



Munich Personal RePEc Archive

Measuring utility without mixing apples and oranges and eliciting beliefs about stock prices

O'Callaghan, Patrick

University of Queensland

9 February 2016

Online at <https://mpra.ub.uni-muenchen.de/69363/>

MPRA Paper No. 69363, posted 15 Feb 2016 16:47 UTC

Measuring utility without mixing apples and oranges and eliciting beliefs about stock prices

Patrick O’Callaghan[†]

University of Queensland

February 15, 2016

Abstract

In day-to-day life we encounter decisions amongst prospects that do not have a convex structure. To address this concern, Herstein and Milnor introduce mixture sets and provide necessary and sufficient conditions for a cardinal and linear utility representation. We derive the same utility representation for partial mixture sets: where the mixture operation is only partially defined. The resulting model has an interesting application to finance. In particular, we use paths instead of events to elicit utility and beliefs about stock prices. This feature is promising for settings where the dimension of the state space is large.

1 Introduction

Many would agree that good apples are crispy and that good oranges are tangy. Experts’ tastes are highly refined along such scales, with procurement and pricing hinging entirely on such comparisons. The ability to measure utility on a single scale that transcends specific characteristics is essential to many fields in economics.

Using their famous example of a glass of tea and a cup of coffee, von Neumann and Morgenstern famously described how this might be achieved using lotteries as a measurement device. Up to some reasonable approximation, experimenters today are able to measure utility up to multiplication by a positive scalar and addition by an arbitrary constant (a cardinal scale). Yet lotteries may lead to distortions that are intrinsic to the uncertainty they introduce. Herstein and Milnor showed that line segments (convex combinations of lotteries) may be replaced with mixture paths, and utility may still be measured.

[†]I thank Simon French, Simon Grant, Peter Hammond, Andrea Isoni, Saul Jacka, Jeff Kline, John Quiggin, Aron Toth and Horst Zank for their helpful suggestions and detailed feedback. I am especially grateful to Simon Grant, Jeff Kline and John Quiggin for their encouragement.

Overview

In the present paper, we allow for the possibility that some of the mixture paths are missing: for instance when there is no obvious mixture path between a crispy apple and a tangy orange. The main theorem identifies necessary and sufficient conditions for a cardinal and linear utility representation when mixtures are partially defined in this way. Our proof is constructive and introduces concepts that are, to our knowledge, new. En route to this result, we provide detailed insight into the nature of mixture paths themselves.

Finally, we show how the model may be applied to elicit beliefs about stock prices in an experimental setting, where prices evolve in continuous time. In the specific setting we describe, the instruments we use to elicit beliefs are (nonconvex) mixture paths known as Brownian bridges. The partial mixture set structure allows us to elicit beliefs using pure subjective moments: there is no need to consider mixtures between moments of different order.

Theoretical framework The idea of measuring utility using lotteries can be traced back to the St. Petersburg paradox and the “moral expectation” of Bernoulli [4]. By considering preferences over lotteries, von Neumann and Morgenstern show that utility can, in principle, be measured in much the same way as temperature (on a cardinal scale). In view of the fact that most prospects in life are not lotteries, HM show that the convex structure of lotteries is unnecessary for the purposes of measuring utility.

The first goal of the present paper is to gain a deeper understanding of the abstraction of convexity that HM introduced and its relationship with their axioms on preferences. HM define preferences over elements in a mixture set. In a mixture set, every pair of prospects is connected by a special kind of path: the subpath between every pair of points along this path is “synchronised” with the path (see fig. 1). HM then require preferences to respect axioms that are necessary and sufficient for a cardinal and linear utility representation.

Our main innovation is to drop the requirement that every pair of prospects is connected by a path of mixtures. In other words, we weaken the “for all” requirement of HM and formally introduce partial mixture sets. This modification has a number of far-reaching implications. For instance preferences satisfying the HM axioms (constrained to hold on the mixture paths that exist) may have no utility representation. Via examples we show that for a cardinal and linear utility representation, stronger axioms are needed to accommodate the more general structure of partial mixture sets. To our knowledge, we are the first to provide conditions for such a utility representation that are both necessary and sufficient.

Partial mixture sets are motivated by our toy example regarding the absence of a path connecting the small world of crispy apples and the small world tangy oranges. But one may argue that the solution is to find a common attribute, perhaps juiciness, that transcends the two goods. But how should we do this when our tool for measuring utility is preferences? The alternatives available to the economist who wishes to measure utility are constrained by physical,

biological or psychological boundaries. Many objects, like houses have inherent discrete properties; not all genetic permutations of fruit are possible; and why assume that the decision maker imagines the same convex set of lotteries we have in mind?

Application A compelling reason for adopting partial mixture sets is that they allow the modeller to tailor the domain of preferences to suit the problem at hand. This is particularly relevant in “large world” decision settings, such as finance, where models are often complicated enough without introducing the additional layer of uncertainty that lotteries would entail.

