*) if it has a right-inverse. In other words, there is a morphism $\sigma \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{B})$ such that $\pi \circ \sigma = I_{\mathcal{C}}$. This implies that, for any other $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$, the function $\vec{\pi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$ is surjective. In other words: for any $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$ there is some $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ making the following diagram commute:

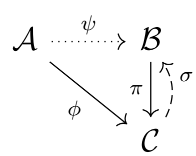
(10)

(The construction is easy: just define $\psi = \sigma \circ \phi$, as suggested by diagram (10).) In fact, it is easily shown that π is a split epimorphism if and only if the function $\vec{\pi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$ is surjective for all $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$.

For example, in the category Set, *every* surjective function is a split epimorphism.[21] For another example, suppose $\boldsymbol{\mathcal{C}}$ is a concrete category where we can define Cartesian products of objects in some natural way (e.g. Meas, Top, Metr, or Diff). Suppose $\mathcal{B} = \mathcal{A} \times \mathcal{C}$ for some other object $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$, and let $\pi : \mathcal{B} \longrightarrow \mathcal{C}$ be the projection onto the second coordinate. Fix $a \in \mathcal{A}$ and define $\sigma : \mathcal{C} \longrightarrow \mathcal{B}$ by $\sigma(c) := (a, c)$ for all $c \in \mathcal{C}$. (For the sake of this example, suppose that σ is a morphism in $\mathcal{C}$.) Then $\pi \circ \sigma = I_{\mathcal{C}}$, so π is a split epimorphism.[22]

7F Weakly directed subcategories. Let $(\mathcal{D}, \rhd)$ be a partially ordered set. Recall that $\mathcal{D}$ is a *directed set* if for any $c, d \in \mathcal{D}$, there is some $e \in \mathcal{D}$ such that $c \lhd e$ and $d \lhd e$. (For example, any linearly ordered set is directed. So is any lattice.)

Let $\boldsymbol{\mathcal{X}}$ be a subcategory of $\boldsymbol{\mathcal{C}}$. We define a partial order $\rhd$ on $\boldsymbol{\mathcal{X}}^\circ$ as follows. First, for any $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, write $\mathcal{X} \rightsquigarrow \mathcal{Y}$ if *either* $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) \neq \emptyset$, *or* there is some $\pi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{Y}, \mathcal{X})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$. (Note that we do *not* require π to be a split epimorphism in $\boldsymbol{\mathcal{X}}$.) Now let $\rhd$ be the transitive closure of $\rightsquigarrow$. We shall say that the subcategory $\boldsymbol{\mathcal{X}}$ is *weakly directed* if $(\boldsymbol{\mathcal{X}}, \rhd)$ is a directed set. For example:

[21] This assumes the Axiom of Choice.

[22] As these examples suggest, in concrete categories, split epimorphisms are surjective.

- Suppose that $\boldsymbol{\mathcal{X}}$ is a *directed category*, in the sense that, for any $\mathcal{X},\mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ there is some $\mathcal{Z} \in \boldsymbol{\mathcal{X}}^\circ$ such that $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X},\mathcal{Z}) \neq \emptyset$ and $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{Y},\mathcal{Z}) \neq \emptyset$. Then $\boldsymbol{\mathcal{X}}$ is weakly directed. (For example: any subcategory with a terminal object is weakly directed. Any subcategory with finite coproducts is weakly directed.)
- Suppose that $\boldsymbol{\mathcal{X}}$ is a connected category, and that for all $\mathcal{X},\mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, there is some $\pi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{Y},\mathcal{X})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$. Then $\boldsymbol{\mathcal{X}}$ is weakly directed.

For instance, the categories from Example 3.1(b,c) (shown in Figure 1) are not weakly directed. But the categories from Example 4.2(b,c) (shown in Figure 2) *are* weakly directed.

7G Structural assumptions. Let $(\boldsymbol{\mathcal{S}},\boldsymbol{\mathcal{X}})$ be a decision environment in a category $\boldsymbol{\mathcal{C}}$. In addition to (CM), our main result requires two other structural assumptions:

(SC) $\boldsymbol{\mathcal{S}}$ is connected.

(XD) $\boldsymbol{\mathcal{X}}$ is weakly directed.

As the earlier examples show, these are mild assumptions, satisfied in most applications.

7H Probabilistic extensions of categories. For all $\mathcal{A},\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, let $\overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$ be the convex space of all finite-support probability measures over $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A},\mathcal{B})$. A typical element will be indicated $\widetilde{\phi} = (\phi_1,p_1;\ \phi_2,p_2;\ldots;\ \phi_N,p_N)$, where $\phi_1,\phi_2,\ldots,\phi_N \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A},\mathcal{B})$ and $(p_1,p_2,\ldots,p_N) \in \mathbb{R}_+^N$ is a vector of non-negative real numbers summing to 1. We shall denote elements of $\overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$ with tildes, as in $\widetilde{\phi},\widetilde{\psi}$, etc. For every $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A},\mathcal{B})$, we shall abuse notation by identifying ϕ with the element of $\overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$ that assigns probability 1 to ϕ. For any $\mathcal{A},\mathcal{B},\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\widetilde{\phi} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$ and $\widetilde{\psi} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{B},\mathcal{C})$, if $\widetilde{\phi} = (\phi_1,p_1;\ \phi_2,p_2;\ldots;\ \phi_N,p_N)$ and $\widetilde{\psi} = (\psi_1,q_1;\ \psi_2,q_2;\ldots;\ \psi_M,q_M)$, then $\widetilde{\psi}\circ\widetilde{\phi} = \widetilde{\xi} := (\xi_1,r_1;\ \xi_2,r_2;\ \ldots;\ \xi_L,r_L) \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{C})$, where $\{\xi_1,\xi_2,\ldots,\xi_L\} := \{\psi_m\circ\psi_n;\ n\in[1\ldots N]$ and $m\in[1\ldots M]\}$ and where, for all $\ell \in [1\ldots L]$, we define

$$r_\ell \quad := \quad \sum\{q_m\cdot p_n\ ;\ n\in[1\ldots N],\ m\in[1\ldots M],\ \text{and}\ \psi_m\circ\psi_n = \xi_\ell\}.\text{[23]} \tag{11}$$

The result is a new category, $\Delta\boldsymbol{\mathcal{C}}$, with the same objects as $\boldsymbol{\mathcal{C}}$, but with more morphisms between them; we shall call it the *probabilistic extension* of $\boldsymbol{\mathcal{C}}$. Note that $\boldsymbol{\mathcal{C}}$ embeds as a subcategory of $\Delta\boldsymbol{\mathcal{C}}$.

An important feature of $\Delta\boldsymbol{\mathcal{C}}$ is that for all $\mathcal{A},\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$ the set $\overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$ has a natural convex structure. Formally, let $\widetilde{\phi},\widetilde{\phi}' \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{A},\mathcal{B})$. Suppose that if $\widetilde{\phi} = (\phi_1,p_1;\ \ldots;\ \phi_N,p_N)$ and $\widetilde{\phi}' = (\phi_1,p_1';\ \ldots;\ \phi_N,p_N')$ for some $\phi_1,\ldots,\phi_N \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A},\mathcal{B})$ and probability vectors $(p_1,\ldots,p_N)$ and $(p_1',\ldots,p_N')$. (Here we allow that $p_n = 0$ or $p_n' = 0$ for some $n \in [1\ldots N]$.) For any $q \in [0,1]$, we define $\widetilde{\phi} \circledast_q \widetilde{\phi}' := (\phi_1,r_1;\ \ldots;\ \phi_N,r_N)$, where, for all $n \in [1\ldots N]$, $r_n :=\ q\,p_n + (1-q)\,p_n'$.

[23] If $\psi_m\circ\psi_n \neq \psi_{m'}\circ\psi_{n'}$ whenever $(m,n)\neq(m',n')$, then $\widetilde{\xi} = (\psi_1\circ\phi_1,q_1p_1;\ \psi_1\circ\phi_2,q_1p_2;\ \ldots;\ \psi_1\circ\phi_N,q_1p_N;\ \psi_2\circ\phi_1,q_2p_1;\ \psi_2\circ\phi_2,q_2p_2;\ \ldots;\ \psi_2\circ\phi_N,q_1p_N;\ \ldots\ldots;\ \psi_M\circ\phi_1,q_Mp_1;\ \psi_M\circ\phi_2,q_Mp_2;\ \ldots;\ \psi_M\circ\phi_N,q_Mp_N)$. But in general, it may be that $\psi_m\circ\psi_n = \psi_{m'}\circ\psi_{n'}$ for some $(m,n)\neq(m',n')$, which is why we use the more complex expression (11).

7I Axioms and main result. Let $(\boldsymbol{\mathcal{S}},\boldsymbol{\mathcal{X}})$ be a decision environment in $\boldsymbol{\mathcal{C}}$, and let $\trianglerighteq^{\rm xa}$ be an ex ante preference structure on $(\boldsymbol{\mathcal{S}},\boldsymbol{\mathcal{X}})$, regarded as a pair of subcategories of $\Delta\boldsymbol{\mathcal{C}}$. Thus, for all $\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X}\in\boldsymbol{\mathcal{X}}^\circ$, there is a preference order $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ on $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$. We require $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ to satisfy the following axioms, based on those of Anscombe and Aumann (1963):

vNM independence. For all $\widetilde{\alpha},\widetilde{\beta},\widetilde{\gamma}\in\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$, if $\widetilde{\alpha}\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}\widetilde{\beta}$, then $\widetilde{\alpha}\circledast_q\widetilde{\gamma}\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}\widetilde{\beta}\circledast_q\widetilde{\gamma}$ for all $q\in[0,1]$.

Continuity. For all $\widetilde{\alpha},\widetilde{\beta},\widetilde{\gamma}\in\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$, the sets $\{q\in[0,1];\ \widetilde{\alpha}\circledast_q\widetilde{\beta}\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}\widetilde{\gamma}\}$ and $\{q\in[0,1];\ \widetilde{\alpha}\circledast_q\widetilde{\beta}\preccurlyeq^{\mathcal{S}}_{\mathcal{X}}\widetilde{\gamma}\}$ are closed in $[0,1]$.

Statewise dominance. For all $\widetilde{\alpha},\widetilde{\beta}\in\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$, if $\underline{\widetilde{\alpha}}(\underline{s})\succcurlyeq^{\rm xp}_{\mathcal{X}}\underline{\widetilde{\beta}}(\underline{s})$ for all $\underline{s}\in\underline{\mathcal{S}}$, then $\widetilde{\alpha}\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}\widetilde{\beta}$.

Boundedness. For all $\mathcal{X}\in\boldsymbol{\mathcal{X}}^\circ$, the set $\underline{\mathcal{X}}$ contains $\succcurlyeq^{\rm xp}_{\mathcal{X}}$-maximal and $\succcurlyeq^{\rm xp}_{\mathcal{X}}$-minimal elements.

Here is the first version of the main result.

Theorem 7.1 *Suppose a concrete category* $\boldsymbol{\mathcal{C}}$ *and decision environment* $(\boldsymbol{\mathcal{S}},\boldsymbol{\mathcal{X}})$ *satisfy structural assumptions* (CM), (SC)*, and* (XD)*. Let* $\trianglerighteq^{\rm xa}$ *be an ex ante preference structure on* $(\boldsymbol{\mathcal{S}},\boldsymbol{\mathcal{X}})$ *in the category* $\Delta\boldsymbol{\mathcal{C}}$*. If* $\trianglerighteq^{\rm xa}$ *satisfies* vNM independence*,* Continuity*,* Statewise Dominance*, and* Boundedness*, then it admits a global SEU representation.*

In this representation, the utility frame L *is a utility subframe of the canonical frame, the belief system* $\{\rho_{\mathcal{S}}\}_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}$ *is unique, and the utility system* (U,A) *is unique up to cardinal equivalence. The affinity functor* A *maps every* $\boldsymbol{\mathcal{X}}$*-morphism to an affine contraction.*

In Theorem 7.1, we need structural assumption (SC) so that the ex post preference orders $\succcurlyeq^{\rm xp}_{\mathcal{X}}$ are well-defined for all $\mathcal{X}\in\boldsymbol{\mathcal{X}}^\circ$ via Lemma 7.1 (so that Statewise Dominance is meaningful). The next example shows why we need (XD).

Example 7.2. Suppose $\boldsymbol{\mathcal{C}}$ is a subcategory of Set. Let $\boldsymbol{\mathcal{S}}^\circ=\{\mathcal{S}\}$ where $\mathcal{S}=\{1,2\}\times\{1,2,3\}$ and the only $\boldsymbol{\mathcal{S}}$-morphism is the identity function on $\mathcal{S}$. Meanwhile, $\boldsymbol{\mathcal{X}}^\circ=\{\mathcal{X},\mathcal{Y},\mathcal{Z}\}$, where $\mathcal{X}=\{x_1,x_2\}$, $\mathcal{Y}=\{y_1,y_2\}$, and $\mathcal{Z}=\{z_1,z_2\}$. Suppose $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X},\mathcal{Y})=\{\gamma\}$ where $\gamma(x_1)=y_1$ and $\gamma(x_2)=y_2$. Suppose $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X},\mathcal{Z})=\{\zeta\}$ where $\zeta(x_1)=z_1$ and $\zeta(x_2)=y_2$. Aside from these morphisms and the identity maps, suppose $\boldsymbol{\mathcal{X}}$ has no other morphisms.

Suppose that $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{Y})$ and $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{Z})$ each contain all $2^6=64$ possible functions. However, $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{X})$ contains only the four possible functions that are constant in the second coordinate. In other words, $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{X})=\{\alpha_1,\alpha_2,\alpha_3,\alpha_4\}$, where

$$\alpha_1=\begin{array}{|c|c|}\hline x_1&x_1\\\hline x_1&x_1\\\hline x_1&x_1\\\hline\end{array}\qquad \alpha_2=\begin{array}{|c|c|}\hline x_1&x_2\\\hline x_1&x_2\\\hline x_1&x_2\\\hline\end{array}\qquad \alpha_3=\begin{array}{|c|c|}\hline x_2&x_1\\\hline x_2&x_1\\\hline x_2&x_1\\\hline\end{array}\quad\text{and}\quad \alpha_4=\begin{array}{|c|c|}\hline x_2&x_2\\\hline x_2&x_2\\\hline x_2&x_2\\\hline\end{array}$$

(Here we represent functions on $\mathcal{S}$ using 2×3 arrays in the obvious way.) Thus, the functions $\vec{\gamma}:\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{X})\longrightarrow\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{Y})$ and $\vec{\zeta}:\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{X})\longrightarrow\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{S},\mathcal{Z})$ are *not* surjective, which means that γ and ζ are *not* split epimorphisms. The morphism digraph of $\boldsymbol{\mathcal{X}}$ has the structure $\mathcal{Y} \xleftarrow{\gamma} \mathcal{X} \xrightarrow{\zeta} \mathcal{Z}$, so $\boldsymbol{\mathcal{X}}$ is *not* weakly directed. So (XD) is not satisfied.

Let $L(\mathcal{S}) := \mathbb{R}^{\mathcal{S}}$ with the obvious partial order, as in Example 6.3(a), let $u_{\mathcal{X}}(x_1) = u_{\mathcal{Y}}(y_1) = u_{\mathcal{Z}}(z_1) = 1$ and $u_{\mathcal{X}}(x_2) = u_{\mathcal{Y}}(y_2) = u_{\mathcal{Z}}(z_2) = 2$, and define $U^{\mathcal{S}}_{\mathcal{X}} : \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{X}) \longrightarrow \mathbb{R}^{\mathcal{S}}$, $U^{\mathcal{S}}_{\mathcal{Y}} : \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{Y}) \longrightarrow \mathbb{R}^{\mathcal{S}}$, and $U^{\mathcal{S}}_{\mathcal{Y}} : \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}, \mathcal{Z}) \longrightarrow \mathbb{R}^{\mathcal{S}}$ as in Example 6.4(a). Thus, if $A(\gamma)$ and $A(\zeta)$ are both the identity map, then (U, A) is a utility system.

Let μ and μ' be the following probability measures on $\mathcal{S}$:

μ :=

1/6	1/6
1/6	1/6
1/6	1/6

and μ' :=

1/3	1/3
1/12	1/12
1/12	1/12

Let $\rho, \rho' : \mathbb{R}^{\mathcal{S}} \longrightarrow \mathbb{R}$ be the corresponding expectation operators.

Let $\underline{\trianglerighteq}^{\text{xa}} = \{\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}, \succcurlyeq^{\mathcal{S}}_{\mathcal{Y}}, \succcurlyeq^{\mathcal{S}}_{\mathcal{Z}}\}$. Suppose that $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ has a local SEU representation (3) given by $u_{\mathcal{X}}$ and ρ, and $\succcurlyeq^{\mathcal{S}}_{\mathcal{Y}}$ has a local SEU representation given by $u_{\mathcal{Y}}$ and ρ, while $\succcurlyeq^{\mathcal{S}}_{\mathcal{Z}}$ has a local SEU representation given by $u_{\mathcal{Z}}$ and ρ'. Then the functions γ and ζ satisfy the property (TP), because ρ and ρ' both have the same marginal $(\frac{1}{2}, \frac{1}{2})$ on the first coordinate. However, there is no single probability measure on $\mathcal{S}$ (hence, no linear functional on $\mathbb{R}^{\mathcal{S}}$) which can provide a local SEU representation for all of $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$, $\succcurlyeq^{\mathcal{S}}_{\mathcal{Y}}$ and $\succcurlyeq^{\mathcal{S}}_{\mathcal{Z}}$. Thus, $\underline{\trianglerighteq}^{\text{xa}}$ does not have a global SEU representation. ♢

A shortcoming of Theorem 7.1 is that the resulting SEU representation is not necessarily "well-adapted" to the category $\boldsymbol{\mathcal{C}}$. For example, in the category Top, we would like an SEU representation of the kind described in Example 6.12(c). That is: L should be the utility frame of bounded *continuous* functions from Example 6.3(c). For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^{\circ}$, the utility functional $U_{\mathcal{X}}$ should arise from a bounded *continuous* utility function $u_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$, as in Example 6.4(c). And for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, the belief $\rho_{\mathcal{S}}$ should arise from a *Borel* probability measure, as in Proposition 6.8(c). Theorem 7.1 does not guarantee an SEU representation with these properties. But we can ensure them with three further axioms:

Ex post Continuity. For all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^{\circ}$, the ex post preference order $\succcurlyeq^{\text{xp}}_{\mathcal{X}}$ is continuous. That is: for all $\underline{y} \in \underline{\mathcal{X}}$, the sets $\{\underline{x} \in \underline{\mathcal{X}};\ \underline{x} \succcurlyeq^{\text{xp}}_{\mathcal{X}} \underline{y}\}$ and $\{\underline{x} \in \underline{\mathcal{X}};\ \underline{x} \preccurlyeq^{\text{xp}}_{\mathcal{X}} \underline{y}\}$ are closed in $\mathcal{X}$.

Certainty Equivalents. For every $\underline{\widetilde{y}} \in \underline{\Delta\mathcal{X}}$, there is some $\underline{x} \in \underline{\mathcal{X}}$ such that $\underline{\widetilde{y}} \approx^{\text{xp}}_{\mathcal{X}} \underline{x}$.

Monotone Path Property. There is some $\mathcal{X}_* \in \boldsymbol{\mathcal{X}}^{\circ}$ admitting a *monotone path*: a continuous function $\gamma : [0, 1] \longrightarrow \mathcal{X}_*$ such that $\gamma(s) \prec^{\mathcal{X}_*}_{\text{xp}} \gamma(t)$ for all $s < t$.

Proposition 7.3 *Let* $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ *be a decision environment in* Top *satisfying conditions* (SC) *and* (XD)*, such that every space in* $\boldsymbol{\mathcal{S}}^{\circ}$ *is normal and Hausdorff, while every space in* $\boldsymbol{\mathcal{X}}^{\circ}$ *is second-countable and path-connected. Let* $\underline{\trianglerighteq}^{\text{xa}}$ *be an ex ante preference structure on* $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ *in the category* ΔTop*. If* $\underline{\trianglerighteq}^{\text{xa}}$ *satisfies* vNM independence*,* Continuity*,* Statewise Dominance*,* Boundedness*,* Ex post Continuity*,* Certainty Equivalents *and the* Monotone Path Property*, then it has a global SEU representation of the kind described in Example* 6.12(c)*.*

To understand why something like the Monotone Path Property is needed, suppose that $\mathcal{S}$ is a connected topological space, while $\mathcal{X}$ is totally disconnected. Then $\overrightarrow{\mathsf{Top}}(\mathcal{S},\mathcal{X})$ contains only constant functions, so the ex ante preference $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$ would simply recapitulate the ex post preference $\succcurlyeq^{\mathcal{X}}_{\mathrm{xp}}$. If this was true for every $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, then we would not have enough information to uniquely determine $\rho_{\mathcal{S}}$ on $\mathfrak{C}_{\mathrm{b}}(\mathcal{S})$.