Consider a trader that has well-formed, but perhaps inaccurate beliefs about the daily evolution of the stock prices of Apple Inc. In seeking to elicit her beliefs from preferences, it would be natural to adopt the benchmark model of [2]. We argue that the difficulty is twofold. First, we must construct a mixture set. This entails setting up a host of derivative securities that yield payoffs that depend on the path that price takes. (These derivatives correspond to “acts” in the model of Anscombe and Aumann.) Derivatives are typically hard to evaluate. The second difficulty persists even if we restrict attention to acts of the simplest form: binary options that pay a dollar if event E occurs and zero otherwise. The issue is that events are highly non-trivial in this setting. The simplest events are commonly known as cylinder sets (see figure 2). Are decision makers able to gauge the likelihood of such objects?

Instead, we propose constructing a partial mixture set using stochastic paths and eliciting Val’s subjective moments (as subjective expectations of powers of random variables) and appealing to our extension of the HM model. Although one would have to elicit a countable infinity of moments to really pin down (the set of) beliefs, the cost of using [2] or [25] still seems higher. This is because there are well-established methods that prescribe what to do in practice when only a finite number of moments are feasible. In particular, by minimising the entropy relative to uniform we obtain a worst case (least informed) estimate of the trader’s beliefs. Or, if the true distribution is available, the best case estimate of beliefs is found by minimising the relative to the truth. Of course, only an experiment would settle the question of whether our proposal improves on the standard approach and that is beyond the scope of the present paper.

Related literature

There are a number of strands that take preferences as primitive and that are in the same spirit as ours. The first strand (see [8] and references therein) remains close to HM by assuming a product of mixture sets and corresponding axioms. Since a product of sets can be rewritten as a discrete or disjoint union, the conditions for a partial mixture set are satisfied, and so our representation theorem generalises these results. In the second strand ([18], [11] and [17]), lottery spaces that form a partial mixture space are adopted. In contrast with the present paper and the first strand, these models provide axioms that are

only sufficient for a representation. We provide a more detailed comparison in the discussion of the representation theorem.

The approach of Krantz et al. [19] also allows for nonconvex spaces. Instead they consider spaces that have objects that are “equally spaced” according to some external measure. (Consider, for instance, the natural numbers.) This property is not a requirement for partial mixture sets.

In a recent paper Richter and Rubinstein [23] discuss abstract convexity in an equilibrium context. The kind of abstraction considered there would allow for discrete sets with no possibility of a cardinal utility representation. On the other hand, if cardinal utility were to feature in their model, then the partial mixture set structure we propose would be a natural place to start.

2 Model

Prospects and preferences Let X denote a nonempty set of prospects. Provided x and y belong to X , the statement “ y is weakly preferred to x ” is summarised by the expression $x \lesssim y$. The collection \lesssim of such statements is the primitive object that we henceforth refer to as *preferences*. Formally, \lesssim is a binary relation on X , so that \lesssim is a subset of $X \times X$. As a consequence, the statement $x \lesssim y$ already implies $x, y \in X$. The following partial converse is the completeness axiom: if $x, y \in X$, then $x \lesssim y$ or $y \lesssim x$. It ensures that every pair of prospects is comparable. Transitivity requires that $x \lesssim y$ and $y \lesssim z$ together imply $x \lesssim z$. These two axioms of HM are standard and thoroughly discussed elsewhere.

Axiom Θ . \lesssim is *transitive* and *complete* on X .

As usual, $<$ denotes the asymmetric, strict subrelation of \lesssim and the symmetric, indifference subrelation is denoted by \sim . In this way, preferences are partitioned, so that $x \lesssim y$ if and only if either $x < y$ or $x \sim y$. An important point to note for what follows is that, when a decision maker’s preferences satisfy Θ , the basic open sets $\{x' : x < x'\}$ and $\{x' : x' < y\}$ such that $x, y \in X$ generate a topology $\tau_{<}$ on X . (Every open set is a union of finite intersections of these basic sets.) We do not assume the decision maker is conscious of this topology. Rather, we view this a part of the modeller’s toolkit and refer to it as the topology generated by preferences or the *preference order topology*. As with HM, no external topological conditions are required of X .

Paths and mixtures Throughout, I will denote the closed unit interval $[0, 1]$ in \mathbb{R} . When X is a convex set, the convex combination of any given pair $x, y \in X$, is a map $\lambda \mapsto (1 - \lambda)x + \lambda y \in X$ for each $\lambda \in I$. This map constitutes a certain kind of *path from x to y* . Seeking minimal conditions for measuring utility, HM introduced a special form of path that substantially generalises the notion of convexity. The definition that now follows coincides with that of HM when every pair of prospects $x, y \in X$ defines a *path of mixtures* ϕ_{xy} . Formally, we weaken the definition of HM by allowing a *partial function* to characterise the

paths on X . Recall that a partial function is one that is not defined throughout its domain.

Definition 1. (X, Φ) is a *partial mixture set* whenever $\Phi : X \times X \times I \rightarrow X$ is a partial function such that if $\phi_{xy} := \Phi(x, y, \cdot)$ is defined, then for every $\lambda, \mu \in I$,

$$\mathcal{P}_1 \quad \phi_{xy}(0) = x,$$

$$\mathcal{P}_2 \quad \phi_{yx} \text{ is defined and } \phi_{yx}(\lambda) = \phi_{xy}(1 - \lambda), \text{ and}$$

$$\mathcal{P}_3 \quad z = \phi_{xy}(\mu) \text{ implies } \phi_{xz} \text{ is defined and } \phi_{xz}(\lambda) = \phi_{xy}(\lambda\mu).$$

The fact that Φ is a partial function ensures that if ϕ_{xy} is defined, then it is uniquely identified by the points x and y . Together, \mathcal{P}_1 and \mathcal{P}_2 confirm that x and y are the *endpoints* of ϕ_{xy} . \mathcal{P}_3 really is the cornerstone of the definition, for as we highlight in the discussion of example 4 below, it also rules out certain ordered sets that are too large to be represented by a real-valued utility function. (Recall this is a function $U : X \rightarrow \mathbb{R}$ such that $x \lesssim y$ if and only if $U(x) \leq U(y)$.)

When there is no ambiguity about the identity of Φ , is common practice to simply refer to X as the partial mixture set with the understanding that this is shorthand for (X, Φ) . With a minor abuse of notation, we let Φ also denote *the collection of mixture paths ϕ_{xy} in X* . The condition for X to be a mixture set is then

$$\{(x, y) : \phi_{xy} \in \Phi\} = X \times X.$$

In the sequel, we will often refer to a path in Φ without reference to its endpoints. For instance, we may consider a given path ϕ or a sequence ϕ_1, ϕ_2, \dots of paths in Φ .

Paths as building blocks for the model When X is a mixture set, the fact that it has a full set of paths makes it a self-contained building block for the model. When X is only a partial mixture set, we must turn to paths in Φ for building blocks. For this reason, the following proposition is a useful place to start. With the aid of fig. 1, it summarises the implications of conditions \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}_3 .

Proposition 1. [†] For each $\phi_{xy} \in \Phi$, the image $\phi_{xy}(I)$ is a mixture set. Moreover, for every $x', y' \in \phi_{xy}(I)$, there exists $\mu, \nu \in I$ such that

$$\phi_{x'y'}(\lambda) = \phi_{xy}((1 - \lambda)\mu + \lambda\nu) \quad \text{for every } \lambda \in I. \quad (1)$$

By proposition 1, we may apply the results of HM to any given path in Φ . Figure 1 and (1) present a very precise relationship between a path in Φ and its subpaths. When two paths are related via (1), we say they are *synchronised*. In what follows we will extend this definition to allow for indifferences to replace equality. When X is a partial mixture set, a primary motivation for the axioms on preferences, is to ensure the building blocks that are the paths in Φ can be

[†]See page 24 for proof.




Figure 1: The graph of ϕ_{xy} is the set of pairs (λ, z) such that $\lambda \in I$ and $z = \phi_{xy}(\lambda)$. Note that this inverse-S shaped curve in $I \times X$ differs from the well-known counterpart in Prospect Theory [16] which belongs to $I \times I$. We discuss this point further in connection with Rank-dependent utility theory of Quiggin [22] in section 4 below. When \mathcal{P}_2 holds, the graph of ϕ_{yx} coincides with the set of pairs mapped out by $\phi_{xy}(1 - \lambda)$. Note that, in this example, the graph of ϕ_{yx} is also upward sloping and inverse S-shaped, provided we place y at the origin. This is not true of $\phi_{xx'}$. When \mathcal{P}_3 holds, the graph of $\phi_{xx'}$ is the set (λ, z) such that $\lambda \in I$ and $z = \phi_{xy}(\lambda\mu)$: the initial, concave segment of ϕ_{xy} . Combining all three conditions, proposition 1 shows that the graph of $\phi_{x'y'}$ is characterised by points in the middle segment, that is (λ, z) such that $z = \phi_{xy}((1 - \lambda)\mu + \lambda\nu)$.

put together in a synchronised manner. Heuristically speaking, a path is synchronised with another if it coincides with another path once it is truncated and the “rate of travel” along it is adjusted accordingly. Of course, with preferences providing the only topological structure on X , notions such as “rate of travel” and “distance” are not well-defined. By defining mixture sets in this abstract way, the point that HM seek to make is that, provided the paths in Φ preserve the structure of the interval I , definition 1 is all we need to measure utility. Indeed, the second motivation for the axioms on preferences that follow is to ensure that the image of a path in Φ has similar properties to those of I .

Two ways of generating partial mixture sets As we will see in example 1, without any further restrictions on preferences, even an arbitrary discrete set can be a mixture set. In the present subsection, we present examples and concepts that complement the axioms that follow.

One way to obtain a partial mixture set is to simply remove points from a mixture set. The paths that remain in tact after this deletion generate a partial mixture set. An intuitive way to summarise the idea is to think of the earth’s surface as a mixture set and the surface of the continents as a partial mixture set. The following proposition and argument formalises this intuition.