More generally, if $\overrightarrow{\mathsf{Top}}(\mathcal{S},\mathcal{X})$ is "small", then its image under the utility functional $U_{\mathcal{X}}$ will be a "small" subspace of $\mathfrak{C}_{\mathrm{b}}(\mathcal{S})$. If this holds for all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, then $\rho_{\mathcal{S}}$ is not determined on a large enough subspace of $\mathfrak{C}_{\mathrm{b}}(\mathcal{S})$ to identify a unique Borel measure. The Monotone Path Property solves this problem.

If $\mathcal{X}$ is a connected topological space, then the conjunction of Ex post Continuity and Certainty Equivalents is equivalent to a single axiom:

Strong ex post Continuity. For every $\underline{\widetilde{y}} \in \underline{\Delta\mathcal{X}}$, the sets $\{\underline{x} \in \underline{\mathcal{X}};\ \ \underline{x} \succcurlyeq^{\mathrm{xp}}_{\mathcal{X}} \underline{\widetilde{y}}\}$ and $\{\underline{x} \in \underline{\mathcal{X}};\ \underline{x} \preccurlyeq^{\mathrm{xp}}_{\mathcal{X}} \underline{\widetilde{y}}\}$ are closed in $\mathcal{X}$.

Since all objects in $\boldsymbol{\mathcal{X}}^\circ$ are connected, we can reformulate Proposition 7.3 by replacing Ex post Continuity and Certainty Equivalents with this single axiom. It is currently unknown whether there are results analogous to Proposition 7.3 for other concrete categories, such as Meas, Metr, or Diff. This is an interesting problem for future research.

8 Existence of global SEU in abstract categories

As noted in Section 1, it is sometimes implausible that a decision problem comes with an obvious or prespecified space of possible "states of nature", or menu of possible "outcomes". This has inspired several models of decision-making which explicitly eschew these ingredients (Skiadas, 1997a,b; Gilboa and Schmeidler, 2004; Karni, 2006; Ahn, 2008; Blume et al., 2021). An agent might be confronted by various sources of uncertainty that she cannot model by specifying state spaces. In each decision problem, she can choose amongst feasible courses of action, and understands that, contingent on the resolution of the uncertainty, these actions may yield various consequences. But she may be unable to specify a complete outcome space for each decision problem. Thus, she lacks the resources to model her decision problems in the standard Savage framework.

Unless we impose some further structure, it is obviously impossible to say anything meaningful about this case. But suppose that the agent can conceptualize the various uncertainty sources and outcome menus as objects in an *abstract category* $\boldsymbol{\mathcal{C}}$, and she can conceptualize the various courses of action as morphisms in $\boldsymbol{\mathcal{C}}$.[24] We shall now show that, if $\boldsymbol{\mathcal{C}}$ has sufficient structure, then the agent's preferences admit a global SEU representation.

8A Initial and terminal objects. Let $\boldsymbol{\mathcal{C}}$ be a category. An object $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$ is *initial* if, for all other objects $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, there is a *unique* morphism $\mathcal{A} \longrightarrow \mathcal{C}$ (usually denoted by the symbol ¡). If $\boldsymbol{\mathcal{C}}$ has an initial object, then it is unique up to isomorphism. So we normally refer to "the" initial object in $\boldsymbol{\mathcal{C}}$; it is generally denoted by $\mathbf{0}_{\boldsymbol{\mathcal{C}}}$. For example, in Set, Meas,

[24] See Appendix C for some examples of abstract categories.

and Top, the initial object is the empty set. Meanwhile, in UPOVS, the initial object is $\mathbb{R}$: for any other unitary POVS $\mathcal{V}$, there is a unique order-preserving linear transformation $\mathbb{R} \xrightarrow{\mathsf{i}} \mathcal{V}$ that maps 1 to the order unit of $\mathcal{V}$.

An object $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$ is *terminal* if, for all other objects $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, there is a *unique* morphism $\mathcal{C} \longrightarrow \mathcal{A}$ (usually denoted by the symbol !). If $\boldsymbol{\mathcal{C}}$ has a terminal object, then it is unique up to isomorphism. So we normally refer to "the" terminal object in $\boldsymbol{\mathcal{C}}$; it is generally denoted by $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$. For example, in Set, Meas, Top, and Diff, the terminal object is the singleton set (regarded as a "one-point space" or a "zero-dimensional manifold", as appropriate).

In many categories of "algebraic" structures (e.g. Vec), the initial and terminal objects are the same; typically this is a "trivial" algebraic structure which contains only the identity element. This is called a *zero object*. For example, in Vec, the zero object is the zero-dimensional vector space $\{0\}$.

8B Global elements. Let $\boldsymbol{\mathcal{C}}$ be a category with a terminal object $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$. Let $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$. A *global element* of $\mathcal{C}$ is a morphism $\phi : \mathbf{1}_{\boldsymbol{\mathcal{C}}} \longrightarrow \mathcal{C}$. Let $\underline{\mathcal{C}} := \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \mathcal{C})$ be the set of global elements of $\mathcal{C}$; we shall denote generic elements by $\underline{a}, \underline{b}, \underline{c}$, etc. For any $\mathcal{C}_1, \mathcal{C}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_1, \mathcal{C}_2)$, we define the function $\underline{\phi} : \underline{\mathcal{C}}_1 \longrightarrow \underline{\mathcal{C}}_2$ by setting $\underline{\phi}(\underline{c}) := \phi \circ \underline{c}$ for all $\underline{c} \in \underline{\mathcal{C}}_1$ (recalling that $\underline{c}$ is itself a morphism from $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$ to $\mathcal{C}_1$). This transformation (sending each object $\mathcal{C}$ to its set $\underline{\mathcal{C}}$ of global elements, and mapping each morphism ϕ to the corresponding function $\underline{\phi}$) is a functor from $\boldsymbol{\mathcal{C}}$ into Set.[25] In many concrete categories (e.g. Set, Meas, Top, Diff, etc.), it is naturally isomorphic to the forgetful functor. In other words: there is a natural bijection between the elements of $\underline{\mathcal{C}}$ and the points of the set underlying $\mathcal{C}$, and the function $\underline{\phi}$ describes the way that the morphism ϕ maps these points. (This justifies using the notation $\underline{\mathcal{C}}$ to refer to both sets.)

Unfortunately, not all categories of interest have terminal objects. So we shall consider a slightly more general construction, which subsumes global elements.

8C Strongly connected categories. The category $\boldsymbol{\mathcal{C}}$ is *strongly connected*, if $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathcal{Y}) \neq \emptyset$ for all non-initial $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{C}}^\circ$.[26] For example: if $\boldsymbol{\mathcal{C}}$ has a terminal object, and $\underline{\mathcal{C}} \neq \emptyset$ for every non-initial object $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, then $\boldsymbol{\mathcal{C}}$ is strongly connected.

8D Constant morphisms. Let $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$. A morphism $\kappa \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ is *constant* if, for all $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$ and all $\phi, \psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$, we have $\kappa \circ \phi = \kappa \circ \psi$. For example, if $\boldsymbol{\mathcal{C}}$ has a terminal object $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$, then any global element $\kappa \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \mathcal{C})$ is a constant morphism. More generally, a morphism $\kappa \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ is constant if there is some global element $\varphi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \mathcal{C})$ such that

[25] Indeed, it is the covariant hom functor $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \bullet)$ from §5C.

[26] We exclude the initial object $\mathbf{0}_{\boldsymbol{\mathcal{C}}}$ even though $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{0}_{\boldsymbol{\mathcal{C}}}, \mathcal{Y}) \neq \emptyset$ for all $\mathcal{Y} \in \boldsymbol{\mathcal{C}}^\circ$, because in many concrete categories $\mathbf{0}_{\boldsymbol{\mathcal{C}}}$ is the empty set, so that $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{X}, \mathbf{0}_{\boldsymbol{\mathcal{C}}}) = \emptyset$ for all (non-initial) $\mathcal{X} \in \boldsymbol{\mathcal{C}}^\circ$.

the following diagram commutes

$$\begin{array}{ccc} \mathcal{B} & & \\ {\scriptstyle !}\downarrow & \searrow^{\kappa} & \\ \mathbf{1}_{\boldsymbol{\mathcal{C}}} & \xrightarrow[\varphi]{} & \mathcal{C} \end{array}$$

But constant morphisms can exist even in categories without terminal objects. For example: let $\boldsymbol{\mathcal{C}}$ be the category of all nonzero-dimensional differentiable manifolds and smooth maps, or let $\boldsymbol{\mathcal{C}}$ be the category of perfect topological spaces and continuous functions. Neither category contains a terminal object. But if $\mathcal{B}$ and $\mathcal{C}$ are objects in these categories, and $\kappa : \mathcal{B} \longrightarrow \mathcal{C}$ is any constant function, then κ is a constant morphism in the category.

If $\kappa \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ is constant, then for any objects $\mathcal{A}, \mathcal{D} \in \boldsymbol{\mathcal{C}}^\circ$ and morphisms $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\delta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$, the morphisms $\kappa \circ \alpha$ and $\delta \circ \kappa$ are also constant. (See Lemma B.1.)

8E Constituents. Suppose $\boldsymbol{\mathcal{C}}$ is strongly connected. Let $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ be non-initial, and let $\psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ be constant morphisms. Write $\psi \sim \phi$ there is some $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ such that $\psi = \phi \circ \alpha$. This is is an equivalence relation on the set of all constant morphisms into $\mathcal{C}$ (see Lemma B.2). The equivalence classes are called ***constituents*** of $\mathcal{C}$. For example, if $\boldsymbol{\mathcal{C}}$ has a terminal object $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$, then every global element in $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \mathcal{C})$ determines a constituent of $\mathcal{C}$. But objects in $\boldsymbol{\mathcal{C}}$ can have constituents even if $\boldsymbol{\mathcal{C}}$ has no terminal object, and hence no global elements, as shown by the examples in §8D involving differentiable manifolds and perfect topological spaces.

Let $\underline{\mathcal{C}}$ be the set of all constituents of $\mathcal{C}$. Generic elements will be denoted $\underline{a}, \underline{b}, \underline{c}$, etc. If a constant morphism ϕ belongs to the constituent (i.e. equivalence class) $\underline{c}$, then we shall say that ϕ has the ***value*** $\underline{c}$. If $\boldsymbol{\mathcal{C}}$ is strongly connected, then for any $\underline{c} \in \underline{\mathcal{C}}$ and any other object $\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, there is a unique constant morphism $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ with value $\underline{c}$ (Lemma B.3). In particular, if $\boldsymbol{\mathcal{C}}$ has a terminal object $\mathbf{1}_{\boldsymbol{\mathcal{C}}}$, then every constituent of $\underline{\mathcal{C}}$ can be realized by a morphism in $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathbf{1}_{\boldsymbol{\mathcal{C}}}, \mathcal{C})$; thus, there is a bijective correspondence between the constituents of $\mathcal{C}$ and the global elements of $\mathcal{C}$. (This justifies using the notation $\underline{\mathcal{C}}$ to refer to both sets.)

8F Morphism values. Let $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$. For every constant morphism $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{C})$, the composition $\phi \circ \alpha$ is also constant. Furthermore, if $\beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ is also constant, and $\alpha \sim \beta$, then $(\phi \circ \alpha) \sim (\phi \circ \beta)$ (see Lemma B.4). Thus, ϕ maps each constituent of $\mathcal{C}$ to a constituent of $\mathcal{D}$; this yields a function $\underline{\phi} : \underline{\mathcal{C}} \longrightarrow \underline{\mathcal{D}}$. For any constituent $\underline{c} \in \underline{\mathcal{C}}$, let us call $\underline{\phi}(\underline{c})$ the ***value*** of ϕ at $\underline{c}$. This transformation (sending each object $\mathcal{C}$ to its set $\underline{\mathcal{C}}$ of constituents, and mapping each morphism ϕ to the corresponding function $\underline{\phi}$) is a functor from $\boldsymbol{\mathcal{C}}$ into Set. (If $\boldsymbol{\mathcal{C}}$ has a terminal object, then this functor is naturally isomorphic to the one in §8B.)

8G Axiomatic characterization in abstract categories. Suppose that $\boldsymbol{\mathcal{C}}$ is a concrete category satisfying condition (CM) from Section 7. For every object $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, there is a bijective correspondence between the constituents of $\mathcal{C}$ and the elements of the set underlying $\mathcal{C}$. (This justifies using the notation $\underline{\mathcal{C}}$ to refer to both sets.) In light of this observation, the concepts introduced in Section 7 extend immediately to any strongly connected category. For any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, we define $L(\mathcal{C}) := \ell^\infty(\underline{\mathcal{C}})$, to obtain a contravariant functor $L : \boldsymbol{\mathcal{C}}^{\text{op}} \Longmapsto \mathsf{UPOVS}$, which will again be called the *canonical frame*. Likewise, if $\underline{\trianglerighteq}^{\text{xa}}$ is an ex ante preference structure on a decision environment $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$, then for any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, we can define an ex post preference order $\succcurlyeq^{\text{xp}}_{\mathcal{X}}$ on $\underline{\mathcal{X}}$ by considering the restriction of $\succeq^{\mathcal{S}}_{\mathcal{X}}$ to constant morphisms from $\mathcal{S}$ into $\mathcal{X}$, and Lemma 7.1 still holds in this setting: if the subcategory $\boldsymbol{\mathcal{S}}$ is connected, then $\succcurlyeq^{\text{xp}}_{\mathcal{X}}$ is defined on $\underline{\mathcal{X}}$ independent of the choice of $\mathcal{S}$. The probabilistic extension $\Delta\boldsymbol{\mathcal{C}}$ of $\boldsymbol{\mathcal{C}}$ is defined exactly as before, and the four axioms of Section 7 have exactly the same formulation. We still require $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ to satisfy structural conditions (SC) and (XD), but instead of (CM), we require $\boldsymbol{\mathcal{C}}$ to satisfy the following condition:

(CC) $\boldsymbol{\mathcal{C}}$ is strongly connected.

Here is the abstract version of Theorem 7.1.

Theorem 8.1 *Suppose a category $\boldsymbol{\mathcal{C}}$ and decision environment $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ satisfy structural assumptions* (CC)*,* (SC)*, and* (XD)*. Let $\underline{\trianglerighteq}^{\text{xa}}$ be an ex ante preference structure on $(\boldsymbol{\mathcal{S}}, \boldsymbol{\mathcal{X}})$ in the category $\Delta\boldsymbol{\mathcal{C}}$. If $\underline{\trianglerighteq}^{\text{xa}}$ satisfies* vNM independence*,* Continuity*,* Statewise Dominance*, and* Boundedness*, then it admits a global SEU representation.*

In this representation, the utility frame L is a utility subframe of the canonical frame, the belief system $\{\rho_{\mathcal{S}}\}_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}$ is unique, and the utility system (U, A) is unique up to cardinal equivalence. The affinity functor A maps every $\boldsymbol{\mathcal{X}}$-morphism to an affine contraction. If $\boldsymbol{\mathcal{X}}$ is also strongly connected, then A maps every $\boldsymbol{\mathcal{X}}$-morphism to the identity map.

There are three ways in which Theorems 7.1 and 8.1 could be extended or improved. First, like the original theorem of Anscombe and Aumann (1963), they rely on a pre-existing "objective" probabilities (in the probabilistic extension of $\boldsymbol{\mathcal{C}}$) in order to derive the existence of the "subjective" beliefs $\{\rho_{\mathcal{S}}\}_{\mathcal{S}\in\boldsymbol{\mathcal{S}}^\circ}$ supporting SEU representation. This is not entirely satisfactory. A better approach would be to follow the path laid by Savage (1954), and derive an SEU representation *ex nihilo*, without any pre-existing probabilistic structure at all. Second, while it is formulated for abstract categories, Theorem 8.1 still depends on *constituents* of objects in the category $\boldsymbol{\mathcal{C}}$; it would be better to shed this dependence. Finally, it is important to also define and axiomatically characterize *non*-expected utility representations for ex ante preference structures. These are interesting avenues for future research.

A Proofs

Proof of Example 1.2. By contradiction, suppose that μ is a probability distribution on $\mathcal{S}_1 \times \mathcal{S}_2 \times \mathcal{S}_3$ that yields the marginals μ_{12}, μ_{23} and μ_{13}. Then we have the following figure, where each three-vertex line represents an equation. For example, the line $\mu_{23}(d, f) \cdots\cdots \mu(a, d, f)$ —— $\mu(b, d, f)$ means that $\mu_{23}(d, f) = \mu(a, d, f) + \mu(b, d, f)$.

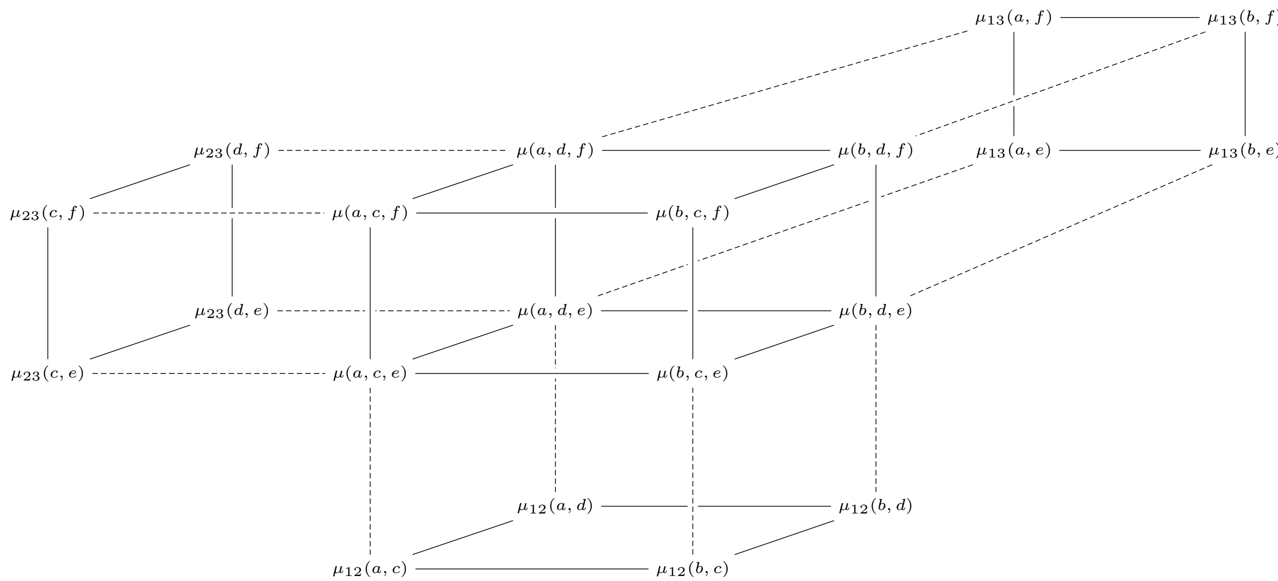

Substituting the marginal values specified in Example 1.2, this reduces to the diagram

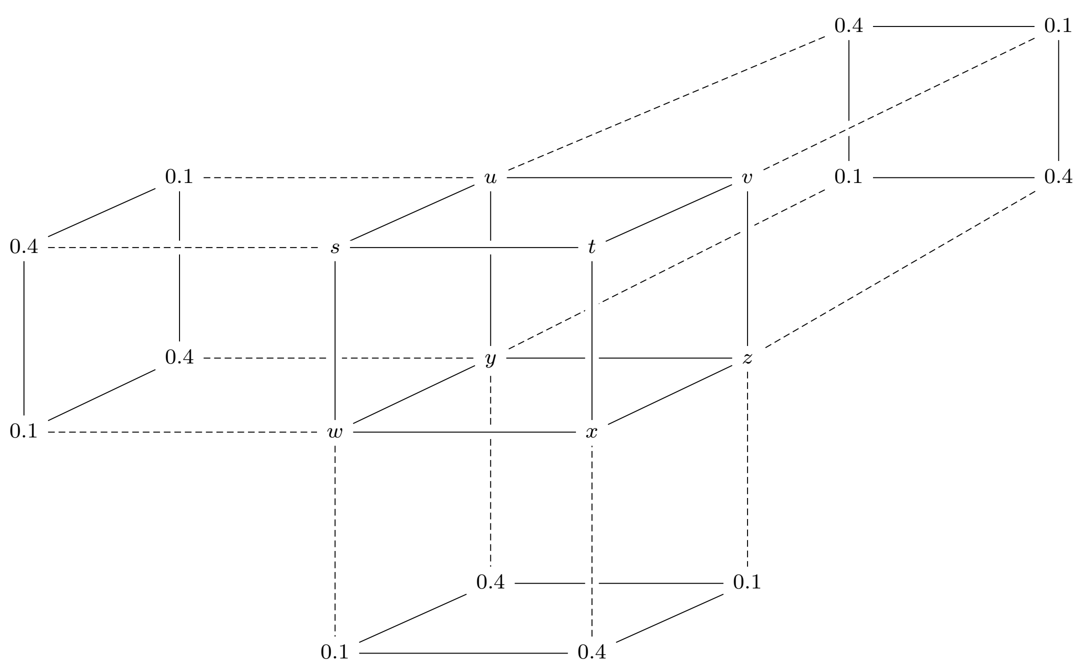

where $s = \mu(a, c, f)$, $t = \mu(b, c, f)$, etc. Thus, the line $0.4 \cdots\cdots s$ —— t means that $0.4 = s + t$. By inspecting the diagram, we see that each of the eight variables s, t, u, v, w, x, y, z is a summand in an equation with value 0.1. Since all these variables are non-negative, we must have $0 \leq s, t, u, v, w, x, y, z \leq 0.1$. But then $s + t + u + v + w + x + y + z \leq 0.8 < 1$, which means that μ cannot be a probability distribution. □

Proof of Proposition 6.5. (a) For any $\mathcal{C}_1, \mathcal{C}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_1, \mathcal{C}_2)$, we must verify that the diagram (1) commutes. To see this, recall that $\overleftarrow{\phi}(\alpha) := \alpha \circ \phi$ for all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_2, \mathcal{X})$. Thus, $U_{\mathcal{X}}^{\mathcal{C}_1} \circ \overleftarrow{\phi}(\alpha) = U_{\mathcal{X}}^{\mathcal{C}_1}[\overleftarrow{\phi}(\alpha)] = U_{\mathcal{X}}^{\mathcal{C}_1}[\alpha \circ \phi] = L(\alpha \circ \phi)[u] = L(\phi) \circ L(\alpha)[u] = L(\phi) \circ U_{\mathcal{X}}^{\mathcal{C}_2}(\alpha)$. This holds for all $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_2, \mathcal{X})$; thus, $U_{\mathcal{X}}^{\mathcal{C}_1} \circ \overleftarrow{\phi} = L(\phi) \circ U_{\mathcal{X}}^{\mathcal{C}_2}$, so diagram (1) commutes.