Proposition 2. [†] Let X be the surface of a sphere and let Φ be the set of all geodesic paths on X . Then (X, Φ) is a mixture set.

Subject to some additional, but reasonable assumptions, we may take X in proposition 2 to be the earth’s surface. Now let X' be the set of points on land, so that X' describes the surface of the continents. Moreover, let $\Phi' = \{\phi \in \Phi : \phi(I) \subseteq X'\}$. Since every path in Φ' automatically satisfies \mathcal{P}_1 , \mathcal{P}_2 and \mathcal{P}_3 , the continents can indeed be written as a partial mixture set.

Another way to generate partial mixture sets is to combine a collection $\{X_a : a \in A\}$ of mixture sets. This is possible provided we note that pairs of endpoints uniquely identify paths in a partial mixture set. Thus, one condition for combining mixture sets is that no pair X_a and X_b such that $a \neq b$ shares more than a single point. Alternatively, we ensure that the paths in the intersection $X_a \cap X_b$ coincide.

In fact, as we discuss further in section 4, the case where the collection is pairwise disjoint already accounts for many models in the literature. When the collection is pairwise disjoint, $\bigsqcup \{X_a : a \in A\}$ denotes the *disjoint union*. For the reasons we have just outlined, it is clear that the following statement holds.

Proposition 3. Every disjoint union of mixture sets is a partial mixture set.

The converse of proposition 3 is not possible in general. Indeed it is only possible if the set Φ of paths induces a partition of X . (Recall that a partition is characterised by an equivalence relation, that is a reflexive, symmetric and transitive binary relation.) We now describe how to test whether Φ induces a partition on X .

[†]See proof on page 25

Consider the set of pairs (x, y) such that $\phi_{xy} \in \Phi$. The only property that follows directly from the definition of a partial mixture set is symmetry: if $\phi_{xy} \in \Phi$, then $\phi_{yx} \in \Phi$. But, if every point in X belongs to some path in Φ , then \mathcal{P}_1 implies $\phi_{xx} \in \Phi$. When this is true, the relation induced by Φ is reflexive. As we have seen in proposition 1, the definition of a partial mixture set is such that when $\phi_{xy} \in \Phi$, then for every $z \in \phi_{xy}(I)$, both ϕ_{xz} and ϕ_{zy} belong to Φ . The final requirement for a partial mixture set to be written as a disjoint union of mixture sets, transitivity, is the converse property: if $\phi_{xz}, \phi_{zy} \in \Phi$, then for some unique $0 < \mu < 1$, ϕ_{xy} belongs to Φ and it can be written as the *concatenation*

$$\phi_{xy}(\lambda) = \begin{cases} \phi_{xz}(\lambda/\mu) & 0 \leq \lambda \leq \mu \\ \phi_{zy}((\lambda - \mu)/(1 - \mu)) & \mu \leq \lambda \leq 1 \end{cases} \quad (2)$$

The justification for (2) lies in a simple inversion of the transformation $\lambda \mapsto (1 - \lambda)\mu + \lambda\nu$ and an appeal to equation (1) of proposition 1.[†]

Clearly, the latter assumption makes sense in some applications. It would, for example, make sense in when X is transport network, where the intuition is that upon arriving at z from x , we could continue on to y . But in many settings, the path from x to y may not be via z . This is precisely the case in the application of section 3. Indeed, the earth-continent example above does not satisfy this property either: simply because the only paths in Φ' are geodesic. This demonstrates that partial mixture sets are more versatile than a disjoint union of mixture sets. Broadly speaking, whenever concerns relating to the design of an experiment lead us to omit certain paths, partial mixture sets may have an important role to play.

Cardinal and linear utility If paths in Φ are to be the building blocks, they had better have a meaningful structure. Since meaning is defined relative to our goal (a cardinal and linear utility representation), let us first define these concepts. A function $U : X \rightarrow \mathbb{R}$ is *linear* if, for every $\phi_{xy} \in \Phi$, the composite function $U \circ \phi_{xy} : I \rightarrow \mathbb{R}$ satisfies

$$(U \circ \phi_{xy})(\lambda) = (1 - \lambda)U(x) + \lambda U(y) \quad \text{for every } \lambda \in I.$$

A utility representation U with a certain property (such as linearity) is *cardinal* if every other utility V with the same property is related via a single *positive affine transformation*. In particular, $V = \theta U + \kappa$, where $\theta > 0$ and $\kappa \in \mathbb{R}$.

The following example highlights that, in absence of any form of continuity, even the requirement that X is a mixture set is very weak indeed.