(b) This follows immediately from the contravariant form of the Yoneda Lemma (Awodey, 2010, Lemma 8.2, p.188).[27] □

In the proofs of Propositions 6.8 and 6.9, we use **AB** to cite Aliprantis and Border (2006).

Proof of Proposition 6.8. A Riesz space $\mathcal{V}$ is a ***Fréchet lattice*** if there is a topology $\mathfrak{T}$ on $\mathcal{V}$ such that:

- $\mathfrak{T}$ makes $\mathcal{V}$ into a topological vector space;
- $\mathfrak{T}$ is generated by a complete metric; and
- $\mathfrak{T}$ is ***locally solid***, meaning that the zero vector $\mathbf{0}$ has a neighbourhood basis consisting of solid sets;[28] equivalently, the lattice operations of $\mathcal{V}$ are uniformly continuous (**AB**, Theorem 8.41, p.334).

For example, any Banach lattice (in particular, any M-space) is a Fréchet lattice. If $\mathcal{V}$ is a Fréchet lattice and $\phi : \mathcal{V} \longrightarrow \mathbb{R}$ is a linear function, then ϕ is continuous with respect to the Fréchet topology if and only if ϕ is order-preserving (**AB**, Thm. 9.11, p.352). All of the unitary POVS which appear in the presheaves of Example 6.3(a-e) are Fréchet lattices. Thus, to characterize the beliefs $\rho_\mathcal{S}$ that appear in belief systems defined on these presheaves, it suffices to look at continuous linear functionals. These are characterized by versions of the Riesz Representation Theorem, as explained below.

(a) For any set $\mathcal{S}$, and every continuous linear functional $\rho_\mathcal{S} : \ell^\infty(\mathcal{S}) \longrightarrow \mathbb{R}$, there is a unique finitely additive signed measure $\mu_\mathcal{S}$ on $\wp(\mathcal{S})$ with bounded variation that satisfies equation (5) (**AB**, Corollary 14.11, p.496). If $\rho_\mathcal{S}$ is order-preserving, then $\mu_\mathcal{S}$ is non-negative. If $\rho_\mathcal{S}$ is uniferent, then $\mu_\mathcal{S}(\mathcal{S}) = 1$, so $\mu_\mathcal{S}$ is a probability measure. (Thus, we can drop the qualifier "with bounded variation", because every probability measure has bounded variation.)

To prove equation (6), let $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$, let $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, and let $\widetilde{\mu} := \phi(\mu_{\mathcal{S}_1})$. For all $v \in \ell^\infty(\mathcal{S}_2)$, we have

$$\int_{\mathcal{S}_2} v \, \mathrm{d}\widetilde{\mu} \underset{(*)}{=} \int_{\mathcal{S}_1} v \circ \phi \, \mathrm{d}\mu_{\mathcal{S}_1} \underset{(\dagger)}{=} \rho_{\mathcal{S}_1}(v \circ \phi) \underset{(\ddagger)}{=} \rho_{\mathcal{S}_1} \circ L(\phi)(v) \underset{(\diamond)}{=} \rho_{\mathcal{S}_2}(v) \underset{(\dagger)}{=} \int_{\mathcal{S}_2} v \, \mathrm{d}\mu_{\mathcal{S}_2},$$

where $(*)$ is by the Change of Variables Theorem (**AB**, Thm.13.46, p.484), both $(\dagger)$ are by equation (5), $(\ddagger)$ is the definition of $L(\phi)$ from Example 6.3, and $(\diamond)$ is by commuting

[27] For *covariant* forms of the Yoneda Lemma, see Riehl (2017, Theorem 2.2.4, p.57) or Adámek et al. (2009, Corollary 6.19, p.88). These become the contravariant form through an extra dualization step.

[28] A subset $\mathcal{B} \subset \mathcal{V}$ is ***solid*** if, for all $b \in \mathcal{B}$ and $v \in \mathcal{V}$, $|v| \leq |b| \implies v \in \mathcal{B}$.

diagram (4). Thus, both $\widetilde{\mu}$ and $\mu_{\mathcal{S}_2}$ satisfy equation (5) for $\rho_{\mathcal{S}_2}$. But the measure with this property is *unique*, so we must have $\widetilde{\mu} = \mu_{\mathcal{S}_2}$. In other words, $\phi(\mu_{\mathcal{S}_1}) = \mu_{\mathcal{S}_2}$. This argument applies to all $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$. This proves equation (6).

(b) For any measurable space $\mathcal{S}$ and continuous linear functional $\rho : \mathfrak{L}^\infty(\mathcal{S}) \longrightarrow \mathbb{R}$, there is a unique finitely additive signed measure $\mu_{\mathcal{S}}$ (with bounded variation) on the sigma-algebra of $\mathcal{S}$ that satisfies equation (5) (**AB**, Thm. 14.4, p.489). Now proceed as in part (a).

(c) For any normal Hausdorff space $\mathcal{S}$, and every continuous positive linear functional $\rho : \mathfrak{C}_{\mathrm{b}}(\mathcal{S}) \longrightarrow \mathbb{R}$, there is a unique finitely additive, normal, positive measure $\mu_{\mathcal{S}}$ on the Borel sigma-algebra of $\mathcal{S}$ that satisfies equation (5) (**AB**, Theorem 14.9, p.491). Furthermore, if $\mathcal{S}$ is a compact Hausdorff space, then $\mu_{\mathcal{S}}$ is countably additive (**AB**, Thm. 14.14, p.409).[29] Now proceed as in part (a).

(d) Let $\mathcal{S}$ be a metric space. The subspace $\mathfrak{LL}_{\mathrm{b}}(\mathcal{S})$ contains the order unit $\mathbf{1}$, so any positive linear functional $\rho : \mathfrak{LL}_{\mathrm{b}}(\mathcal{S}) \longrightarrow \mathbb{R}$ extends to a positive linear functional $\widetilde{\rho} : \mathfrak{C}_{\mathrm{b}}(\mathcal{S}) \longrightarrow \mathbb{R}$ (**AB**, Corollary 8.33, p.331). As explained above, $\widetilde{\rho}$ is continuous in the uniform norm (**AB**, Thm. 9.11, p.352). Furthermore, if $\mathcal{S}$ is compact, then $\mathfrak{LL}_{\mathrm{b}}(\mathcal{S}) = \mathfrak{LL}(\mathcal{S})$ is uniformly dense in $\mathfrak{C}_{\mathrm{b}}(\mathcal{S}) = \mathfrak{C}(\mathcal{S})$, by the Stone-Weierstrass Theorem (**AB**, Thm. 9.12, p.352). Thus, the extension $\widetilde{\rho}$ is unique. But for any compact metric space $\mathcal{S}$ and any continuous linear functional $\widetilde{\rho} : \mathfrak{C}(\mathcal{S}) \longrightarrow \mathbb{R}$, there is a unique countably additive signed Borel measure $\mu_{\mathcal{S}}$ that satisfies equation (5) (**AB**, Thm. 14.15, p.491). Now proceed as in part (a). □

Proof of Proposition 6.9. For any unitary Archimedean Riesz space $\mathcal{V}$, there is a compact Hausdorff space $\mathcal{V}^\S$ (unique up to homeomorphism) and an injective, uniferent, order-preserving linear function $ß_{\mathcal{V}} : \mathcal{V} \longrightarrow \mathfrak{C}(\mathcal{V}^\S)$ whose image is both order-dense and uniformly dense in $\mathfrak{C}(\mathcal{V}^\S)$ (Fremlin, 2012, Theorem 353M). If $\mathcal{W}$ is another unitary Archimedean Riesz space, and $\phi : \mathcal{V} \longrightarrow \mathcal{W}$ is a uniferent, order-preserving linear function, then there is a (unique) continuous function $\phi^\S : \mathcal{W}^\S \longrightarrow \mathcal{V}^\S$ such that, for all $v \in \mathcal{V}$, if $w := \phi(v)$ and $v' := ß_{\mathcal{V}}(v)$, then $ß_{\mathcal{W}}(w) = v' \circ \phi^\S$ (Pivato, 2020, Proposition B1(c)). Furthermore, if $\mathcal{U}$ is a third Riesz space and $\psi : \mathcal{U} \longrightarrow \mathcal{V}$ is another uniferent, order-preserving linear function, then $(\phi \circ \psi)^\S = \psi^\S \circ \phi^\S$.[30] Thus, the mappings $\mathcal{V} \mapsto \mathcal{V}^\S$ and $\phi \mapsto \phi^\S$ determine a contravariant functor $\S : \mathsf{UARiesz}^{\mathrm{op}} \Longrightarrow \mathsf{CHS}$.[31] Finally, if $\mathcal{V}$ is an M-space, then $ß_{\mathcal{V}}$ is bijective, hence an isomorphism (see **AB**, Theorem 9.32, or Fremlin 2012, Corollary 354L, or Meyer-Nieberg 1991, Theorem 2.1.3).

---

[29] This result actually applies to linear functionals on the space of continuous real-valued functions with *compact support*; this justifies Footnote 14. But if $\mathcal{S}$ is compact, then *every* function has compact support.

[30] This follows from the construction of these objects. The space $\mathcal{V}^\S$ is the set of uniferent order-preserving linear functions from $\mathcal{V}$ into $\mathbb{R}$, endowed with the topology of pointwise convergence. (Likewise for $\mathcal{U}^\S$ and $\mathcal{W}^\S$). The function $\phi^\S : \mathcal{W}^\S \longrightarrow \mathcal{V}^\S$ is defined: $\phi^\S(\eta) = \eta \circ \phi$ for all $\eta \in \mathcal{W}^\S$. (Likewise for $\psi^\S$). It follows that $(\phi \circ \psi)^\S = \psi^\S \circ \phi^\S$. See the proof of Proposition B1(c) in Pivato (2020) for further explanation.

[31] It is essentially the contravariant hom functor $\overrightarrow{\mathsf{UARiesz}}(\bullet, \mathbb{R})$, but it maps into CHS instead of Set.

(a) Define $H := \S \circ L : \boldsymbol{\mathcal{S}} \Longmapsto \mathsf{CHS}$. This is a composition of two contravariant functors, so it is a covariant functor. For every $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, we get a compact Hausdorff space $H(\mathcal{S}) := L(\mathcal{S})^\S$. For every morphism $\phi \in \vec{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, we get a continuous function $H(\phi) := L(\phi)^\S : H(\mathcal{S}_1) \longrightarrow H(\mathcal{S}_2)$.

(b) Let $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{C}}^\circ$, and let $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}_1, \mathcal{S}_2)$. We must show that this diagram commutes:

$$\begin{array}{ccc} L(\mathcal{S}_1) & \xleftarrow{L(\phi)} & L(\mathcal{S}_2) \\ {\scriptstyle ß_{\mathcal{S}_1}}\downarrow & & \downarrow{\scriptstyle ß_{\mathcal{S}_2}} \\ \mathfrak{C}[H(\mathcal{S}_1)] & \xleftarrow[\mathfrak{C}[H(\phi)]]{} & \mathfrak{C}[H(\mathcal{S}_2)] \end{array} \tag{A1}$$

This is just a question of working through the definitions. For $j = 1, 2$, recall that $H(\mathcal{S}_j) = L(\mathcal{S}_j)^\S$ is the set of all uniferent order-preserving linear functions from $L(\mathcal{S}_j)$ into $\mathbb{R}$; we will denote a generic function of this type by η_j. The function $ß_j : L(\mathcal{S}_j) \longrightarrow \mathfrak{C}[H(\mathcal{S}_j)]$ is defined as follows: for any $v_j \in L(\mathcal{S}_j)$, the function $ß_j(v_j) : H(\mathcal{S}_j) \longrightarrow \mathbb{R}$ is defined by

$$ß_j(v_j)(\eta_j) \quad := \quad \eta_j(v_j), \qquad \text{for all } \eta_j \in H(\mathcal{S}_j). \tag{A2}$$

Meanwhile, for any uniferent order-preserving linear function $\Phi : L(\mathcal{S}_2) \longrightarrow L(\mathcal{S}_1)$, the induced map $\Phi^\S : H(\mathcal{S}_1) \longrightarrow H(\mathcal{S}_2)$ is defined by

$$\Phi^\S(\eta_1) \quad := \quad \eta_1 \circ \Phi, \qquad \text{for all } \eta_1 \in H(\mathcal{S}_1). \tag{A3}$$

In particular, let $\Phi := L(\phi)$. Then $H(\phi) = L(\phi)^\S = \Phi^\S$ is defined by (A3).

Now, for any continuous $\Psi : H(\mathcal{S}_1) \longrightarrow H(\mathcal{S}_2)$, the function $\mathfrak{C}[\Psi] : \mathfrak{C}[H(\mathcal{S}_2)] \longrightarrow \mathfrak{C}[H(\mathcal{S}_1)]$ is defined as follows: for any continuous function $f_2 : H(\mathcal{S}_2) \longrightarrow \mathbb{R}$, we have $\mathfrak{C}[\Psi](f_2) := f_2 \circ \Psi$ (a continuous function from $H(\mathcal{S}_1)$ to $\mathbb{R}$). Setting $\Psi := \Phi^\S$, we see that $\mathfrak{C}[\Phi^\S] : \mathfrak{C}[H(\mathcal{S}_2)] \longrightarrow \mathfrak{C}[H(\mathcal{S}_1)]$ is defined by

$$\mathfrak{C}[\Phi^\S](f_2) \quad = \quad f_2 \circ \Phi^\S, \qquad \text{for all } f_2 \in \mathfrak{C}[H(\mathcal{S}_2)]. \tag{A4}$$

Thus, for any $\eta_1 \in H(\mathcal{S}_1)$, we have

$$\mathfrak{C}[\Phi^\S](f_2)(\eta_1) \underset{(*)}{=} f_2 \circ \Phi^\S(\eta_1) = f_2[\Phi^\S(\eta_1)] \underset{(\dagger)}{=} f_2[\eta_1 \circ \Phi], \tag{A5}$$

where $(*)$ is by (A4) and $(\dagger)$ is by (A3).

In particular, let $v_2 \in L(\mathcal{S}_2)$. Then $\mathfrak{C}[\Phi^\S] \circ ß_{\mathcal{S}_2}(v_2)$ is the function defined by

$$\mathfrak{C}[\Phi^\S] \circ ß_{\mathcal{S}_2}(v_2)(\eta_1) \underset{(*)}{=} ß_{\mathcal{S}_2}(v_2)[\eta_1 \circ \Phi] \underset{(\dagger)}{=} \eta_1 \circ \Phi(v_2) \underset{(\ddagger)}{=} ß_{\mathcal{S}_1}[\Phi(v_2)](\eta_1), \tag{A6}$$

for all $\eta_1 \in H(\mathcal{S}_1)$. Here, $(*)$ is obtained by substituting $f_2 := ß_{\mathcal{S}_2}(v_2)$ into (A5), while $(\dagger)$ comes from applying equation (A2) to $ß_{\mathcal{S}_2}$ and $(\ddagger)$ comes from applying (A2) to $ß_{\mathcal{S}_1}$.

Equation (A6) holds for all $\eta_1 \in H(\mathcal{S}_1)$. Thus, $\mathfrak{C}[\Phi^\S] \circ ß_{\mathcal{S}_2}(v_2) = ß_{\mathcal{S}_1}[\Phi(v_2)]$. *This* equation holds for any $v_2 \in L(\mathcal{S}_2)$. Thus, $\mathfrak{C}[\Phi^\S] \circ ß_{\mathcal{S}_2} = ß_{\mathcal{S}_1} \circ \Phi$. Recalling that $\Phi = L(\phi)$ and $\Phi^\S = H(\phi)$, we get $\mathfrak{C}[H(\phi)] \circ ß_{\mathcal{S}_2} = ß_{\mathcal{S}_1} \circ L(\phi)$. Thus, diagram (A1) commutes.

This argument works for any $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{S}_1, \mathcal{S}_2)$. So ß is a natural transformation from L to $\mathfrak{C} \circ H$.

(c) Let $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and consider the belief $\rho_\mathcal{S} : L(\mathcal{S}) \longrightarrow \mathbb{R}$. Let $\text{ß}_\mathcal{S} : L(\mathcal{S}) \longrightarrow \mathfrak{C}(H(\mathcal{S}), \mathbb{R})$ be the uniferent, order-preserving linear injection described above, and let $\mathcal{V} := \text{ß}_\mathcal{S}[L(\mathcal{S})]$; then $\mathcal{V}$ is a uniformly dense subspace of $\mathfrak{C}(H(\mathcal{S}), \mathbb{R})$, and it contains the constant 1-valued function (because $\text{ß}_\mathcal{S}$ is uniferent). Furthermore, $\text{ß}_\mathcal{S} : L(\mathcal{S}) \longrightarrow \mathcal{V}$ is an isomorphism. Define $\widetilde{\rho} := \rho_\mathcal{S} \circ \text{ß}_\mathcal{S}^{-1} : \mathcal{V} \longrightarrow \mathbb{R}$; then $\widetilde{\rho}$ is a uniferent, order-preserving linear function, because it is a composition of two such functions. Thus, it can be extended to a uniferent, order-preserving linear function $\widehat{\rho} : \mathfrak{C}(H(\mathcal{S}), \mathbb{R}) \longrightarrow \mathbb{R}$ (**AB**, Corollary 8.33, p.331). Now, $\mathfrak{C}(H(\mathcal{S}), \mathbb{R})$ is a Banach lattice, so $\widehat{\rho}$ is continuous with respect to the uniform norm (**AB**, Theorem 9.6, p.350). Since $\mathcal{V}$ is uniformly dense in $\mathfrak{C}(H(\mathcal{S}), \mathbb{R})$, we conclude that the extension $\widehat{\rho}$ is unique. Since $H(\mathcal{S})$ is a compact Hausdorff space, there is a unique Borel probability measure $\mu_\mathcal{S}$ on $H(\mathcal{S})$ such that $\widehat{\rho}(f) = \int_{H(\mathcal{S})} f \, \mathrm{d}\mu_\mathcal{S}$ for all $f \in \mathfrak{C}(H(\mathcal{S}), \mathbb{R})$ (**AB**, Theorem 14.14, p.497). But if $f := \text{ß}_\mathcal{S}(v)$, then $\widehat{\rho}(f) = \widetilde{\rho}(f) = \rho \circ \text{ß}_\mathcal{S}^{-1}(f) = \rho(v)$. This proves equation (7). The proof of equation (8) is very similar to the proof of equation (6) in Proposition 6.8. □

Proof of Lemma 7.1. Let $\mathcal{S}, \mathcal{S}' \in \boldsymbol{\mathcal{S}}$. We must show that $\overset{\text{xp}}{\succcurlyeq}_{\mathcal{S},\mathcal{X}}$ and $\overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}',\mathcal{X}}$ are identical.

First suppose that $\mathcal{S} \sim \mathcal{S}'$. Then either $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}, \mathcal{S}') \neq \emptyset$ or $\overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}', \mathcal{S}) \neq \emptyset$. Without loss of generality, assume the former. Let $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}, \mathcal{S}')$. Then $\kappa_{\underline{x}}^{\mathcal{S}'} \circ \phi = \kappa_{\underline{x}}^{\mathcal{S}}$ and $\kappa_{\underline{y}}^{\mathcal{S}'} \circ \phi = \kappa_{\underline{y}}^{\mathcal{S}}$, by the definition of constant morphisms and their values. Thus,

$$\begin{aligned} \underline{x} \overset{\text{xp}}{\succcurlyeq}_{\mathcal{S},\mathcal{X}} \underline{y} \quad &\underset{(*)}{\Longleftrightarrow} \quad \kappa_{\underline{x}}^{\mathcal{S}} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \kappa_{\underline{y}}^{\mathcal{S}} \quad \Longleftrightarrow \quad \kappa_{\underline{x}}^{\mathcal{S}'} \circ \phi \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \kappa_{\underline{y}}^{\mathcal{S}'} \circ \phi \\ &\underset{(\dagger)}{\Longleftrightarrow} \quad \kappa_{\underline{x}}^{\mathcal{S}'} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}'} \kappa_{\underline{y}}^{\mathcal{S}'} \quad \underset{(*)}{\Longleftrightarrow} \quad \underline{x} \overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}',\mathcal{X}} \underline{y}, \end{aligned}$$

as desired. Here, both $(*)$ are by the definition of $\overset{\text{xp}}{\succcurlyeq}_{\mathcal{S},\mathcal{X}}$ and $\overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}',\mathcal{X}}$, while $(\dagger)$ is by the property (BP).

Now let $\mathcal{S}$ and $\mathcal{S}'$ be arbitrary objects in $\boldsymbol{\mathcal{S}}^\circ$. Since $\boldsymbol{\mathcal{S}}$ is connected, there is a path $\mathcal{S} \sim \mathcal{S}_1 \sim \mathcal{S}_2 \sim \ldots \sim \mathcal{S}_N \sim \mathcal{S}'$ in the graph $(\boldsymbol{\mathcal{S}}, \sim)$. By the previous case, we have $\overset{\text{xp}}{\succcurlyeq}_{\mathcal{S},\mathcal{X}} = \overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}_1,\mathcal{X}} = \cdots = \overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}_N,\mathcal{X}} = \overset{\text{xp}}{\succcurlyeq}_{\mathcal{S}',\mathcal{X}}$. □

Let $\mathcal{X}$ be a convex set. A function $v : \mathcal{X} \longrightarrow \mathbb{R}$ is *mixture-preserving* if $v(x \circledast_q y) = q\, v(x) + (1-q)\, v(y)$ for all $x, y \in \mathcal{X}$ and $q \in [0,1]$.[32] The proof of Theorems 7.1 and 8.1 will use the classic theorem of von Neumann and Morgenstern, which we restate for reference.

Proposition A.1 *Let $\mathcal{X}$ be a convex subset of a vector space, and let $\succcurlyeq$ be a weak order on $\mathcal{X}$. Suppose $\succcurlyeq$ satisfies the following axioms:*

Independence. *For all $x, y \in \mathcal{X}$, if $x \succcurlyeq y$, then $x \circledast_q z \succcurlyeq y \circledast_q z$ for all $z \in \mathcal{X}$ and $q \in [0,1]$.*

Continuity. *For all $x, y, z \in \mathcal{X}$, the sets $\{q \in [0,1];\ x \circledast_q y \succcurlyeq z\}$ and $\{q \in [0,1];\ x \circledast_q y \preccurlyeq z\}$ are closed.*

[32] These are also called *affine* functions. We eschew this term to avoid confusion with elements of Aff.

Then there is a mixture-preserving function $v : \mathcal{X} \longrightarrow \mathbb{R}$ such that,

$$\text{for all } x, y \in \mathcal{X}, \qquad x \succcurlyeq y \quad \Longleftrightarrow \quad v(x) \geq v(y). \tag{A7}$$

Furthermore v is unique up to positive affine transformations. If $v' : \mathcal{X} \longrightarrow \mathbb{R}$ is another mixture-preserving function satisfying (A7), then there is some $\phi \in \mathsf{Aff}$ such that $v' = \phi \circ v$.

If a concrete category $\boldsymbol{\mathcal{C}}$ satisfies (CM), then it satisfies (CC), and the constituents of each object $\mathcal{C}$ in $\boldsymbol{\mathcal{C}}^\circ$ correspond bijectively with the elements of the underlying set $\underline{\mathcal{C}}$. Thus, Theorem 7.1 is just a special case of Theorem 8.1. So it suffices to prove Theorem 8.1.

Proof of Theorem 8.1. For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, Proposition A.1 yields a mixture-preserving function $V_{\mathcal{X}}^{\mathcal{S}} : \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X}) \longrightarrow \mathbb{R}$ satisfying the following instance of (A7):

$$\text{for all } \widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X}), \qquad \widetilde{\alpha} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta} \quad \Longleftrightarrow \quad V_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) \geq V_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta}). \tag{A8}$$

Furthermore, this function is unique up to positive affine transformation.

For any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, let $\underline{\Delta\mathcal{X}}$ denote the set of constituents of $\mathcal{X}$ in the category $\Delta\boldsymbol{\mathcal{C}}$. For any $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ the elements of $\underline{\Delta\mathcal{X}}$ can be identified with the constant morphisms in $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$ (see Lemma B.3 in Appendix B below). Generic elements will be denoted $\underline{\widetilde{x}}, \underline{\widetilde{y}}, \underline{\widetilde{z}}$, etc. By definition, the ex post preference order $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$ is the restriction of $\succcurlyeq_{\mathcal{X}}^{\mathcal{S}}$ to $\underline{\Delta\mathcal{X}}$. (By Lemma 7.1 and property (SC), this is true for any $\mathcal{S}$.) Thus, if $v_{\mathcal{X}}^{\mathcal{S}}$ is the restriction of $V_{\mathcal{X}}^{\mathcal{S}}$ to $\underline{\Delta\mathcal{X}}$, then we have the utility representation:

$$\text{for all } \underline{\widetilde{x}}, \underline{\widetilde{y}} \in \underline{\Delta\mathcal{X}}, \qquad \underline{\widetilde{x}} \succcurlyeq_{\mathcal{X}}^{\text{xp}} \underline{\widetilde{y}} \quad \Longleftrightarrow \quad v_{\mathcal{X}}^{\mathcal{S}}(\underline{\widetilde{x}}) \geq v_{\mathcal{X}}^{\mathcal{S}}(\underline{\widetilde{y}}). \tag{A9}$$

Now, $\underline{\Delta\mathcal{X}}$ is a mixture space, and $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$ satisfies the axioms vNM independence and Continuity (because it is the restriction of $\succcurlyeq_{\mathcal{X}}^{\mathcal{S}}$, which satisfies these axioms). Furthermore, $v_{\mathcal{X}}^{\mathcal{S}}$ is a mixture-preserving function. So (A9) says that $v_{\mathcal{X}}^{\mathcal{S}}$ is a vNM utility representation of $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$. This can be repeated for any $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$. Thus, by the uniqueness part of Proposition A.1, we conclude that there is a *single* mixture-preserving utility function $v_{\mathcal{X}} : \underline{\Delta\mathcal{X}} \longrightarrow \mathbb{R}$ such that, for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, there exist constants $A_{\mathcal{S}} > 0$ and $B_{\mathcal{S}} \in \mathbb{R}$ such that $v_{\mathcal{X}}^{\mathcal{S}} = A_{\mathcal{S}}\, v_{\mathcal{X}} + B_{\mathcal{S}}$. By replacing $V_{\mathcal{X}}^{\mathcal{S}}$ with $\widehat{V}_{\mathcal{X}}^{\mathcal{S}} := (V_{\mathcal{X}}^{\mathcal{S}} - B_{\mathcal{S}})/A_{\mathcal{S}}$ for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ if necessary, we can assume without loss of generality that the functions $V_{\mathcal{X}}^{\mathcal{S}}$ which appear in the representation (A8) are such that $V_{\mathcal{X}}^{\mathcal{S}}$ is equal to $v_{\mathcal{X}}$ when restricted to $\underline{\Delta\mathcal{X}}$, for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$.

By Boundedness, $\underline{\mathcal{X}}$ contains a $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$-minimal constituent $\underline{o}$ and $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$-maximal element $\underline{l}$. By applying a positive affine transformation to $v_{\mathcal{X}}$ if necessary, we can assume without loss of generality that $v_{\mathcal{X}}(\underline{o}) = 0$ and $v_{\mathcal{X}}(\underline{l}) = 1$. Thus, $0 \leq v_{\mathcal{X}}(\underline{x}) \leq 1$ for all $\underline{x} \in \underline{\mathcal{X}}$. So $v_{\mathcal{X}}$ is bounded.

Claim 1: *There is a canonical bijection from $\underline{\Delta\mathcal{X}}$ to the set $\Delta(\underline{\mathcal{X}})$ of finite-support probability distributions on $\underline{\mathcal{X}}$.*

Proof. Let $\underline{\widetilde{x}} \in \underline{\Delta\mathcal{X}}$. There is some $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and constant morphism $\widetilde{\phi} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ such that $\underline{\widetilde{x}}$ is the value of $\widetilde{\phi}$. Suppose that $\widetilde{\phi} = (p_1, \phi_1;\ p_2, \phi_2;\ \ldots;\ p_N, \phi_N)$.

First note that $\phi_1, \ldots, \phi_N$ must themselves be constant morphisms in the category $\boldsymbol{\mathcal{C}}$. To see this, suppose that ϕ_n is *not* a constant. Then there is some $\mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$ and morphisms $\alpha, \beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ such that $\phi_n \circ \alpha \neq \phi_n \circ \beta$. But $\widetilde{\phi} \circ \alpha = (p_1, \phi_1 \circ \alpha;\ \ldots;\ p_N, \phi_N \circ \alpha)$ and $\widetilde{\phi} \circ \beta = (p_1, \phi_1 \circ \beta;\ \ldots;\ p_N, \phi_N \circ \beta)$. If $\phi_n \circ \alpha \neq \phi_n \circ \beta$, then $\widetilde{\phi} \circ \alpha \neq \widetilde{\phi} \circ \beta$, contradicting the fact that $\widetilde{\phi}$ is constant.

Let $\underline{x}_1, \ldots, \underline{x}_N \in \underline{\mathcal{X}}$ be the values of $\phi_1, \ldots, \phi_N$. Then $(p_1, \underline{x}_1;\ p_2, \underline{x}_2;\ \ldots;\ p_N, \underline{x}_N)$, an element of $\Delta(\underline{\mathcal{X}})$. In this way, every element of $\underline{\Delta\mathcal{X}}$ determines an element of $\Delta(\underline{\mathcal{X}})$. Conversely, given any element $(p_1, \underline{x}_1;\ p_2, \underline{x}_2;\ \ldots;\ p_N, \underline{x}_N)$ of $\Delta(\underline{\mathcal{X}})$, let $\phi_1, \ldots, \phi_N \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$ be constant morphisms with values $\underline{x}_1, \ldots, \underline{x}_N$ (where $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ is an arbitrary object). Let $\widetilde{\phi} = (p_1, \phi_1;\ p_2, \phi_2;\ \ldots;\ p_N, \phi_N)$. Then $\widetilde{\phi} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ is a constant. So, let $\underline{\widetilde{x}} \in \underline{\Delta\mathcal{X}}$ be its value. ◇ `Claim 1`

In light of this claim, we can treat $\succeq^{\text{xp}}_{\mathcal{X}}$ as a preference order on $\Delta(\underline{\mathcal{X}})$ satisfying the axioms of Proposition A.1, and we can treat $v_{\mathcal{X}}$ as a mixture-preserving utility function on $\Delta(\underline{\mathcal{X}})$ that represents $\succeq^{\text{xp}}_{\mathcal{X}}$.

Let L be the canonical frame on $\boldsymbol{\mathcal{C}}$, so that $L(\mathcal{C}) := \ell^\infty(\underline{\mathcal{C}})$ for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$. For any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, define $U^{\mathcal{C}}_{\mathcal{X}} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \mathbb{R}^{\underline{\mathcal{C}}}$ as follows:

$$\text{for all } \alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}), \text{ let} \quad U^{\mathcal{C}}_{\mathcal{X}}(\alpha) \ := \ v_{\mathcal{X}} \circ \underline{\alpha}. \tag{A10}$$

(Recall that $\underline{\alpha} : \underline{\mathcal{C}} \longrightarrow \underline{\mathcal{X}}$ and $v_{\mathcal{X}} : \underline{\mathcal{X}} \longrightarrow \mathbb{R}$.) As observed above, $v_{\mathcal{X}}$ is a bounded function. Thus, $U^{\mathcal{C}}_{\mathcal{X}}(\alpha)$ is also bounded. Thus, $U^{\mathcal{C}}_{\mathcal{X}}(\alpha) \in \ell^\infty(\underline{\mathcal{C}})$. We thus obtain a function $U^{\mathcal{C}}_{\mathcal{X}} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \ell^\infty(\underline{\mathcal{C}})$.

Claim 2: *For any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, the system $U_{\mathcal{X}} := (U^{\mathcal{C}}_{\mathcal{X}};\ \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ)$ is a utility functional.*

Proof. We must show that this collection of functions defines a natural transformation $U_{\mathcal{X}} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\bullet, \mathcal{X}) \Longrightarrow \underline{L}$. The argument is similar to the proof of Proposition 6.5(a). Let $\mathcal{C}_1, \mathcal{C}_2 \in \boldsymbol{\mathcal{C}}^\circ$ and let $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_1, \mathcal{C}_2)$. We must show that the diagram (1) commutes. So, let $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}_2, \mathcal{X})$. We must show that $U^{\mathcal{C}_1}_{\mathcal{X}} \circ \overleftarrow{\phi}(\alpha) = \underline{L}(\phi) \circ U^{\mathcal{C}_2}_{\mathcal{X}}(\alpha)$.

Now, $U^{\mathcal{C}_1}_{\mathcal{X}} \circ \overleftarrow{\phi}(\alpha) = U^{\mathcal{C}_1}_{\mathcal{X}}(\alpha \circ \phi) = v_{\mathcal{X}} \circ \underline{\alpha \circ \phi} = v_{\mathcal{X}} \circ \underline{\alpha} \circ \underline{\phi}$. Meanwhile, $U^{\mathcal{C}_2}_{\mathcal{X}}(\alpha) = v_{\mathcal{X}} \circ \underline{\alpha}$, so $\underline{L}(\phi) \circ U^{\mathcal{C}_2}_{\mathcal{X}}(\alpha) = \underline{L}(\phi)(v_{\mathcal{X}} \circ \underline{\alpha}) = v_{\mathcal{X}} \circ \underline{\alpha} \circ \underline{\phi}$. Thus, $U^{\mathcal{C}_1}_{\mathcal{X}} \circ \overleftarrow{\phi}(\alpha) = \underline{L}(\phi) \circ U^{\mathcal{C}_2}_{\mathcal{X}}(\alpha)$, as desired. ◇ `Claim 2`

Claim 3: *There is a functor $A : \boldsymbol{\mathcal{X}} \Longmapsto \mathsf{Aff}$ making the collection $\{U^{\mathcal{C}}_{\mathcal{X}};\ \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ\}$ into a utility system.*

Proof. For any $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, we must construct an affine function $\widehat{\phi} \in \mathsf{Aff}$ that makes diagram (9) commute for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$.

First, let $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, and recall that $\vec{\phi} : \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X}) \longrightarrow \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{Y})$ is defined by $\vec{\phi}(\widetilde{\alpha}) := \phi \circ \widetilde{\alpha}$ for all $\widetilde{\alpha} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$. This is a mixture-preserving transformation, by defining formula (11). Thus, if we define $\widehat{V}^{\mathcal{S}}_{\mathcal{X}} := V^{\mathcal{S}}_{\mathcal{Y}} \circ \vec{\phi} : \overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X}) \longrightarrow \mathbb{R}$, then $\widehat{V}^{\mathcal{S}}_{\mathcal{X}}$ is another mixture-preserving function on $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$. For any $\widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$, we have:

$$\widetilde{\alpha} \succsim^{\mathcal{S}}_{\mathcal{X}} \widetilde{\beta} \underset{(*)}{\iff} \phi \circ \widetilde{\alpha} \succsim^{\mathcal{S}}_{\mathcal{Y}} \phi \circ \widetilde{\beta} \underset{(\dagger)}{\iff} \vec{\phi}(\widetilde{\alpha}) \succsim^{\mathcal{S}}_{\mathcal{Y}} \vec{\phi}(\widetilde{\beta})$$
$$\underset{(\ddagger)}{\iff} V^{\mathcal{S}}_{\mathcal{Y}} \circ \vec{\phi}(\widetilde{\alpha}) \geq V^{\mathcal{S}}_{\mathcal{Y}} \circ \vec{\phi}(\widetilde{\beta}) \underset{(\diamond)}{\iff} \widehat{V}^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\alpha}) \geq V^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\beta}),$$

where $(*)$ is by property (TP), $(\dagger)$ is the definition of $\vec{\phi}$, $(\ddagger)$ is by the utility representation (A8), and $(\diamond)$ is the definition of $\widehat{V}^{\mathcal{S}}_{\mathcal{X}}$. Thus, $\widehat{V}^{\mathcal{S}}_{\mathcal{X}}$ yields another mixture-preserving utility representation for $\succsim^{\mathcal{S}}_{\mathcal{X}}$. By the uniqueness part of Proposition A.1, there exists $\widehat{\phi}_{\mathcal{S}} \in \mathsf{Aff}$ such that $\widehat{V}^{\mathcal{S}}_{\mathcal{X}} = \widehat{\phi}_{\mathcal{S}} \circ V^{\mathcal{S}}_{\mathcal{X}}$. In other words,

$$V^{\mathcal{S}}_{\mathcal{Y}} \circ \vec{\phi} \quad = \quad \widehat{\phi}_{\mathcal{S}} \circ V^{\mathcal{S}}_{\mathcal{X}}. \tag{A11}$$

Claim 3.1: *$\widehat{\phi}_{\mathcal{S}}$ does not depend on the choice of $\mathcal{S}$.*

Proof. Recall that the constant morphisms in $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$ are in bijective correspondence with the constituents of $\underline{\Delta\mathcal{X}}$. If $\widetilde{\alpha} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{X})$ is a constant morphism with value $\underline{\widetilde{x}}$, then $\vec{\phi}(\widetilde{\alpha}) = \phi \circ \widetilde{\alpha}$ is a constant morphism in $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S},\mathcal{Y})$ with value $\underline{\phi}(\underline{\widetilde{x}})$. For any $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, recall that $V^{\mathcal{S}}_{\mathcal{X}}$, restricted to $\underline{\Delta\mathcal{X}}$, agrees with $v_{\mathcal{X}}$, while $V^{\mathcal{S}}_{\mathcal{Y}}$, restricted to $\underline{\Delta\mathcal{Y}}$, agrees with $v_{\mathcal{Y}}$. Thus, for any $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, equation (A11) yields

$$v_{\mathcal{Y}} \circ \underline{\phi} \quad = \quad \widehat{\phi}_{\mathcal{S}} \circ v_{\mathcal{X}}. \tag{A12}$$

Likewise, for any other $\mathcal{S}' \in \boldsymbol{\mathcal{S}}^{\circ}$, we also have

$$v_{\mathcal{Y}} \circ \underline{\phi} \quad = \quad \widehat{\phi}_{\mathcal{S}'} \circ v_{\mathcal{X}}. \tag{A13}$$

The left hand side of equations (A12) and (A13) agree, so $\widehat{\phi}_{\mathcal{S}} \circ v_{\mathcal{X}} = \widehat{\phi}_{\mathcal{S}'} \circ v_{\mathcal{X}}$. Thus, $\widehat{\phi}_{\mathcal{S}}(\mathbf{v}) = \widehat{\phi}_{\mathcal{S}'}(\mathbf{v})$ for all $\mathbf{v} \in v_{\mathcal{X}}(\underline{\Delta\mathcal{X}})$. Since $\widehat{\phi}_{\mathcal{S}}$ and $\widehat{\phi}_{\mathcal{S}'}$ are affine functions, and the set $v_{\mathcal{X}}(\underline{\Delta\mathcal{X}})$ contains more than one point, this implies that $\widehat{\phi}_{\mathcal{S}} = \widehat{\phi}_{\mathcal{S}'}$. ▽ `Claim 3.1`

Thus, there is a single affine function $\widehat{\phi} \in \mathsf{Aff}$ (independent of the choice of $\mathcal{S}$) such that

$$v_{\mathcal{Y}} \circ \underline{\phi} \quad = \quad \widehat{\phi} \circ v_{\mathcal{X}}. \tag{A14}$$

Now, for any $\mathcal{C} \in \boldsymbol{\mathcal{C}}^{\circ}$, recall that $L(\mathcal{C}) = \ell^{\infty}(\underline{\mathcal{C}})$, and $\widehat{\phi}_{L(\mathcal{C})} : L(\mathcal{C}) \longrightarrow L(\mathcal{C})$ is defined as follows: for any $\mathbf{v} \in L(\mathcal{C})$, $\widehat{\phi}_{L(\mathcal{C})}(\mathbf{v}) := \mathbf{v}'$ where $\mathbf{v}'(\underline{c}) := \widehat{\phi}[\mathbf{v}(\underline{c})]$ for all $\underline{c} \in \underline{\mathcal{C}}$.