Example 1 (A discrete mixture set). Let $X = \{x, y\}$ and suppose $x < y$. Let $\phi_{xy}(\lambda) = x$ if $\lambda < 1$ and $\phi_{xy}(1) = y$. Then, for every $\lambda \in I$, let $\phi_{yx}(\lambda) =$

[†]In fact, for the first line, we may appeal to \mathcal{P}_3 directly: for each $\lambda \in I$, $\phi_{xz}(\lambda) = \phi_{xy}(\lambda\mu)$ if and only if $\phi_{xz}(\lambda/\mu) = \phi_{xy}(\lambda)$ when $\mu > 0$. For the second, simply take $\nu = 1$ and, since $\mu < 1$, the resulting inversion yields $\lambda \mapsto (\lambda - \mu)/(1 - \mu)$ for $\mu \leq \lambda \leq 1$.

$\phi_{xy}(1 - \lambda)$, $\phi_{xx}(\lambda) = x$ and $\phi_{yy}(\lambda) = y$. It is not hard to check that these four paths are uniquely defined in such a way that X is a mixture set.

Clearly no utility representation of these preferences can be linear. The trouble is the fact that $\phi_{xy}(\lambda) = x$ for every $\lambda < 1$ combined with the fact that $x < y$. The first equality implies $(U \circ \phi_{xy})(\lambda) = U(x)$ for every $\lambda < 1$. Then the fact that U is a utility representation ensures that $(U \circ \phi_{xy})(\lambda)$ does not tend to $U(y)$ as λ tends to 1 even though $\phi_{xy}(1) = y$. This is in stark contrast with $(1 - \lambda)U(x) + \lambda U(y)$ which varies continuously between $U(x)$ and $U(y)$.

The essence of example 1 is that $U \circ \phi_{xy}$ is discontinuous at 1 and since X is discrete, the source of discontinuity is clearly ϕ_{xy} . We now clarify what is meant by continuity of ϕ_{xy} .

Continuous paths When X is a topological space, a path f is normally required to be continuous. That is, if F is closed in X , then the preimage $f^{-1}(F)$ is closed in I . Given \mathcal{O} , basic closed sets in the preference order topology are of the form $\{x : x \lesssim z\}$ and $\{x : z \lesssim x\}$. So in order to check f is continuous it suffices to check that $\{\lambda : f(\lambda) \lesssim z\}$ and $\{\lambda : z \lesssim f(\lambda)\}$ are closed subsets of I . For the case where X is a mixture set, this is precisely the form of continuity axiom that HM introduced. Our axiom is generalised only so as to accommodate the partial nature of Φ .

Axiom \mathcal{C} . For every $\phi \in \Phi$ and every $z \in X$, the sets $\{\lambda : \phi(\lambda) \lesssim z\}$ and $\{\lambda : z \lesssim \phi(\lambda)\}$ are closed in I .

Consequences of \mathcal{O} and \mathcal{C} The first and most basic consequence of \mathcal{O} and \mathcal{C} is that the image of each path ϕ in Φ is connected and compact in the preference order topology on X . Given that I is connected and compact and $\phi : I \rightarrow X$ is continuous, it is not surprising that $\phi(I)$ shares these properties. Since $\phi(I)$ is connected, it cannot be written as a pair of nonempty, disjoint closed sets of the form $\{x' : x' \lesssim x\}$ and $\{x' : y \lesssim x'\}$ such that $x < y$. This means that, for each $x, y \in \phi(I)$, if $x < y$, then $\phi(I)$ contains z such that $x < z < y$. This ensures that example 1 is ruled out by \mathcal{C} .

Building on this connectedness property, the following lemma shows that we are not confined to $\phi(I)$ such that $\phi \in \Phi$. The indifference relation allows us to make our first step towards putting the building blocks that are the paths in Φ together. This lemma corresponds to theorem 1 of HM.

Lemma 1. [†] Let preferences on X satisfy \mathcal{O} and \mathcal{C} . If $\phi_{xy} \in \Phi$, $z \in X$ and $x < z < y$, then there exists $0 < \mu < 1$ such that $z \sim \phi_{xy}(\mu)$.

Lemma 1 provides a basic existence requirement that is necessary for any linear representation U . Indeed, if $x < z < y$, then $(1 - \mu)U(x) + \mu U(y) = U(z)$ for some unique $0 < \mu < 1$. When, for any such z , there is a unique μ satisfying lemma 1, the set $\phi_{xy}(I)$ is *linearly ordered*: there is just one $z \in \phi_{xy}(I)$ such

[†]See page 25 for proof.

that $\phi_{xy}(\mu) \sim z$. Continuity of ϕ_{xy} together with the fact that $x < y$ then imply that $\phi_{xy}(\lambda) < \phi_{xy}(\lambda')$ if and only if $\lambda < \lambda'$. In this case, $\phi_{xy}(I)$ is a *linear continuum*: a linearly ordered, connected subset of X such that the infimum (and supremum) according to \preceq of any subset is uniquely defined. In particular, $\inf \phi_{xy} = x$ and $\sup \phi_{xy} = y$.

If $\phi_{xy}(I)$ is a linear continuum, then so is the image of $U \circ \phi_{xy}$. But since there is nothing in the axioms introduced so far to ensure that μ of lemma 1 is unique, further axioms are needed. In the setting where X is a partial mixture set, this turns out to be the most important consequence of the usual independence axiom.

Independence The following condition coincides with axiom 3 of HM when X is a mixture set: a concise form of the well-known *independence* axiom.

Axiom \mathcal{I} . If $\phi, \gamma \in \Phi$, $\phi(0) = \gamma(0)$ and $\phi(1) \sim \gamma(1)$, then $\phi(1/2) \sim \gamma(1/2)$.

Together \mathcal{O} , \mathcal{C} and \mathcal{I} are enough to ensure that a path ϕ_{xy} such that $x \sim y$ satisfies $\phi_{xy}(\lambda) \sim x$ for every $\lambda \in I$. This follows directly from theorem 2d of HM. Clearly, if $\phi \in \Phi$ is any other path such that $\phi(0) \sim \phi(1) \sim x$, then the same argument implies $\phi(\lambda) \sim \phi_{xy}(\lambda)$ for every $\lambda \in I$. Since this is a necessary condition for a linear utility representation of preferences, we see that \mathcal{I} has the intended effect: but only if we restrict attention to paths with endpoints that belong to a single indifference set. As we shall see, this is far from the case when $x < y$.

Before exploring this key issue, we use \mathcal{I} to improve on lemma 1 and show that $\phi_{xy}(I)$ is indeed a linear continuum when $x < y$. This lemma combines theorems 4 and 6 of HM.

Lemma 2. [†] Let preferences on X satisfy \mathcal{O} , \mathcal{C} and \mathcal{I} . If $\phi_{xy} \in \Phi$ and $x < y$, then $x < z < y$ if and only if there is a unique $0 < \mu < 1$ such that $z \sim \phi_{xy}(\mu)$.

When X is a mixture set lemma 2 is enough to yield a cardinal, linear utility representation. But, as the following example highlights, it is easy to see that \mathcal{I} is too weak when X is a partial mixture set.

Example 2 (Preferences satisfying \mathcal{I} with no linear utility representation). [‡]

Let $\phi, \gamma \in \Phi$, where ϕ is a path from x to x' and γ a path from y to y' . Moreover, suppose that $x \sim y < x' \sim y'$. Clearly every utility representation U satisfies $U(x) = U(y) = r$ and $U(x') = U(y') = r'$ for some $r < r' \in \mathbb{R}$. For U to be linear, we need ϕ and γ to be synchronised in such a way that $\phi(1/2) \sim \gamma(1/2)$. Only then do we have

$$(U \circ \phi)(1/2) = (U \circ \gamma)(1/2) = (r + r')/2.$$

The trouble is that, \mathcal{I} only applies when at least one of the endpoints of ϕ and γ coincide. In this case, $x = y$ or $x' = y'$. If Φ were to contain the path $\phi_{x'y}$,

[†]See page 25 for proof.

[‡]A similar example can be found at Karni and Safra [18, p.324].

then two applications of \mathcal{G} would ensure that each of ϕ and γ is synchronised with $\phi_{x'y}$. When X is a partial mixture set, there is no guarantee that such a path exists. As such, \mathcal{G} is compatible with preferences satisfying $\phi(1/2) < \gamma(1/2)$.

Example 2 highlights that the restriction of preferences satisfying \mathcal{G} to any subset of X that is a mixture set has a linear representation. It also shows that, in general, \mathcal{G} is insufficient for a linear representation on the whole of X . This is relevant to the empirical setting where data on two or more distinct mixture sets is collected separately. Even though \mathcal{G} holds on each of these, there is no guarantee that there is a linear utility representation on the union. In fact, this example provides a glimpse of a much deeper problem. One that goes to the heart of the relationship between the independence axiom and the mixture set structure. In the next example we define a partial mixture set upon which preferences may satisfy \mathcal{O} , \mathcal{C} and \mathcal{G} and have no utility representation.

Example 3 (A union of short lines). Let \mathbb{A}_+ denote a well-ordered set of the form $\{0, 1, 2, \dots\}$. Each element $a \in \mathbb{A}_+$ indexes a potential level of awareness of a trader, Val. We assume that \mathbb{A}_+ is the set of all countable ordinal numbers. Thus, for each $a \in \mathbb{A}_+$, the set $\{b \in \mathbb{A}_+ : b < a\}$ is countable, even though \mathbb{A}_+ itself is uncountable. (The situation is similar to the way that \mathbb{Z}_+ is countably infinite even though every one of its elements is finite.)

Each level of awareness is associated with its own mixture set X_a and $X = \bigsqcup\{X_a : a \in \mathbb{A}_+\}$. Since X is a disjoint union of mixture sets, proposition 3 ensures that it is a partial mixture set. The awareness structure we have in mind resembles that of Heifetz, Meier, and Schipper [12]. Let preferences on X satisfy \mathcal{O} , \mathcal{C} and \mathcal{G} .

Take each X_a to be the set of all functions from a state space S_a into a set of consequences C_a that satisfies the following properties for each $a \in \mathbb{A}_+$.