Claim 3.2: *For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^{\circ}$, the diagram (9) commutes.*

Proof. Let $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C},\mathcal{X})$. By formula (A10), $U^{\mathcal{C}}_{\mathcal{X}}(\alpha) = \mathbf{v} \in L(\mathcal{C})$ is the function defined by $\mathbf{v}(\underline{c}) := v_{\mathcal{X}}[\underline{\alpha}(\underline{c})]$ for all $\underline{c} \in \underline{\mathcal{C}}$. Thus, $\widehat{\phi}_{L(\mathcal{C})} \circ U^{\mathcal{C}}_{\mathcal{X}}(\alpha) = \mathbf{v}'$ is the function defined:

$$\mathbf{v}'(\underline{c}) \quad = \quad \widehat{\phi} \circ v_{\mathcal{X}} \circ \underline{\alpha}(\underline{c}), \qquad \text{for all } \underline{c} \in \underline{\mathcal{C}}. \tag{A15}$$

Meanwhile, $\vec{\phi}(\alpha) = \phi \circ \alpha$, so $U_{\mathcal{Y}}^{\mathcal{C}} \circ \vec{\phi}(\alpha) = U_{\mathcal{Y}}^{\mathcal{C}}(\phi \circ \alpha) = \mathbf{v}''$ is the function defined by

$$\mathbf{v}''(\underline{c}) \quad = \quad v_{\mathcal{Y}} \circ \underline{\phi} \circ \underline{\alpha}(\underline{c}), \quad \text{for all } \underline{c} \in \underline{\mathcal{C}}. \tag{A16}$$

Applying equation (A14) to compare the right-hand sides of formulae (A15) and (A16), we see that $\mathbf{v}' = \mathbf{v}''$. In other words, $\widehat{\phi}_{L(\mathcal{C})} \circ U_{\mathcal{X}}^{\mathcal{C}}(\alpha) = U_{\mathcal{Y}}^{\mathcal{C}} \circ \vec{\phi}(\alpha)$. This holds for any $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$. Thus, $\widehat{\phi}_{L(\mathcal{C})} \circ U_{\mathcal{X}}^{\mathcal{C}} = U_{\mathcal{Y}}^{\mathcal{C}} \circ \vec{\phi}$. In other words, diagram (9) commutes. ▽ `Claim 3.2`

Define $A_{\mathcal{X},\mathcal{Y}}(\phi) := \widehat{\phi}$. Repeating this argument for all $\phi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, we obtain a function $A_{\mathcal{X},\mathcal{Y}} : \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) \longrightarrow \mathsf{Aff}$. Repeat this for all pairs of objects $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$. It is easily verified that these functions preserve composition. In other words: for any $\mathcal{X}, \mathcal{Y}, \mathcal{Z} \in \boldsymbol{\mathcal{X}}^\circ$ and $\phi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$ and $\psi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{Y}, \mathcal{Z})$, we have $A_{\mathcal{X},\mathcal{Z}}(\psi \circ \phi) = A_{\mathcal{Y},\mathcal{Z}}(\psi) \circ A_{\mathcal{X},\mathcal{Y}}(\phi)$. Thus, these functions together define a functor $A : \boldsymbol{\mathcal{X}} \Longmapsto \mathsf{Aff}$. ◇ `Claim 3`

Claim 4: *The affinity functor A maps all $\boldsymbol{\mathcal{X}}$-morphisms to contractions.*

Proof. Let $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, let $\phi \in \vec{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, and let $\widehat{\phi} = A(\phi)$. To see that $\widehat{\phi}$ is a contraction, recall that $\underline{\mathcal{X}}$ contains $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$-minimal and $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$-maximal constituents $\underline{o}_{\mathcal{X}}$ and $\underline{l}_{\mathcal{X}}$ such that $v_{\mathcal{X}}(\underline{o}_{\mathcal{X}}) = 0$ and $v_{\mathcal{X}}(\underline{l}_{\mathcal{X}}) = 1$, by $\mathsf{Boundedness}$. Likewise, $\underline{\mathcal{Y}}$ contains $\succcurlyeq_{\mathcal{Y}}^{\text{xp}}$-minimal and $\succcurlyeq_{\mathcal{Y}}^{\text{xp}}$-maximal constituents $\underline{o}_{\mathcal{Y}}$ and $\underline{l}_{\mathcal{Y}}$ such that $v_{\mathcal{Y}}(\underline{o}_{\mathcal{Y}}) = 0$ and $v_{\mathcal{Y}}(\underline{l}_{\mathcal{Y}}) = 1$. Thus, $\underline{o}_{\mathcal{Y}} \preccurlyeq_{\mathcal{Y}}^{\text{xp}} \underline{\phi}(\underline{o}_{\mathcal{X}}) \preccurlyeq_{\mathcal{Y}}^{\text{xp}} \underline{\phi}(\underline{l}_{\mathcal{X}}) \preccurlyeq_{\mathcal{Y}}^{\text{xp}} \underline{l}_{\mathcal{Y}}$. Thus,

$$0 \;=\; v_{\mathcal{Y}}(\underline{o}_{\mathcal{Y}}) \underset{(*)}{\leq} v_{\mathcal{Y}} \circ \underline{\phi}(\underline{o}_{\mathcal{X}}) \underset{(\dagger)}{=} \widehat{\phi} \circ v_{\mathcal{X}}(\underline{o}_{\mathcal{X}}) \underset{(*)}{\leq} \widehat{\phi} \circ v_{\mathcal{X}}(\underline{l}_{\mathcal{X}}) \underset{(\dagger)}{=} v_{\mathcal{Y}} \circ \underline{\phi}(\underline{l}_{\mathcal{X}}) \underset{(*)}{\leq} v_{\mathcal{Y}}(\underline{l}_{\mathcal{Y}}) \;=\; 1,$$

where each $(*)$ is by statement (A9) and both $(\dagger)$ are by equation (A14). Thus, $0 \leq \widehat{\phi}(0) \leq \widehat{\phi}(1) \leq 1$. Thus, $\widehat{\phi}(x) = B\,x + C$ for some $C \geq 0$ and $B \leq 1$. ◇ `Claim 4`

Claim 5: *For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, the function $U_{\mathcal{X}}^{\mathcal{C}} : \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X}) \longrightarrow \ell^\infty(\underline{\mathcal{C}})$ is mixture-preserving.*

Proof. Let $\widetilde{\alpha}, \widetilde{\gamma} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$, let $q \in [0, 1]$, and let $\widetilde{\beta} := q\,\widetilde{\alpha} + (1-q)\,\widetilde{\gamma}$. We must show that $U_{\mathcal{X}}^{\mathcal{C}}(\widetilde{\beta}) = q\,U_{\mathcal{X}}^{\mathcal{C}}(\widetilde{\alpha}) + (1-q)\,U_{\mathcal{X}}^{\mathcal{C}}(\widetilde{\gamma})$. So let $\underline{c} \in \underline{\mathcal{C}}$. According to defining formula (A10), we must show that $v_{\mathcal{X}} \circ \underline{\widetilde{\beta}}(\underline{c}) = q\,v_{\mathcal{X}} \circ \underline{\widetilde{\alpha}}(\underline{c}) + (1-q)\,v_{\mathcal{X}} \circ \underline{\widetilde{\gamma}}(\underline{c})$.

Now, let $\phi_1, \ldots, \phi_N \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$, and suppose that $\widetilde{\alpha} = (a_1, \phi_1;\ \ldots; a_N, \phi_N)$, $\widetilde{\beta} = (b_1, \phi_1;\ \ldots; b_N, \phi_N)$, and $\widetilde{\gamma} = (g_1, \phi_1;\ \ldots; g_N, \phi_N)$, where $\mathbf{a} = (a_1, \ldots, a_N)$, $\mathbf{b} = (b_1, \ldots, b_N)$, and $\mathbf{g} = (g_1, \ldots, g_N)$ are probability vectors (some of whose entries may be zero). Thus, for any $\underline{c} \in \mathcal{C}$, $\underline{\widetilde{\alpha}}(\underline{c}) = (a_1, \underline{\phi}_1(\underline{c});\ \ldots; a_N, \underline{\phi}_N(\underline{c}))$, $\underline{\widetilde{\beta}}(\underline{c}) = (b_1, \underline{\phi}_1(\underline{c});\ \ldots; b_N, \underline{\phi}_N(\underline{c}))$, and $\underline{\widetilde{\gamma}}(\underline{c}) = (g_1, \underline{\phi}_1(\underline{c});\ \ldots; g_N, \underline{\phi}_N(\underline{c}))$, where $\underline{\phi}_1(\underline{c}), \ldots, \underline{\phi}_N(\underline{c})$ are constituents of $\underline{\mathcal{X}}$, so that the expressions for $\underline{\widetilde{\alpha}}(\underline{c})$, $\underline{\widetilde{\beta}}(\underline{c})$ and $\underline{\widetilde{\gamma}}(\underline{c})$ represent elements of $\Delta(\underline{\mathcal{X}})$.

However, if $\widetilde{\beta} := q\,\widetilde{\alpha} + (1-q)\,\widetilde{\gamma}$, then $\mathbf{b} = q\,\mathbf{a} + (1-q)\,\mathbf{g}$. Thus, $\underline{\widetilde{\beta}}(\underline{c}) = q\,\underline{\widetilde{\alpha}}(\underline{c}) + (1-q)\,\underline{\widetilde{\gamma}}(\underline{c})$. Thus, since $v_{\mathcal{X}}$ is a mixture-preserving function on $\Delta(\underline{\mathcal{X}})$, we get $v_{\mathcal{X}} \circ \underline{\widetilde{\beta}}(\underline{c}) = v_{\mathcal{X}}\left(q\,\underline{\widetilde{\alpha}}(\underline{c}) + (1-q)\,\underline{\widetilde{\gamma}}(\underline{c})\right) = q\,v_{\mathcal{X}} \circ \underline{\widetilde{\alpha}}(\underline{c}) + (1-q)\,v_{\mathcal{X}} \circ \underline{\widetilde{\gamma}}(\underline{c})$, as desired.

This works for all $\underline{c} \in \underline{\mathcal{C}}$, so by (A10) we conclude that $U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\beta}) = q\, U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\alpha}) + (1-q)\, U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\gamma})$. ◇ Claim 5

For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, let $\mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$ be the image of $\overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ under $U^{\mathcal{C}}_{\mathcal{X}}$ (a subset of $\ell^\infty(\underline{\mathcal{C}})$). Let $\mathbf{0}_{\mathcal{C}}$ be the all-zero element of $\ell^\infty(\underline{\mathcal{C}})$, and let $\mathbf{1}_{\mathcal{C}}$ be the all-one element.

Claim 6: *For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, the image set $\mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$ is a convex subset of $\ell^\infty(\underline{\mathcal{C}})$, and contains $\mathbf{0}_{\mathcal{C}}$ and $\mathbf{1}_{\mathcal{C}}$.*

Proof. $\mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$ is convex by Claim 5, because $\overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ is convex. Recall that $\underline{\mathcal{X}}$ has constituents $\underline{o}$ and $\underline{l}$ such that $v_{\mathcal{X}}(\underline{o}) = 0$ and $v_{\mathcal{X}}(\underline{l}) = 1$. Thus, if $\mathbf{o}, \mathbf{l} \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X})$ are constant morphisms with the values $\underline{o}$ and $\underline{l}$, then $U^{\mathcal{C}}_{\mathcal{X}}(\mathbf{o}) = \mathbf{0}_{\mathcal{C}}$ and $U^{\mathcal{C}}_{\mathcal{X}}(\mathbf{l}) = \mathbf{1}_{\mathcal{C}}$. ◇ Claim 6

For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, let $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$ be the linear subspace of $\ell^\infty(\underline{\mathcal{C}})$ generated by $\mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$. Then $\mathbf{1}_{\mathcal{C}} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$, because $\mathbf{1}_{\mathcal{C}} \in \mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$ by Claim 6.

Claim 7: *Let $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, let $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, and suppose that $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) \neq \emptyset$. Then*

(a) *$\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$.*

(b) *If there is some $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$, then $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} = \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$.*

Proof. (a) Let $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$, and let $\widehat{\phi} = A(\phi)$. Then $\widehat{\phi}(x) = B\,x + C$ for some $C \geq 0$ and $B \leq 1$, by Claim 4. For any $\mathbf{a} \in \mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$, there exists $\widetilde{\alpha} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ such that $\mathbf{a} = U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\alpha})$. Then

$$B\,\mathbf{a} + C\mathbf{1}_{\mathcal{C}} \;=\; \widehat{\phi}_{L(\mathcal{C})}(\mathbf{a}) \;=\; \widehat{\phi}_{L(\mathcal{C})} \circ U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\alpha}) \underset{(*)}{=} U^{\mathcal{C}}_{\mathcal{Y}} \circ \vec{\phi}(\widetilde{\alpha}) \;=\; U^{\mathcal{C}}_{\mathcal{Y}}(\phi \circ \widetilde{\alpha}),$$

where $(*)$ is by Claim 3 and commuting diagram (9). Thus, $B\,\mathbf{a} + C\mathbf{1}_{\mathcal{C}}$ is an element of $\mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$, and hence, $\mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$. But $\mathbf{1}_{\mathcal{C}} \in \mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$ also. Thus, $\mathbf{a} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$.

This argument works for all $\mathbf{a} \in \mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$. Thus, $\mathcal{V}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$ and hence $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$.

(b) If $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$ is a split epimorphism in $\boldsymbol{\mathcal{C}}$, then the function $\vec{\phi} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{X}) \longrightarrow \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{Y})$ is surjective. Meanwhile, the function $U^{\mathcal{C}}_{\mathcal{Y}} : \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{Y}) \longrightarrow \mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$ is surjective, by definition. Thus, the composition $U^{\mathcal{C}}_{\mathcal{Y}} \circ \vec{\phi}$ is surjective onto $\mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$. By Claim 3 and commuting diagram (9), we deduce that $\widehat{\phi}_{L(\mathcal{C})} \circ U^{\mathcal{C}}_{\mathcal{X}}$ must also be surjective onto $\mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$, which means $\widehat{\phi}_{L(\mathcal{C})} : \mathcal{V}^{\mathcal{C}}_{\mathcal{X}} \longrightarrow \mathcal{V}^{\mathcal{C}}_{\mathcal{Y}}$ is surjective. But $\widehat{\phi}_{L(\mathcal{C})}(\mathcal{V}^{\mathcal{C}}_{\mathcal{X}}) \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$, because $\widehat{\phi}$ is an affine transformation and $\mathbf{1}_{\mathcal{C}} \in \mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$. Thus, $\mathcal{V}^{\mathcal{C}}_{\mathcal{Y}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$, which means $\mathcal{W}^{\mathcal{C}}_{\mathcal{Y}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$. Since we have already established the reverse inclusion in part (a), we conclude that $\mathcal{W}^{\mathcal{C}}_{\mathcal{Y}} = \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$. ◇ Claim 7

Claim 8: *For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\mathcal{X}, \mathcal{Z} \in \boldsymbol{\mathcal{X}}$, if $\mathcal{X} \lhd \mathcal{Z}$, then $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$.*

Proof. Recall that $\rhd$ is the transitive closure of the relation $\rightsquigarrow$. Thus, it suffices to show that, if $\mathcal{X} \rightsquigarrow \mathcal{Z}$, then $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$. But if $\mathcal{X} \rightsquigarrow \mathcal{Z}$, then either (1) $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Z}) \neq \emptyset$, *or* there is some $\pi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{Z}, \mathcal{X})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$. In case (1), Claim 7(a) says that $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$, while in case (2), Claim 7(b) says that $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} = \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$. ◇ Claim 8

Now, for all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, define

$$L'(\mathcal{C}) \quad := \quad \bigcup_{\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ} \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}. \tag{A17}$$

Claim 9: *$L'(\mathcal{C})$ is a linear subspace of $\ell^\infty(\underline{\mathcal{C}})$, and contains $\mathbf{1}_{\mathcal{C}}$.*

Proof. (a) First note that $\mathbf{1}_{\mathcal{C}} \in L'(\mathcal{C})$ because $\mathbf{1}_{\mathcal{C}} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$ for all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, by Claim 6.

Now let $\mathbf{v}, \mathbf{w} \in L'(\mathcal{C})$ and $r \in \mathbb{R}$. Then defining formula (A17) yields $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$ such that $\mathbf{v} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$ and $\mathbf{w} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{Y}}$. By condition (XD), there is some $\mathcal{Z} \in \boldsymbol{\mathcal{X}}^\circ$ such that $\mathcal{X} \lhd \mathcal{Z}$ and $\mathcal{Y} \lhd \mathcal{Z}$. Thus, Claim 8 says that $\mathcal{W}^{\mathcal{C}}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$ and $\mathcal{W}^{\mathcal{C}}_{\mathcal{Y}} \subseteq \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$. Thus, $\mathbf{v}, \mathbf{w} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$, so that $\mathbf{v} + r\,\mathbf{w} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$ (because $\mathcal{W}^{\mathcal{C}}_{\mathcal{Z}}$ is a linear subspace). Thus, $\mathbf{v} + r\,\mathbf{w} \in L'(\mathcal{C})$. This works for all $\mathbf{v}, \mathbf{w} \in L'(\mathcal{C})$ and $r \in \mathbb{R}$, so $L'(\mathcal{C})$ is a linear subspace of $\ell^\infty(\underline{\mathcal{C}})$. $\diamond$ `Claim 9`

Claim 10: *L' is a utility subframe of L.*

Proof. For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, Claim 9 says that $L'(\mathcal{C})$ is a linear subspace of $L(\mathcal{C})$, containing the order unit $\mathbf{1}_{\mathcal{C}}$. It becomes a unitary partially ordered vector space by restricting the order from $L(\mathcal{C})$.

Now let $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and let $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, so that $\phi^* := L(\phi)$ is a UPOVS morphism from $L(\mathcal{C})$ to $L(\mathcal{B})$. We must show that ϕ^* maps $L'(\mathcal{C})$ into $L'(\mathcal{B})$.