- (i) C_a is order isomorphic to \mathbb{R}_+ , where the linear ordering \lesssim of C_a is induced[†] from preferences over constant functions in X_a ;
- (ii) $\inf C_a \sim \inf C_{a+1}$;
- (iii) there exists $c_{a+1} \in C_{a+1}$ such that $c < c_{a+1}$ for every $c \in C_a$.

Property (i) ensures that X_a is a mixture set, where mixtures between functions are taken pointwise in C_a . From property (ii), we see that all the mixture sets contain a common lower bound.

It is property (iii) that makes this problem interesting. Yet it is also well-motivated since it seems reasonable to expect that awareness levels are payoff relevant, especially given the competitive nature of Val's work. Since \mathbb{A}_+ is uncountable there is no possibility of a utility representation U of preferences. This is because the image of U is a subset of \mathbb{R} : thus every collection of nonempty pairwise disjoint open subsets is countable.

For a and b that are separated by a limit ordinal in \mathbb{A}_+ , property (iii) and the assumption that X is a mixture set together imply that paths in X_a and X_b

[†]See Schmeidler [26] for an exposition of this step.

necessarily travel at different rates. The fact that X_a and X_b are disjoint renders \mathcal{G} impotent. We return to this example and explore this point in proposition 5.

The preferences identified in example 3 have no (real-valued) utility representation because they do not satisfy the *countable chain condition*: every pairwise disjoint collection of nonempty open preference intervals is countable. The surprising fact is that this is true even though preferences satisfy \mathcal{O} , \mathcal{C} and \mathcal{G} on each mixture set X_a , so that, by HM, on any such subset a cardinal, linear utility representation exists.

A stronger independence axiom The above difficulties may be overcome by strengthening \mathcal{G} . The condition we need must be capable of synchronising paths with endpoints that belong to the same indifference set.

Axiom \mathcal{S} . If $\phi, \gamma \in \Phi$, $\phi(0) \sim \gamma(0)$ and $\phi(1) \sim \gamma(1)$, then $\phi(1/2) \sim \gamma(1/2)$.

\mathcal{S} implies \mathcal{G} provided indifference is reflexive (for then $x = y$ implies $x \sim y$). As the proof of the following proposition shows, the converse is also true when indifference is transitive and X is a mixture set. This provides some justification for the claim that \mathcal{S} is a natural extension of the standard independence axiom to settings where Φ is only partially defined.

Proposition 4. [†] Let X be a mixture set and let preferences satisfy \mathcal{O} . Then \mathcal{G} holds if and only if \mathcal{S} does.

The next lemma shows that the preferences of example 2 are excluded by axiom \mathcal{S} . It may also be viewed as an improvement on lemma 2. In particular it goes some way towards showing that overlapping paths are synchronised.

Lemma 3. [‡] Let preferences on X satisfy \mathcal{O} , \mathcal{C} and \mathcal{S} . If $\phi_{xy} \in \Phi$ and $x < z < y$, then there is a unique $0 < \lambda < 1$ such that $\gamma(\lambda) \sim z$ for every $\gamma \in \Phi$ such that $\gamma(0) \sim x$ and $\gamma(1) \sim y$.

The implications of \mathcal{S} for synchronising paths go much further than lemma 3, indeed we will show that it rules out example 3. But both for this purpose and our derivation of a cardinal and linear utility, we need to know how to generate new paths from those in Φ .

Continuous concatenations of paths in Φ Condition \mathcal{P}_3 and proposition 1 tell us how a smaller path can be written in terms of a larger path. We now provide a way to construct a larger path, that is a concatenation of paths in Φ .

The basic concept of concatenation was introduced in (2). This allows us to write a path in Φ in terms of two subpaths. We now extend this idea to generate new paths in X from those in Φ . This extension is only possible if we let adjoining endpoints of paths in a concatenation be different. Whereas in (2) ϕ_{xz} and ϕ_{zy} were subpaths of ϕ_{xy} , we wish to concatenate ϕ_{xz} and $\phi_{z'y}$

[†]See page 25 for proof.

[‡]See page 26 for proof.

provided the weaker condition $z \sim z'$ holds. This generalisation clearly parallels the one that led us to give up \mathcal{G} in favour of \mathcal{S} .

A minor obstacle arises when $z \neq z'$. If f is the resulting concatenation, then since it is a function, it cannot satisfy $f(\mu) = z$ and $f(\mu) = z'$. To simplify the exposition, we adopt the convention that f coincides with ϕ_{xz} on the interval $[0, \mu)$ and with $\phi_{z'y}$ on $[\mu, 1]$.

Definition 2. $f : I \rightarrow X$ is a *concatenation in Φ* if there exists $\phi_0, \dots, \phi_m \in \Phi$ such that $\phi_n(1) \sim \phi_{n+1}(0)$, $f(1) = \phi_m(1)$ and

$$f(\nu) = \phi_n((\nu - \mu_n)/(\mu_{n+1} - \mu_n)) \quad (3)$$

for $0 = \mu_0 < \dots < \mu_{m+1} = 1$ and $\nu \in [\mu_n, \mu_{n+1})$.