Let $\mathbf{w} \in L'(\mathcal{C})$. Then defining formula (A17) says that $\mathbf{w} \in \mathcal{W}^{\mathcal{C}}_{\mathcal{X}}$ for some $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. Thus, $\mathbf{w} = \sum_{n=1}^{N} r_n\,\mathbf{v}_n$ for some $\mathbf{v}_1, \ldots, \mathbf{v}_N \in \mathcal{V}^{\mathcal{C}}_{\mathcal{X}}$ and $r_1, \ldots, r_N \in \mathbb{R}$. For all $n \in [1 \ldots N]$, let $\widetilde{\alpha}_n \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{C}, \mathcal{X})$ be such that $U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\alpha}_n) = \mathbf{v}_n$. Now define $\widetilde{\beta}_n := \overleftarrow{\phi}(\widetilde{\alpha}_n) = \widetilde{\alpha}_n \circ \phi$; so $\widetilde{\beta}_n \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{B}, \mathcal{X})$. Then

$$U^{\mathcal{B}}_{\mathcal{X}}(\widetilde{\beta}_n) \quad = \quad U^{\mathcal{B}}_{\mathcal{X}} \circ \overleftarrow{\phi}(\widetilde{\alpha}_n) \quad \underset{(*)}{=} \quad \phi^* \circ U^{\mathcal{C}}_{\mathcal{X}}(\widetilde{\alpha}_n) \quad = \quad \phi^*(\mathbf{v}_n), \tag{A18}$$

where $(*)$ is by Claim 2 and commuting diagram (1) (with $\mathcal{C}_1 := \mathcal{B}$ and $\mathcal{C}_2 := \mathcal{C}$). Thus,

$$\phi^*(\mathbf{w}) \quad = \quad \phi^*\left(\sum_{n=1}^{N} r_n\,\mathbf{v}_n\right) \quad = \quad \sum_{n=1}^{N} r_n\,\phi^*(\mathbf{v}_n) \quad \underset{(*)}{=} \quad \sum_{n=1}^{N} r_n\,U^{\mathcal{B}}_{\mathcal{X}}(\widetilde{\beta}_n), \tag{A19}$$

where $(*)$ is by equation (A18). The right hand side of (A19) is a linear combination of elements of $\mathcal{V}^{\mathcal{B}}_{\mathcal{X}}$, and hence an element of $\mathcal{W}^{\mathcal{B}}_{\mathcal{X}}$. Thus, it is an element of $L'(\mathcal{B})$, as desired.

Thus, for any $\phi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$, we can define $L'(\phi)$ to be the restriction of $L(\phi)$ to $L'(\mathcal{C})$ to obtain a map $L'(\phi) : L'(\mathcal{C}) \longrightarrow L'(\mathcal{B})$. Since this is a restriction of a uniferent, order-preserving linear map, it is itself a uniferent, order-preserving linear map. Thus, L' is a utility frame. $\diamond$ `Claim 10`

Having established that L' is a utility frame, we shall from now on regard U as a utility system taking values in L'.

Claim 11: *Let $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, and $\widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$. If $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) = U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta})$, then $\widetilde{\alpha} \approx_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta}$.*

Proof. If $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) = U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta})$, then for all $\underline{s} \in \underline{\mathcal{S}}$, we have $v_{\mathcal{X}} \circ \underline{\widetilde{\alpha}}(\underline{s}) = v_{\mathcal{X}} \circ \underline{\widetilde{\beta}}(\underline{s})$ and thus $\underline{\widetilde{\alpha}}(\underline{s}) \approx_{\mathcal{X}}^{\text{xp}} \underline{\widetilde{\beta}}(\underline{s})$. Thus, Statewise Dominance implies that $\widetilde{\alpha} \approx_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta}$. ◇ Claim 11

For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, we define a weak order $\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$ as follows: for any $\mathbf{a}, \mathbf{b} \in \mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$, stipulate that

$$\left(\mathbf{a} \,\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}\, \mathbf{b}\right) \Leftrightarrow \left(\widetilde{\alpha} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta} \text{ for some } \widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X}) \text{ with } \mathbf{a} = U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) \text{ and } \mathbf{b} = U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta})\right). \tag{A20}$$

Claim 11 implies that this is a well-defined weak order on $\mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$.

Claim 12: *For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, there is a mixture-preserving function $\rho_{\mathcal{X}}^{\mathcal{S}} : \mathcal{V}_{\mathcal{X}}^{\mathcal{S}} \longrightarrow \mathbb{R}$ that represents $\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}$. It is unique up to positive affine transformations.*

Proof. $U_{\mathcal{X}}^{\mathcal{S}}$ is mixture-preserving by Claim 5, and $\succcurlyeq_{\mathcal{X}}^{\mathcal{S}}$ satisfies the vNM independence and Continuity on $\overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$. Thus, $\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}$ satisfies the analogous axioms on $\mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$, which is convex by Claim 6. Now apply Proposition A.1. ◇ Claim 12

Claim 13: *For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$ and $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, $\rho_{\mathcal{X}}^{\mathcal{S}}$ is nondecreasing in every coordinate.*

Proof. Let $\mathbf{a}, \mathbf{b} \in \mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$, and suppose that $a_{\underline{s}} \geq b_{\underline{s}}$ for all $\underline{s} \in \mathcal{S}$. Let $\widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$, such that $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) = \mathbf{a}$ and $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta}) = \mathbf{b}$. Then for all $\underline{s} \in \mathcal{S}$, we have $v_{\mathcal{X}} \circ \underline{\widetilde{\alpha}}(\underline{s}) \geq v_{\mathcal{X}} \circ \underline{\widetilde{\beta}}(\underline{s})$, hence $\underline{\widetilde{\alpha}}(\underline{s}) \succcurlyeq_{\mathcal{X}}^{\text{xp}} \underline{\widetilde{\beta}}(\underline{s})$ (because $v_{\mathcal{X}}$ represents $\succcurlyeq_{\mathcal{X}}^{\text{xp}}$, by statement (A9)). Thus, $\widetilde{\alpha} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta}$, by Statewise Dominance. Thus, $\mathbf{a} \,\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}\, \mathbf{b}$, by defining formula (A20). Thus, $\rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{a}) \geq \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{b})$, by Claim 12. ◇ Claim 13

By applying a positive affine transformation if necessary, we can assume without loss of generality that $\rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{0}_{\mathcal{S}}) = 0$ and $\rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) = 1$. Recall that $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$ is the linear subspace of $\ell^\infty(\underline{\mathcal{S}})$ generated by $\mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$. Thus, $\rho_{\mathcal{X}}^{\mathcal{S}}$ extends to a unique linear functional on all of $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$.

Claim 14: *Let $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, let $\mathcal{X}, \mathcal{Y} \in \boldsymbol{\mathcal{X}}^\circ$, and suppose that $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y}) \neq \emptyset$. Then*

(a) *$\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Y}}^{\mathcal{S}}$, and $\rho_{\mathcal{Y}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$.*

(b) *If there is some $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Y})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$, then $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} = \mathcal{W}_{\mathcal{Y}}^{\mathcal{S}}$ and $\rho_{\mathcal{Y}}^{\mathcal{S}} = \rho_{\mathcal{X}}^{\mathcal{S}}$.*

Proof. (a) $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Y}}^{\mathcal{S}}$ by Claim 7(a). Let $\widehat{\phi} := A(\phi)_{L(\mathcal{S})}$. Let $\mathbf{a}, \mathbf{b} \in \mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$. Let $\mathbf{a}' := \widehat{\phi}(\mathbf{a})$ and $\mathbf{b}' := \widehat{\phi}(\mathbf{b})$. Find $\widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$ such that $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\alpha}) = \mathbf{a}$ and $U_{\mathcal{X}}^{\mathcal{S}}(\widetilde{\beta}) = \mathbf{b}$. Let

$\widetilde{\alpha}' := \overleftarrow{\phi}(\widetilde{\alpha}) = \phi \circ \widetilde{\alpha}$ and $\widetilde{\beta}' := \overleftarrow{\phi}(\widetilde{\beta}) = \phi \circ \widetilde{\beta}$. Then $\mathbf{a}' = U_{\mathcal{Y}}^{\mathcal{S}}(\widetilde{\alpha}')$ and $\mathbf{b}' = U_{\mathcal{Y}}^{\mathcal{S}}(\widetilde{\beta}')$, by Claim 3 and commuting diagram (9). Thus,

$$\begin{aligned} \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{a}) \geq \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{b}) &\underset{(*)}{\Longleftrightarrow} \mathbf{a} \mathrel{\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}} \mathbf{b} \underset{(\dagger)}{\Longleftrightarrow} \widetilde{\alpha} \succcurlyeq_{\mathcal{X}}^{\mathcal{S}} \widetilde{\beta} \underset{(\diamond)}{\Longleftrightarrow} \phi \circ \widetilde{\alpha} \succcurlyeq_{\mathcal{Y}}^{\mathcal{S}} \phi \circ \widetilde{\beta} \\ &\underset{(\flat)}{\Longleftrightarrow} \widetilde{\alpha}' \succcurlyeq_{\mathcal{Y}}^{\mathcal{S}} \widetilde{\beta}' \underset{(\dagger)}{\Longleftrightarrow} \mathbf{a}' \mathrel{\widehat{\succcurlyeq}_{\mathcal{Y}}^{\mathcal{S}}} \mathbf{b}' \underset{(*)}{\Longleftrightarrow} \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{a}') \geq \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{b}') \\ &\underset{(\sharp)}{\Longleftrightarrow} \rho_{\mathcal{Y}}^{\mathcal{S}} \circ \widehat{\phi}(\mathbf{a}) \geq \rho_{\mathcal{Y}}^{\mathcal{S}} \circ \widehat{\phi}(\mathbf{b}). \end{aligned} \tag{A21}$$

Here, both $(*)$ are by Claim 12, both $(\dagger)$ are by defining formula (A20), $(\diamond)$ is by (TP), $(\flat)$ is by the definitions of $\widetilde{\alpha}'$ and $\widetilde{\beta}'$, and $(\sharp)$ is by the definitions of $\mathbf{a}'$ and $\mathbf{b}'$.

Finally, recall that $\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{0}_{\mathcal{S}}) = 0$ and $\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) = 1$. Suppose $A(\phi) = B\,x + C$, for some $B > 0$ and $C \in \mathbb{R}$. Then $\widehat{\phi}(\mathbf{v}) = B\,\mathbf{v} + C\,\mathbf{1}_{\mathcal{S}}$ for all $\mathbf{v} \in L(\mathcal{S})$. Then $\rho_{\mathcal{Y}}^{\mathcal{S}}(\widehat{\phi}(\mathbf{v})) = \rho_{\mathcal{Y}}^{\mathcal{S}}(B\,\mathbf{v} + C\mathbf{1}_{\mathcal{S}}) = B\,\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{v}) + C\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) = B\,\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{v}) + C$. Thus, we get:

$$\rho_{\mathcal{Y}}^{\mathcal{S}} \circ \widehat{\phi}(\mathbf{a}) \geq \rho_{\mathcal{Y}}^{\mathcal{S}} \circ \widehat{\phi}(\mathbf{b}) \;\Leftrightarrow\; B\,\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{a}) + C \;\geq\; B\,\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{b}) + C \;\Leftrightarrow\; \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{a}) \;\geq\; \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{b}). \tag{A22}$$

Combining (A21) and (A22), we deduce that $\rho_{\mathcal{Y}}^{\mathcal{S}}$, restricted to $\mathcal{V}_{\mathcal{X}}^{\mathcal{S}}$, is another mixture-preserving utility representation of $\widehat{\succcurlyeq}_{\mathcal{X}}^{\mathcal{S}}$. The uniqueness part of Proposition A.1 implies that $\rho_{\mathcal{Y}}^{\mathcal{S}}$ is a positive affine transformation of $\rho_{\mathcal{X}}^{\mathcal{S}}$. But $\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{0}_{\mathcal{S}}) = 0 = \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{0}_{\mathcal{S}})$ and $\rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) = 1 = \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{1}_{\mathcal{S}})$, so we conclude that $\rho_{\mathcal{Y}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$.

(b) $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} = \mathcal{W}_{\mathcal{Y}}^{\mathcal{S}}$ by Claim 7(b). Thus, part (a) implies that $\rho_{\mathcal{Y}}^{\mathcal{S}} = \rho_{\mathcal{X}}^{\mathcal{S}}$. ◇ Claim 14

Claim 15: *For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$ and $\mathcal{X}, \mathcal{Z} \in \boldsymbol{\mathcal{X}}$, if $\mathcal{X} \lhd \mathcal{Z}$, then $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Z}}^{\mathcal{S}}$, and $\rho_{\mathcal{Z}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$.*

Proof. $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Z}}^{\mathcal{S}}$ by Claim 8. Recall that $\rhd$ is the transitive closure of the relation $\rightsquigarrow$. Thus, it suffices to show that, if $\mathcal{X} \rightsquigarrow \mathcal{Z}$, then $\rho_{\mathcal{Z}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$.

If $\mathcal{X} \rightsquigarrow \mathcal{Z}$, then either (1) $\overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{X}, \mathcal{Z}) \neq \emptyset$, *or* (2) there is some $\phi \in \overrightarrow{\boldsymbol{\mathcal{X}}}(\mathcal{Z}, \mathcal{X})$ that is a split epimorphism in $\boldsymbol{\mathcal{C}}$. In case (1), Claim 14(a) says that $\rho_{\mathcal{Z}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$. In case (2), Claim 14(b) says that $\rho_{\mathcal{Z}}^{\mathcal{S}} = \rho_{\mathcal{X}}^{\mathcal{S}}$. ◇ Claim 15

For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, define $\rho_{\mathcal{S}} : L'(\mathcal{S}) \longrightarrow \mathbb{R}$ as follows. For any $\mathbf{v} \in L'(\mathcal{S})$, find some $\mathcal{X} \in \boldsymbol{\mathcal{X}}^{\circ}$ such that $\mathbf{v} \in \mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$, and then define $\rho_{\mathcal{S}}(\mathbf{v}) := \rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{v})$.

Claim 16: *For all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^{\circ}$, the function $\rho_{\mathcal{S}}$ is well-defined, and is a uniferent, order-preserving linear functional on $L'(\mathcal{S})$.*

Proof. (*Well-defined*) Let $\mathbf{v} \in L'(\mathcal{S})$. Defining formula (A17) yields some $\mathcal{X} \in \boldsymbol{\mathcal{X}}^{\circ}$ such that $\mathbf{v} \in \mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$. Suppose there is some other $\mathcal{Y} \in \boldsymbol{\mathcal{X}}^{\circ}$ such that $\mathbf{v} \in \mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$ also. We must show that $\rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{v}) = \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{v})$. By condition (XD) there is some $\mathcal{Z} \in \boldsymbol{\mathcal{X}}^{\circ}$ such that $\mathcal{X} \lhd \mathcal{Z}$ and $\mathcal{Y} \lhd \mathcal{Z}$. Thus, Claim 15 says that $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Z}}^{\mathcal{S}}$ and $\mathcal{W}_{\mathcal{Y}}^{\mathcal{S}} \subseteq \mathcal{W}_{\mathcal{Z}}^{\mathcal{S}}$. Furthermore, $\rho_{\mathcal{Z}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{X}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{X}}^{\mathcal{S}}$ and $\rho_{\mathcal{Z}}^{\mathcal{S}}$ agrees with $\rho_{\mathcal{Y}}^{\mathcal{S}}$ on $\mathcal{W}_{\mathcal{Y}}^{\mathcal{S}}$. Thus, $\mathbf{v} \in \mathcal{W}_{\mathcal{Z}}^{\mathcal{S}}$, and we have $\rho_{\mathcal{X}}^{\mathcal{S}}(\mathbf{v}) = \rho_{\mathcal{Z}}^{\mathcal{S}}(\mathbf{v}) = \rho_{\mathcal{Y}}^{\mathcal{S}}(\mathbf{v})$, as desired.

(*Linear*) Let $\mathbf{v}, \mathbf{w} \in L'(\mathcal{S})$ and $r \in \mathbb{R}$. Just as in the proof of Claim 9, there is some $\mathcal{Z} \in \boldsymbol{\mathcal{X}}^\circ$ such that $\mathbf{v}, \mathbf{w} \in \mathcal{W}^{\mathcal{S}}_{\mathcal{Z}}$, so that $\mathbf{v} + r\,\mathbf{w} \in \mathcal{W}^{\mathcal{S}}_{\mathcal{Z}}$. Thus, by the definition of $\rho_{\mathcal{S}}$, we have $\rho_{\mathcal{S}}(\mathbf{v} + r\,\mathbf{w}) = \rho^{\mathcal{S}}_{\mathcal{Z}}(\mathbf{v} + r\,\mathbf{w}) = \rho^{\mathcal{S}}_{\mathcal{Z}}(\mathbf{v}) + r\,\rho^{\mathcal{S}}_{\mathcal{Z}}(\mathbf{w}) = \rho_{\mathcal{S}}(\mathbf{v}) + r\,\rho_{\mathcal{S}}(\mathbf{w})$.

(*Order-preserving*) Let $\mathbf{v}, \mathbf{w} \in L'(\mathcal{S})$, with $\mathbf{v} < \mathbf{w}$. Just as in the proof of Claim 9, there is some $\mathcal{Z} \in \boldsymbol{\mathcal{X}}^\circ$ such that $\mathbf{v}, \mathbf{w} \in \mathcal{W}^{\mathcal{S}}_{\mathcal{Z}}$. Thus, by the definition of $\rho_{\mathcal{S}}$, we have $\rho_{\mathcal{S}}(\mathbf{v}) = \rho^{\mathcal{S}}_{\mathcal{Z}}(\mathbf{v}) \leq \rho^{\mathcal{S}}_{\mathcal{Z}}(\mathbf{v}) = \rho_{\mathcal{S}}(\mathbf{v})$.

(*Uniferent*) For any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, we have $\mathbf{1}_{\mathcal{S}} \in \mathcal{V}^{\mathcal{S}}_{\mathcal{X}} \subset \mathcal{W}^{\mathcal{S}}_{\mathcal{X}}$ (by Claim 6), and thus $\rho_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) = \rho^{\mathcal{S}}_{\mathcal{X}}(\mathbf{1}_{\mathcal{S}}) = 1$. ◇ `Claim 16`

Claim 17: *The collection $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ is a belief structure.*

Proof. Let $L'_{\restriction\boldsymbol{\mathcal{S}}}$ be the restriction of L' to a contravariant functor $L'_{\restriction\boldsymbol{\mathcal{S}}} : \boldsymbol{\mathcal{S}}^{\mathrm{op}} \Longmapsto \mathsf{UPOVS}$. We must show that $\{\rho_{\mathcal{S}}\}_{\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ}$ is a co-cone to $\mathbb{R}$ from this functor . In other words, for all $\mathcal{S}_1, \mathcal{S}_2 \in \boldsymbol{\mathcal{S}}^\circ$ and $\phi \in \overrightarrow{\boldsymbol{\mathcal{S}}}(\mathcal{S}_1, \mathcal{S}_2)$, we must show that the diagram (4) commutes.

Let $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$; then $\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}} \subseteq \mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}} \subseteq L'(\mathcal{S}_2)$, by defining formula (A17). So, let $\mathbf{a}_2, \mathbf{b}_2 \in \mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}$; then $\mathbf{a}_2 = U^{\mathcal{S}_2}_{\mathcal{X}}(\widetilde{\alpha}_2)$ and $\mathbf{b}_2 = U^{\mathcal{S}_2}_{\mathcal{X}}(\widetilde{\beta}_2)$ for some $\widetilde{\alpha}_2, \widetilde{\beta}_2 \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}_2, \mathcal{X})$. Let $\widetilde{\alpha}_1 := \overleftarrow{\phi}(\widetilde{\alpha}_2) = \widetilde{\alpha}_2 \circ \phi$ and $\widetilde{\beta}_1 := \overleftarrow{\phi}(\widetilde{\beta}_2) = \widetilde{\beta}_2 \circ \phi$. Then $\widetilde{\alpha}_1, \widetilde{\beta}_1 \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}_1, \mathcal{X})$. Let $\mathbf{a}_1 := U^{\mathcal{S}_1}_{\mathcal{X}}(\widetilde{\alpha}_1)$ and $\mathbf{b}_1 := U^{\mathcal{S}_1}_{\mathcal{X}}(\widetilde{\beta}_1)$. For notational simplicity, let $\phi^* := L'(\phi)$. Then $\phi^* : L'(\mathcal{S}_2) \longrightarrow L'(\mathcal{S}_1)$ is a linear function, and $\mathbf{a}_1 = \phi^*(\mathbf{a}_2)$ and $\mathbf{b}_1 = \phi^*(\mathbf{b}_2)$, by Claim 2 and diagram (1). Thus,

$$\begin{aligned} \rho_{\mathcal{S}_1} \circ \phi^*(\mathbf{a}_2) \geq \rho_{\mathcal{S}_1} \circ \phi^*(\mathbf{b}_2) \quad &\iff \quad \rho_{\mathcal{S}_1}(\mathbf{a}_1) \geq \rho_{\mathcal{S}_1}(\mathbf{b}_1) \quad \underset{(*)}{\iff} \quad \mathbf{a}_1 \mathbin{\widehat{\succcurlyeq}^{\mathcal{S}_1}_{\mathcal{X}}} \mathbf{b}_1 \\ &\underset{(\dagger)}{\iff} \quad \widetilde{\alpha}_1 \succcurlyeq^{\mathcal{S}_1}_{\mathcal{X}} \widetilde{\beta}_1 \quad \underset{(\diamond)}{\iff} \quad \widetilde{\alpha}_2 \succcurlyeq^{\mathcal{S}_2}_{\mathcal{X}} \widetilde{\beta}_2 \\ &\underset{(\dagger)}{\iff} \quad \mathbf{a}_2 \mathbin{\widehat{\succcurlyeq}^{\mathcal{S}_2}_{\mathcal{X}}} \mathbf{b}_2 \quad \underset{(*)}{\iff} \quad \rho_{\mathcal{S}_2}(\mathbf{a}_2) \geq \rho_{\mathcal{S}_2}(\mathbf{b}_2). \end{aligned}$$

Here, both $(*)$ are by Claim 12, both $(\dagger)$ are by defining formula (A20), and $(\diamond)$ is by property (BP).

Thus, $(\rho_{\mathcal{S}_1} \circ \phi^*)_{\restriction\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}}$ and $(\rho_{\mathcal{S}_2})_{\restriction\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}}$ are both mixture-preserving utility representations for $\widehat{\succcurlyeq}^{\mathcal{S}_2}_{\mathcal{X}}$, so by the uniqueness part of Proposition A.1 they are equal up to positive affine transformation. Recall that $\rho_{\mathcal{S}_1}$ and $\rho_{\mathcal{S}_2}$ are normalized so that $\rho_{\mathcal{S}_1}(\mathbf{0}_{\underline{\mathcal{S}}_1}) = 0$ and $\rho_{\mathcal{S}_1}(\mathbf{1}_{\underline{\mathcal{S}}_1}) = 1$, while $\rho_{\mathcal{S}_2}(\mathbf{0}_{\underline{\mathcal{S}}_2}) = 0$ and $\rho_{\mathcal{S}_2}(\mathbf{1}_{\underline{\mathcal{S}}_2}) = 1$. However, $\phi^*(\mathbf{0}_{\underline{\mathcal{S}}_2}) = \mathbf{0}_{\underline{\mathcal{S}}_1}$ (because ϕ^* is linear), and $\phi^*(\mathbf{1}_{\underline{\mathcal{S}}_2}) = \mathbf{1}_{\underline{\mathcal{S}}_1}$ (because $\phi^*(\mathbf{v}) = \mathbf{v} \circ \phi$ for any $\mathbf{v} \in L(\mathcal{S}_2)$). Thus, $\rho_{\mathcal{S}_1} \circ \phi^*(\mathbf{0}_{\mathcal{S}_2}) = \rho_{\mathcal{S}_1}(\mathbf{0}_{\underline{\mathcal{S}}_1}) = 0$ and $\rho_{\mathcal{S}_1} \circ \phi^*(\mathbf{1}_{\mathcal{S}_2}) = \rho_{\mathcal{S}_1}(\mathbf{1}_{\underline{\mathcal{S}}_1}) = 0$. It follows that $(\rho_{\mathcal{S}_1} \circ \phi^*)_{\restriction\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}} = (\rho_{\mathcal{S}_2})_{\restriction\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}}$.

Now, $(\rho_{\mathcal{S}_2})_{\restriction\mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}}}$ is the unique extension of $(\rho_{\mathcal{S}_2})_{\restriction\mathcal{V}^{\mathcal{S}_2}_{\mathcal{X}}}$ to $\mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}}$, and likewise for $(\rho_{\mathcal{S}_1} \circ \phi^*)_{\restriction\mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}}}$. Thus, we deduce that $(\rho_{\mathcal{S}_1} \circ \phi^*)_{\restriction\mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}}} = (\rho_{\mathcal{S}_2})_{\restriction\mathcal{W}^{\mathcal{S}_2}_{\mathcal{X}}}$. This argument works for any $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. Thus, by defining formula (A17), we conclude that $\rho_{\mathcal{S}_1} \circ \phi^* = \rho_{\mathcal{S}_2}$. Thus, the diagram (4) commutes, as desired. ◇ `Claim 17`

It remains to show that this data yields a local SEU representation (3) for each preference order in $\underline{\trianglerighteq}^{\mathrm{xa}}$. So, let $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$, and $\widetilde{\alpha}, \widetilde{\beta} \in \overrightarrow{\Delta\mathcal{C}}(\mathcal{S}, \mathcal{X})$. Then

$$\widetilde{\alpha} \succcurlyeq^{\mathcal{S}}_{\mathcal{X}} \widetilde{\beta} \quad \underset{(*)}{\iff} \quad U^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\alpha}) \mathbin{\widehat{\succcurlyeq}^{\mathcal{S}}_{\mathcal{X}}} U^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\beta}) \quad \underset{(\diamond)}{\iff} \quad \rho_{\mathcal{S}}[U^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\alpha})] \geq \rho_{\mathcal{S}}[U^{\mathcal{S}}_{\mathcal{X}}(\widetilde{\beta})], \qquad \text{(A23)}$$

as desired. Here, $(*)$ is by the defining formula (A20), while and $(\diamond)$ is by Claim 12. □

Remark. Statements (A8) and (A23) are two vNM utility representations for the same preference order $\succcurlyeq^{\mathcal{S}}_{\mathcal{X}}$. By the uniqueness part of Proposition A.1, $V^{\mathcal{S}}_{\mathcal{X}}$ must be a positive affine transformation of $\rho_{\mathcal{S}} \circ U^{\mathcal{S}}_{\mathcal{X}}$. Since we have normalized these functions so that $V^{\mathcal{S}}_{\mathcal{X}}(\mathbf{0}_{\mathcal{S}}) = 0 = \rho_{\mathcal{S}} \circ U^{\mathcal{S}}_{\mathcal{X}}(\mathbf{0}_{\mathcal{S}})$ and $V^{\mathcal{S}}_{\mathcal{X}}(\mathbf{1}_{\mathcal{S}}) = 0 = \rho_{\mathcal{S}} \circ U^{\mathcal{S}}_{\mathcal{X}}(\mathbf{1}_{\mathcal{S}})$, we conclude that in fact, $V^{\mathcal{S}}_{\mathcal{X}} = \rho_{\mathcal{S}} \circ U^{\mathcal{S}}_{\mathcal{X}}$. But this fact is never used in the above proof.

Proof of Proposition 7.3. From Theorem 7.1, we obtain a global SEU representation. Let $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. In the first part of the proof of Theorems 7.1 and 8.1, we obtained a mixture-preserving function $v_{\mathcal{X}} : \underline{\Delta\mathcal{X}} \longrightarrow \mathbb{R}$ satisfying formula (A9). We then defined the utility functional $U_{\mathcal{X}}$ by equation (A10). Let $u_{\mathcal{X}}$ be the restriction of $v_{\mathcal{X}}$ to $\underline{\mathcal{X}}$.

Claim 1: *$u_{\mathcal{X}}(\underline{\mathcal{X}}) = v_{\mathcal{X}}(\underline{\Delta\mathcal{X}})$, which is an interval in $\mathbb{R}$.*

Proof. $u_{\mathcal{X}}(\underline{\mathcal{X}}) = v_{\mathcal{X}}(\underline{\Delta\mathcal{X}})$ by Certainty Equivalents. Meanwhile, $v_{\mathcal{X}}(\underline{\Delta\mathcal{X}})$ is an interval because $\underline{\Delta\mathcal{X}}$ is a convex set and $v_{\mathcal{X}}$ is mixture-preserving. ◇ Claim 1

By hypothesis, $\mathcal{X}$ is second-countable. Thus, given Ex post Continuity, Debreu's Theorem says that there is a continuous function $w_{\mathcal{X}} : \mathcal{X} \longrightarrow \mathbb{R}$ representing $\succcurlyeq^{\text{xp}}_{\mathcal{X}}$.[33] Let $\mathcal{W} := w_{\mathcal{X}}(\underline{\mathcal{X}}) \subseteq \mathbb{R}$. Since $u_{\mathcal{X}}$ and $w_{\mathcal{X}}$ both represent $\succcurlyeq^{\text{xp}}_{\mathcal{X}}$, there is an increasing function $\phi : \mathcal{W} \longrightarrow \mathbb{R}$ such that $u_{\mathcal{X}} = \phi \circ w_{\mathcal{X}}$.

Claim 2: *ϕ is continuous.*

Proof. (by contradiction) Suppose ϕ is *not* continuous. Any discontinuity of ϕ is a jump discontinuity, because ϕ is increasing. Any jump discontinuity induces a gap in the image $\phi(\mathcal{W})$. But $\phi(\mathcal{W}) = \phi \circ w_{\mathcal{X}}(\underline{\mathcal{X}}) = u_{\mathcal{X}}(\underline{\mathcal{X}})$, which is an interval by Claim 1. Contradiction. ◇ Claim 2

We conclude that $u_{\mathcal{X}}$ is also continuous, so that $U_{\mathcal{X}}$ is as described in Example 6.4(c). Thus, for all $\mathcal{S} \in \mathsf{Top}^\circ$ and all $\alpha \in \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X})$, the composite $u_{\mathcal{X}} \circ \alpha$ is continuous and bounded. In other words $U^{\mathcal{S}}_{\mathcal{X}}(\alpha)$ is continuous and bounded. Thus, $U^{\mathcal{S}}_{\mathcal{X}}$ maps all of $\overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X})$ into a subset of $\mathfrak{C}_{\text{b}}(\mathcal{S})$.

By the Monotone Path Property, there is a space $\mathcal{X}_* \in \boldsymbol{\mathcal{X}}^\circ$ and a continuous function $\gamma : [0,1] \longrightarrow \mathcal{X}_*$ such that $\gamma(s) \prec^{\mathcal{X}_*}_{\text{xp}} \gamma(t)$ whenever $s < t$. Let $v := u_{\mathcal{X}_*} \circ \gamma : [0,1] \longrightarrow \mathbb{R}$. Then v is continuous and strictly increasing. Let $a := v(0)$ and $b := v(1)$. Then $v : [0,1] \longrightarrow [a,b]$ is a homeomorphism. For any $\mathcal{S} \in \mathsf{Top}^\circ$, let $\mathfrak{C}(\mathcal{S}, [a,b])$ be the set of all continuous functions from $\mathcal{S}$ into $[a,b]$ (a subset of $\mathfrak{C}_{\text{b}}(\mathcal{S})$).

Claim 3: *For all $\mathcal{S} \in \mathsf{Top}^\circ$, the set $\{U^{\mathcal{S}}_{\mathcal{X}_*}(\alpha);\ \alpha \in \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X}_*)\}$ contains $\mathfrak{C}(\mathcal{S}, [a,b])$.*

Proof. Let $f \in \mathfrak{C}(\mathcal{S}, [a,b])$. Let $\alpha := \gamma \circ v^{-1} \circ f : \mathcal{S} \longrightarrow \mathcal{X}_*$. This is a composition of continuous functions, hence continuous. So $\alpha \in \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X}_*)$. But $U^{\mathcal{S}}_{\mathcal{X}_*}(\alpha) = u_{\mathcal{X}_*} \circ \alpha = u_{\mathcal{X}_*} \circ \gamma \circ v^{-1} \circ f = v \circ v^{-1} \circ f = f$. This works for any $f \in \mathfrak{C}(\mathcal{S}, [a,b])$. ◇ Claim 3

[33]See Theorem 5.6 of Mehta (1998) or p.631 of Bosi and Herden (2008).

For any $\mathcal{S} \in \mathsf{Top}^\circ$, defining equation (A17) says that $L(\mathcal{S})$ contains the linear span of the set $\{U_\mathcal{X}^\mathcal{S}(\alpha);\quad \alpha \in \overrightarrow{\mathsf{Top}}(\mathcal{S}, \mathcal{X})\}$, for all $\mathcal{X} \in \boldsymbol{\mathcal{X}}^\circ$. Thus, Claim 3 implies that $L(\mathcal{S}) = \mathfrak{C}_\mathrm{b}(\mathcal{S})$. In other words, L is the utility frame from Example 6.3(c). But in this case, if all objects in $\boldsymbol{\mathcal{S}}^\circ$ are normal Hausdorff spaces, then Proposition 6.8(c) yields a finitely additive normal Borel probability measure $\mu_\mathcal{S}$ on $\mathcal{S}$ for all $\mathcal{S} \in \boldsymbol{\mathcal{S}}^\circ$, satisfying equations (5) and (6). Thus, we have a global SEU representation of the kind described in Example 6.12(c). □

B Constants and constituents

This appendix gathers some supplementary results that justify the claims made about constant morphisms and constituents in Section 8. Let $\boldsymbol{\mathcal{C}}$ be a category. For any $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, let $\mathcal{K}(\mathcal{A}, \mathcal{B})$ be the set of constant morphisms from $\mathcal{A}$ to $\mathcal{B}$.

Lemma B.1 *If $\kappa \in \mathcal{K}(\mathcal{B}, \mathcal{C})$, then for any other objects $\mathcal{A}, \mathcal{D} \in \boldsymbol{\mathcal{C}}^\circ$ and morphisms $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\delta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$, the morphisms $\kappa \circ \alpha$ and $\delta \circ \kappa$ are also constant.*

Proof. Let $\mathcal{Z} \in \boldsymbol{\mathcal{C}}^\circ$, and let $\phi, \psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{Z}, \mathcal{A})$. Then $(\kappa \circ \alpha) \circ \phi = \kappa \circ (\alpha \circ \phi) \underset{(*)}{=} \kappa \circ (\alpha \circ \psi) = (\kappa \circ \alpha) \circ \psi$, where $(*)$ is because κ is constant. This holds for any $\mathcal{Z}$, ϕ and ψ. Thus, $\kappa \circ \alpha$ is constant.

Now let $\phi, \psi \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{Z}, \mathcal{B})$. Then $(\delta \circ \kappa) \circ \phi = \delta \circ (\kappa \circ \phi) \underset{(*)}{=} \delta \circ (\kappa \circ \psi) = (\delta \circ \kappa) \circ \psi$. where $(*)$ is because κ is constant. This holds for any $\mathcal{Z}$, ϕ and ψ. Thus, $\delta \circ \kappa$ is constant. □

Let $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, and let $\mathcal{K}(\bullet, \mathcal{C})$ be the set of all constant morphisms into $\mathcal{C}$ from *any* other object. Recall: for any noninitial $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$ and $\psi \in \mathcal{K}(\mathcal{A}, \mathcal{C})$ and $\phi \in \mathcal{K}(\mathcal{B}, \mathcal{C})$, we write $\psi \sim \phi$ there is some $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ such that $\psi = \phi \circ \alpha$.

Lemma B.2 *If $\boldsymbol{\mathcal{C}}$ is strongly connected, then $\sim$ is an equivalence relation on $\mathcal{K}(\bullet, \mathcal{C})$, for each $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$.*

Proof. We must show that $\sim$ is reflexive, symmetric, and transitive.

Reflexive. Set $\alpha = I_\mathcal{A}$ to conclude that $\phi \sim \phi$.

Symmetric. Let $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, let $\psi \in \mathcal{K}(\mathcal{A}, \mathcal{C})$ and $\phi \in \mathcal{K}(\mathcal{B}, \mathcal{C})$, and suppose $\psi \sim \phi$. Thus, there exists $\alpha \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ such that $\psi = \phi \circ \alpha$. Recall that $\overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{A}) \neq \emptyset$, because $\boldsymbol{\mathcal{C}}$ is strongly connected. Let $\beta \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{A})$ be arbitrary. Then $\psi \circ \beta = (\phi \circ \alpha) \circ \beta = \phi \circ (\alpha \circ \beta) \underset{(*)}{=} \phi \circ I_\mathcal{B} = \phi$, where $(*)$ is because ϕ is constant. Thus, $\psi \circ \beta = \phi$, so $\phi \sim \psi$, as desired.

Transitive. Let $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3 \in \boldsymbol{\mathcal{C}}^\circ$, let $\phi_n \in \mathcal{K}(\mathcal{A}_n, \mathcal{C})$ for $n = 1, 2, 3$, and suppose that $\phi_1 \sim \phi_2$ and $\phi_2 \sim \phi_3$. Thus, there exist $\alpha_1 \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}_1, \mathcal{A}_2)$ and $\alpha_2 \in \overrightarrow{\boldsymbol{\mathcal{C}}}(\mathcal{A}_2, \mathcal{A}_3)$ such that $\phi_1 = \phi_2 \circ \alpha_1$ and $\phi_2 = \phi_3 \circ \alpha_2$. Let $\alpha := \alpha_2 \circ \alpha_1$. Then $\phi_3 \circ \alpha = \phi_3 \circ \alpha_2 \circ \alpha_1 = \phi_2 \circ \alpha_1 = \phi_1$. Thus, $\phi_3 \sim \phi_1$, as desired. □

For all $\mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$, recall that $\underline{\mathcal{C}} := \mathcal{K}(\bullet, \mathcal{C})/_\sim$ is the set of *constituents* of $\mathcal{C}$. For any $\underline{c} \in \underline{\mathcal{C}}$ and $\phi \in \mathcal{K}(\mathcal{B}, \mathcal{C})$, we say that ϕ has the *value* $\underline{c}$ if ϕ belongs to the equivalence class $\underline{c}$.

Lemma B.3 *Suppose $\boldsymbol{\mathcal{C}}$ is strongly connected. For all $\mathcal{B}, \mathcal{C} \in \boldsymbol{\mathcal{C}}^\circ$ and $\underline{c} \in \underline{\mathcal{C}}$, there is a unique $\phi \in \mathcal{K}(\mathcal{B}, \mathcal{C})$ with value $\underline{c}$.*

Proof. (Existence) If $\underline{c} \in \underline{\mathcal{C}}$, then there is some $\mathcal{A} \in \boldsymbol{\mathcal{C}}^\circ$ and some $\kappa \in \mathcal{K}(\mathcal{A}, \mathcal{C})$ such that $\underline{c}$ is the equivalence class of κ. Let $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{A})$ be arbitrary (this exists because $\boldsymbol{\mathcal{C}}$ is strongly connected). Let $\phi := \kappa \circ \alpha$. Then ϕ is also constant (by Lemma B.1), and $\phi \sim \kappa$ (by definition) so ϕ has the value $\underline{c}$.

(Uniqueness) Let $\phi, \psi \in \mathcal{K}(\mathcal{B}, \mathcal{C})$ and suppose they both have the value $\underline{c}$. Then $\phi \sim \psi$. Thus, there is some $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{B})$ such that $\psi \circ \alpha = \phi$. But then $\phi = \psi \circ \alpha \underset{(*)}{=} \psi \circ I_{\mathcal{B}} = \psi$, where $(*)$ is because ψ is constant. We conclude that $\phi = \psi$. □

Lemma B.4 *Let $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D} \in \boldsymbol{\mathcal{C}}^\circ$ and let $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{C}, \mathcal{D})$. For all $\kappa_1 \in \mathcal{K}(\mathcal{A}, \mathcal{C})$ and $\kappa_2 \in \mathcal{K}(\mathcal{B}, \mathcal{C})$, if $\kappa_1 \sim \kappa_2$, then $(\phi \circ \kappa_1) \sim (\phi \circ \kappa_2)$.*

Proof. First note that $\phi \circ \kappa_1$ and $\phi \circ \kappa_2$ are themselves constants, by Lemma B.1. If $\kappa_1 \sim \kappa_2$, then there is some $\alpha \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ such that $\kappa_2 \circ \alpha = \kappa_1$. But then $(\phi \circ \kappa_2) \circ \alpha = \phi \circ (\kappa_2 \circ \alpha) = \phi \circ \kappa_1$. Thus, $(\phi \circ \kappa_1) \sim (\phi \circ \kappa_2)$. □

C Some abstract categories

This appendix briefly describes some non-concrete categories, to provide some context for Section 8. In some categories, morphisms are "generalized functions". For example:

- Objects in $\boldsymbol{\mathcal{C}}^\circ$ are sets. For any $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, the morphisms in $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ are *correspondences* —i.e. functions from $\mathcal{A}$ into the set $\wp^+(\mathcal{B})$ of nonempty subsets of $\mathcal{B}$. For any morphisms $\phi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\psi \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ and any $a \in \mathcal{A}$, $\psi \circ \phi(a) = \bigcup_{b \in \phi(a)} \psi(b)$.[34]

- Objects in $\boldsymbol{\mathcal{C}}^\circ$ are finite or countable sets. For any $\mathcal{A}, \mathcal{B} \in \boldsymbol{\mathcal{C}}^\circ$, the morphisms in $\vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ are *stochastic matrices* —i.e. arrays $\boldsymbol{\phi} = [\phi_b^a]_{b \in \mathcal{B}}^{a \in \mathcal{A}}$ of non-negative real numbers such that, for any $a \in \mathcal{A}$, $\sum_{b \in \mathcal{B}} \phi_b^a = 1$. For any morphisms $\boldsymbol{\phi} \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{A}, \mathcal{B})$ and $\boldsymbol{\psi} \in \vec{\boldsymbol{\mathcal{C}}}(\mathcal{B}, \mathcal{C})$ and any $a \in \mathcal{A}$ and $c \in \mathcal{C}$, $(\boldsymbol{\psi} \circ \boldsymbol{\phi})_c^a = \sum_{b \in \mathcal{B}} \phi_b^a \cdot \psi_c^b$.

[34] Ghirardato (2001) studied a version of the Savage model in which acts are represented by correspondences; he characterized a Choquet expected utility representations in this setting. In the present paper, the probabilistic extension of the category just described yields a version of the Anscombe-Aumann model in which acts are (mixtures of) correspondences. By applying Theorem 8.1, we obtain an SEU representation in this setting. For any $\mathcal{X} \in \boldsymbol{\mathcal{C}}^\circ$, the set of global elements $\underline{\mathcal{X}}$ in this category is in naturally isomorphic to $\wp^+(\mathcal{X})$. So via Proposition 6.5(b), a utility functional is determined by a function $u_{\mathcal{X}} : \wp^+(\mathcal{X}) \longrightarrow \mathbb{R}$.

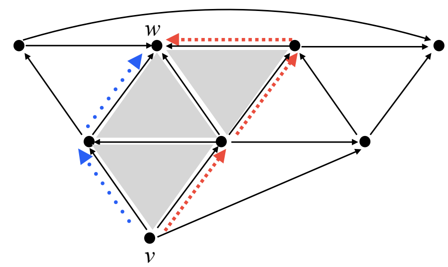

Figure 3: A 2-dimensional oriented simplicial complex. The two paths from v to w are homotopic.

In other categories, morphisms are *equivalence classes* of functions. For example:

- Objects in $\mathcal{C}^\circ$ are *based topological spaces* (i.e. topological spaces with one point selected as a "base point"). For any $\mathcal{A}, \mathcal{B} \in \mathcal{C}^\circ$, morphisms in $\overrightarrow{\mathcal{C}}(\mathcal{A}, \mathcal{B})$ are homotopy equivalence classes of basepoint-preserving continuous maps form $\mathcal{A}$ to $\mathcal{B}$. (A function $\phi : \mathcal{A} \longrightarrow \mathcal{B}$ is *basepoint-preserving* if it sends the basepoint of $\mathcal{A}$ to the basepoint of $\mathcal{B}$. Two basepoint-preserving continuous maps $\phi_0, \phi_1 : \mathcal{A} \longrightarrow \mathcal{B}$ are *homotopy equivalent* if there is a continuous function $\Phi : [0,1] \times \mathcal{A} \longrightarrow \mathcal{B}$ such that $\Phi(0, \bullet) = \phi_0$, $\Phi(1, \bullet) = \phi_1$, and $\Phi(s, \bullet) : \mathcal{A} \longrightarrow \mathcal{B}$ is basepoint-preserving for all $s \in [0,1]$.)

In other categories, objects are not sets, and morphisms have no resemblance to functions.

- Fix a topological space $\mathcal{T}$. The objects of $\mathcal{C}^\circ$ are the points in $\mathcal{T}$. For any objects $a, b \in \mathcal{C}^\circ$, the morphisms in $\overrightarrow{\mathcal{C}}(a, b)$ are homotopy equivalence classes of continuous paths from a to b. (A *path* from a to b is a continuous function $\phi : [0,1] \longrightarrow \mathcal{T}$ such that $\phi(0) = a$ and $\phi(1) = b$. Two such paths ϕ_0 and ϕ_1 are *homotopy-equivalent* if there is a continuous function $\Phi : [0,1] \times [0,1] \longrightarrow \mathcal{T}$ such that $\Phi(0, \bullet) = \phi_0$, $\Phi(1, \bullet) = \phi_1$, and $\Phi(s, \bullet) : [0,1] \longrightarrow \mathcal{T}$ is a path from a to b for all $s \in [0,1]$.)

  Composition of morphisms is obtained by *concatenation* of paths. If $\phi : [0,1] \longrightarrow \mathcal{T}$ is a path from a to b, and $\psi : [0,1] \longrightarrow \mathcal{T}$ is a path from b to c, then we define $(\psi * \phi) : [0,1] \longrightarrow \mathcal{T}$ by setting $(\psi * \phi)(t) := \phi(2\,t)$ for $t \in [0, \frac{1}{2}]$ and $(\psi * \phi)(t) := \psi(2\,t - 1)$ for $t \in [0, \frac{1}{2}]$; this is a path from a to c. Finally, given two morphisms (i.e. homotopy equivalence classes of paths) $[\phi]$ and $[\psi]$, we define $[\phi] \circ [\psi] := [\phi * \psi]$.

- Fix a logical language $\mathcal{L}$. The objects in $\mathcal{C}^\circ$ are sentences in $\mathcal{L}$. For any objects $a, b \in \mathcal{C}^\circ$, a morphism in $\overrightarrow{\mathcal{C}}(a, b)$ is a *proof*—i.e. a sequence of sentences in $\mathcal{L}$ which logically derives b from a. Composition of morphisms is just concatenation of proofs: if ϕ is a proof that derives b from a, and ψ is a proof that derives c from b, then the concatenation $\phi\psi$ is a proof that derives c from a.

We conclude with a general class of examples. A *2-dimensional simplicial complex* consists of a collection $\mathcal{V}$ of zero-dimensional *vertices*, a collection $\mathcal{E}$ of one-dimensional *edges* which connect pairs of vertices, and a collection $\mathcal{F}$ of two-dimensional triangular *faces*, each of which spans three edges. (There can be many edges between any pair of vertices, and the same edge can be part of many faces.) We will say that a simplicial complex $(\mathcal{V}, \mathcal{E}, \mathcal{F})$ is ***oriented*** if each of the edges in $\mathcal{E}$ is given an orientation, so that it goes from one vertex to another; see Figure 3 for an example.

Given to vertices $v, w \in \mathcal{V}$, a ***directed path*** is a sequence $\zeta = (v_0, e_1, v_1, e_2, v_2, \ldots, v_{N-1}, e_N, v_N)$ with $v_0, v_1, \ldots, v_N \in \mathcal{V}$ and $e_1, e_2, \ldots, e_N \in \mathcal{E}$, where $v_0 = v$, $v_N = w$, and for all $n \in [1 \ldots N]$, e_n is a directed edge from v_{n-1} to v_n.

Now let $u, v, w \in \mathcal{V}$, and suppose that $\zeta = (v_0, e_1, \ldots, e_N, v_N)$ is a directed path from u to v and $\zeta' = (v'_0, e'_1, \ldots, e'_M, v'_M)$ is a directed path from v to w. Then in particular, $v_N = v = v'_0$. We define the ***concatenation*** of ζ and ζ' to be the directed path $\zeta' * \zeta := (v_0, e_1, \ldots, e_N, v_N = v'_0, e'_1, \ldots, e'_M, v'_M)$; this is a path from u to w.

Now fix $v, w \in \mathcal{V}$, and suppose that $\xi = (v_0, e_1, \ldots, e_N, v_N)$ and $\xi' = (v'_0, e'_1, \ldots, e'_M, v'_M)$ are two different directed paths from v to w. We will say that ξ' is a ***reduction*** of ξ if $M = N - 1$ and there is some $m \in [1 \ldots N]$ such that $v'_n = v_n$ and $e'_n = e_n$ for all $n \leq m$, $v'_n = v_{n+1}$ for all $n \geq m$ and $e'_n = e_{n+1}$ for all $n \geq m + 1$, and the three edges e_n, e_{n+1} and e'_n span a face in $\mathcal{F}$. Finally, given two directed paths ζ and ζ' from v to w, we will say that ζ' is ***homotopic*** to ζ if there is a sequence of paths $\xi_0, \xi_1, \ldots, \xi_L$ from v to w such that $\xi_0 = \zeta$, $\xi_L = \zeta'$, and for all $\ell \in [1 \ldots L]$, either ξ_ℓ is a reduction of $\xi_{\ell-1}$, or $\xi_{\ell-1}$ is a reduction of ξ_ℓ. For instance, in Figure 3, the two paths from v to w are homotopic.

It is easily verified that homotopy is an equivalence relation on the set of paths. Furthermore, if for any $u, v, w \in \mathcal{V}$ and paths ζ, ζ' from u to v and paths ξ, ξ' from v to w, if ζ is homotopic to ζ' and ξ is homotopic to ξ', then $\xi * \zeta$ is homotopic to $\xi' * \zeta'$.

Now, given an oriented 2-dimensional simplicial complex $(\mathcal{V}, \mathcal{E}, \mathcal{F})$, consider the category $\boldsymbol{\mathcal{C}}$ defined as follows. The objects in $\boldsymbol{\mathcal{C}}^\circ$ are the vertices in $\mathcal{V}$, and for any $v, w \in \mathcal{V}$, the morphisms in $\overrightarrow{\boldsymbol{\mathcal{C}}}(v, w)$ are the homotopy equivalence classes of directed paths from v to w. Composition of morphisms is obtained by concatenation of the underlying paths; as we have observed in the previous paragraph, this is well-defined because homotopy equivalence is compatible with the concatenation operation.

This example is fully general, in the sense that *any* category can be represented as a category of this kind, for a sufficiently large simplicial complex $(\mathcal{V}, \mathcal{E}, \mathcal{F})$.[35]

References

Abramsky, S., 2015. Arrow's theorem by arrow theory. In: Hirvonen, Å., Kontinen, J., Kossak, R., Villaveces, A. (Eds.), Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics. De Gruyter, Ch. 2, pp. 15–30.

Abramsky, S., Winschel, V., 2017. Coalgebraic analysis of subgame-perfect equilibria in infinite games without discounting. Mathematical Structures in Computer Science 27 (5), 751–761.

Adachi, T., 2014. Toward categorical risk measure theory. Theory and Applications of Categories 29 (14), 389–405.

Adámek, J., Herrlich, H., Strecker, G. E., 2009. Abstract and concrete categories: the joy of cats. Dover, Mineola, NY.

Ahn, D. S., 2008. Ambiguity without a state space. Review of Economic Studies 75.1 (262), 3–28.

[35]However, in general, $\mathcal{V}$, $\mathcal{E}$ and $\mathcal{F}$ are not sets; they are proper classes.

Ahn, D. S., Ergin, H., 2010. Framing contingencies. Econometrica 78 (2), 655–695.

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

Amarante, M., 2015. Analogy in decision making. Mathematics of Operations Research 40 (4), 1027–1041.

Anscombe, F. J., Aumann, R. J., 1963. A definition of subjective probability. Annals of Mathematical Statistics 34, 199–205.

Awodey, S., 2010. Category theory, 2nd Edition. Oxford University Press, Oxford.

Blume, L., Easley, D., Halpern, J. Y., 2021. Constructive decision theory. Journal of Economic Theory 196, 105306.

Bosi, G., Herden, G., 2008. Continuous utility representation theorems in arbitrary concrete categories. Applied Categorical Structures 16 (5), 629–651.

Cockett, R., Seely, R., 2007. Polarized category theory, modules, and game semantics. Theory and Applications of Categories 18 (2), 4–101.

Cohen, M., Jaffray, J.-Y., 1980. Rational behavior under complete ignorance. Econometrica, 1281–1299.

De Oliveira, H., 2016. Blackwell's informativeness theorem using category theory, (preprint).

Dekel, E., Lipman, B. L., Rustichini, A., 2001. Representing preferences with a unique subjective state space. Econometrica 69 (4), 891–934.

Dietrich, F., 2018. Savage's theorem under changing awareness. Journal of Economic Theory 176, 1–54.

Dominiak, A., Tserenjigmid, G., 2018. Ambiguity under growing awareness, (preprint).

Eberhart, C., Hirschowitz, T., 2018. What's in a game?: A theory of game models. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science. ACM, pp. 374–383.

Eklund, P., Fedrizzi, M., Helgesson, R., 2013. Monadic social choice. Multicriteria and Multiagent Decision Making with Applications to Economics and Social Sciences, 143–150.

Eklund, P., Fedrizzi, M., Nurmi, H., 2010. A categorical approach to the extension of social choice functions. In: Information Processing and Management of Uncertainty in Knowledge-Based Systems. Applications: 13th International Conference, IPMU 2010, Dortmund, Germany, June 28–July 2, 2010. Proceedings, Part II 13. Springer, pp. 261–270.

Fishburn, P. C., 1970. Utility theory for decision making. John Wiley & Sons, New York.

Folland, G. B., 1984. Real Analysis. John Wiley & Sons, New York.

Fremlin, D. H., 2012. Measure theory. Vol. 3: Measure Algebras, 2nd Edition. Torres Fremlin, Colchester, U.K.

Fukuda, S., 2021. Topology-free constructions of a universal type space as coherent belief hierarchies, (working paper).

Galeazzi, P., Marti, J., 2023. Choice structures in games. Games and Economic Behavior 140, 431–455.

Ghirardato, P., 2001. Coping with ignorance: unforeseen contingencies and non-additive uncertainty. Economic Theory 17, 247–276.

Gilboa, I., Schmeidler, D., 2004. Subjective distributions. Theory and Decision 56 (4), 345–357.

Guarino, P., 2022. Topology-free type structures with conditioning events. arXiv preprint arXiv:2212.07246.

Hammond, P. J., 2022. Quantum probability with measurable spaces as random contexts I: Two preliminary examples.

Hayashi, T., 2012. Expanding state space and extension of beliefs. Theory and Decision 73 (4), 591–604.

Heifetz, A., Meier, M., Schipper, B. C., 2006. Interactive unawareness. J. Econom. Theory 130 (1), 78–94.

Heinsalu, S., 2014. Universal type structures with unawareness. Games and Economic Behavior 83, 255–266.

Hirschowitz, M., Hirschowitz, A., Hirschowitz, T., 2007. A theory for game theories. In: International Conference on Foundations of Software Technology and Theoretical Computer Science. Springer, pp. 192–203.

Honsell, F., Lenisa, M., 2011. Conway games, algebraically and coalgebraically. Logical Methods in Computer Science 7 (3).

Honsell, F., Lenisa, M., Pellarini, D., 2014. Categories of coalgebraic games with selective sum. Fundamenta Informaticae 134 (3-4), 395–414.

Honsell, F., Lenisa, M., Redamalla, R., 2012a. Categories of coalgebraic games. In: International Symposium on Mathematical Foundations of Computer Science. Springer, pp. 503–515.

Honsell, F., Lenisa, M., Redamalla, R., 2012b. Equivalences and congruences on infinite Conway games. RAIRO-Theoretical Informatics and Applications 46 (2), 231–259.

Jaynes, E. T., 1968. Prior probabilities. IEEE Transactions on systems science and cybernetics 4 (3), 227–241.

Jaynes, E. T., 2003. Probability theory: The logic of science. Cambridge university press.

Joyal, A., 1977. Remarques sur la théorie des jeux à deux personnes. Gazette des sciences mathématiques du Québec 1 (4), 46–52.

Karni, E., 2006. Subjective expected utility theory without states of the world. J. Math. Econom. 42 (3), 325–342.

Karni, E., 2007. A new approach to modeling decision-making under uncertainty. Econom. Theory 33 (2), 225–242.

Karni, E., 2011. A theory of Bayesian decision making with action-dependent subjective probabilities. Econom. Theory 48 (1), 125–146.

Karni, E., 2013. Bayesian decision theory with action-dependent probabilities and risk attitudes. Econom. Theory 53 (2), 335–356.

Karni, E., 2017. States of nature and the nature of states. Economics & Philosophy 33 (1), 73–90.

Karni, E., Schmeidler, D., 1991. Utility theory with uncertainty. Handbook of mathematical economics IV, 1763–1831.

Karni, E., Vierø, M.-L., 2013. "Reverse Bayesianism": a choice-based theory of growing awareness. The American Economic Review 103 (7), 2790–2810.

Karni, E., Vierø, M.-L., 2015. Probabilistic sophistication and reverse Bayesianism. Journal of Risk and Uncertainty 50 (3), 189–208.

Karni, E., Vierø, M.-L., 2017. Awareness of unawareness: a theory of decision making in the face of ignorance. J. Econom. Theory 168, 301–328.

Katznelson, Y., 2004. An introduction to harmonic analysis. Cambridge University Press.

Keiding, H., 1981. The categorical approach to social choice theory. Mathematical Social Sciences 1 (2), 177–191.

Keynes, J. M., 1921. A treatise on probability. MacMillan, London.

Kijima, K., Takahara, Y., 1987. Category-theoretic basis for social decision-making theory. International journal of systems science 18 (2), 251–267.

Kreps, D. M., 1979. A representation theorem for "preference for flexibility". Econometrica 47 (3), 565–577.

Laplace, P. S., 1820. Théorie analytique des probabilités. Courcier.

Leinster, T., 2014. Basic category theory. Cambridge University Press, Cambridge.

Lenisa, M., 2015. Multigames and strategies, coalgebraically. Theoretical Computer Science 604, 46–62.

Lescanne, P., Perrinel, M., 2012. "Backward" coinduction, Nash equilibrium and the rationality of escalation. Acta informatica 49 (3), 117–137.

Mac Lane, S., 1998. Categories for the working mathematician, 2nd Edition. Springer-Verlag, New York.

Machina, M. J., 2003. States of the world and the state of decision theory. In: Meyer, D. J. (Ed.), The Economics of Risk. W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, Ch. 2, pp. 17–49.

Machover, M., Terrington, S. D., 2014. Mathematical structures of simple voting games. Mathematical Social Sciences 71, 61–68.

Mehta, G. B., 1998. Preferences and utility. In: Hammond, P. J., Seidl, C., Barberà, S. (Eds.), Handbook of utility theory. Vol. 1. Kluwer Academic Publishers, Ch. 1, pp. 1–48.

Meyer-Nieberg, P., 1991. Banach lattices. Universitext. Springer-Verlag, Berlin.

Moss, L. S., Viglizzo, I. D., 2004. Harsanyi type spaces and final coalgebras constructed from satisfied theories. Electronic notes in theoretical computer science 106, 279–295.

Moss, L. S., Viglizzo, I. D., 2006. Final coalgebras for functors on measurable spaces. Information and Computation 204 (4), 610–636.

Pintér, M., 2010. The non-existence of a universal topological type space. Journal of Mathematical Economics 46 (2), 223–229.

Pivato, M., 2020. Subjective expected utility with a spectral state space. Economic Theory 69 (2), 249–313.

Pivato, M., Vergopoulos, V., 2020. Subjective expected utility with imperfect perception. Journal of Mathematical Economics 88, 104–122.

Riehl, E., 2017. Category theory in context. Courier Dover Publications.

Rozen, V., Zhitomirski, G., 2006. A category theory approach to derived preference relations in some decision making problems. Mathematical Social Sciences 51 (3), 257–273.

Savage, L. J., 1954. The foundations of statistics. John Wiley & Sons, Inc., New York; Chapman & Hill, Ltd., London.

Schipper, B. C., 2013. Awareness-dependent subjective expected utility. Internat. J. Game Theory 42 (3), 725–753.

Schmeidler, D., Wakker, P., 1990. Expected utility and mathematical expectation. In: Utility and Probability. Springer, pp. 70–78.

Simmons, H., 2011. An introduction to category theory. Cambridge University Press, Cambridge.

Skiadas, C., 1997a. Conditioning and aggregation of preferences. Econometrica 65 (2), 347–367.

Skiadas, C., 1997b. Subjective probability under additive aggregation of conditional preferences. Journal of Economic Theory 76 (2), 242–271.

Spivak, D. I., 2014. Category theory for the sciences. MIT Press, Cambridge, MA.

Streufert, P. A., 2018. The category of node-and-choice preforms for extensive-form games. Studia Logica 106 (5), 1001–1064.

Streufert, P. A., 2020. The category of node-and-choice forms, with subcategories for choice-sequence forms and choice-set forms. In: Nonclassical Logics and Their Applications: Post-proceedings of the 8th International Workshop on Logic and Cognition. Springer, pp. 15–66.

Tohmé, F., Caterina, G., Gangle, R., 2017. Local and global optima in decision-making: a sheaf-theoretical analysis of the difference between classical and behavioral approaches. International Journal of General Systems 46 (8), 879–897.

Tohmé, F., Viglizzo, I., 2023. A categorical representation of games, (working paper).

Vassilakis, S., 1992. Some economic applications of Scott domains. Mathematical Social Sciences 24 (2-3), 173–208.

Vorob'ev, N. N., 1962. Consistent families of measures and their extensions. Theory of Probability & Its Applications 7 (2), 147–163.

Zhou, L., 1999. Subjective probability theory with continuous acts. J. Math. Econom. 32 (1), 121–130.