When f is a concatenation such that $f(0) \sim x$ and $f(1) \sim y$, we simply say that f is a concatenation from x to y . If f is a concatenation of just one path, then $m = 1$ and $f \in \Phi$. But, regardless of whether $f \in \Phi$, the fact that f maps I into X implies that f is a path in X . Moreover, in the presence of \mathcal{O} and \mathcal{C} , the concatenations of definition 2 are continuous. In addition to the axioms, this conclusion follows from the restriction to concatenations with consecutive components that satisfy $\phi_n(1) \sim \phi_{n+1}(0)$. (As well as the fact that the union of finitely many closed sets is closed.)

We now refine our concept of (path) concatenation with a view to identifying those f that are synchronised with every path in Φ .

Synchronising concatenations A concatenation f from x to y is *synchronising* whenever every $\phi \in \Phi$ such that $x \lesssim \inf \phi$ and $\sup \phi \lesssim y$ satisfies

$$\phi(\lambda) \sim f((1 - \lambda)\mu + \lambda\nu) \quad \text{for every } \lambda \in I$$

for some $\mu, \nu \in I$ that are unique whenever $x < y$. Note that every $\phi \in \Phi$ is itself a concatenation. As such, the following lemma is a generalisation of proposition 1 with indifference replacing equality.

Lemma 4. [†] Let preferences on X satisfy \mathcal{O} , \mathcal{C} and \mathcal{S} . If Φ generates a concatenation from x to y , then it generates a synchronising concatenation from x to y .

Lemma 4 extends lemmas 1 to 3 from points z such that $x < z < y$ to paths ϕ such that $x \lesssim \phi(\lambda) \lesssim y$ for every $\lambda \in I$. Moreover, lemma 4 only requires a sequence of paths in Φ that connect x and y , so that ϕ_{xy} need not belong to Φ . It turns out that lemma 4 is the most important step in the proof of the main theorem. To this end, it remains for us to identify minimal conditions on preferences and Φ that guarantee synchronising concatenations between every $x, y \in X$ exist and are, up to indifference, unique. Before identifying these conditions, we end this section by using this new instrument to show that \mathcal{S} rules out example 3.

[†]See page 13 for proof.

Proposition 5. The partial mixture set and preferences of example 3 fail to satisfy \mathcal{S} .

PROOF OF PROPOSITION 5. Recall that in example 3 the order isomorphisms f_w were arbitrary. When \mathcal{S} holds, this is not the case. This axiom forces them to be synchronised. Our proof consists of assuming that \mathcal{S} holds and deriving contradiction. In particular, we show that contrary to the construction in example 3, either \mathcal{O} fails to hold, or, for some $w \in \mathbb{W}$, L_w is not a mixture set.[†] \square

An Archimedean, richness condition A very similar argument to the proof of proposition 5 can be used to show that there is no collection Φ of paths that makes the “long line” a mixture set. This ordered set is closely related to the construction of example 3 and the proof of proposition 5 and is formally introduced in the next example.

The long line is a well known example of a set that is path-connected. That is, for every pair x, y of its elements, there is a continuous path from x to y . In this way, the long line serves to distinguish mixture sets from path-connected sets. For if the long line were a mixture set, then the representation of HM would only be a local one: only once for small enough subsets of X do we have a (partially defined) linear utility representation.

The key to showing the long line is not a mixture set is to use condition \mathcal{P}_3 to obtain synchronising concatenations (with equality replacing indifference) that span the order. The proof has nothing to do with the axioms since $x \sim y$ if and only if $x = y$ when X is a line, and both \mathcal{C} and \mathcal{S} then hold trivially. When X is a partial mixture set, condition \mathcal{P}_3 is too weak to rule out the long line.

Example 4 (The long line as a partial mixture set). Consider the set \mathbb{A}_+ of example 3. Let $\mathbb{L}_+ = \mathbb{A}_+ \times_{\text{lex}} [0, 1)$ be the lexicographically ordered product where the first dimension is dominant. It is straightforward to find a paths such that \mathbb{L}_+ is a partial mixture set: let ϕ_{xy} be the convex combination on \mathbb{L}_+ that is defined if and only if, for some $a \in \mathbb{A}_+$ and $r, s \in [0, 1)$, both $x = a \times r$ and $y = a \times s$, and in this case $\phi_{xy}(\lambda) = a \times t$, where $t = (1 - \lambda)r + \lambda s$. It is clear that each order interval $[a \times 0, a \times 1)$ is a mixture set, and that \mathbb{L}_+ is a disjoint union of mixture sets indexed by the uncountable set \mathbb{A}_+ . It is straightforward to check that, when preferences coincide with \leq_{lex} , \mathcal{O} , \mathcal{C} and \mathcal{S} hold.

Example 4 confirms that we still do not have sufficient conditions for a utility representation, let alone one that is cardinal and linear. The key reason that the proof of proposition 5 does not apply is that for any x and y such that $x = w \times 0$ and $y = v \times 0$ and $w < v$, there is no concatenation from x to y . Even if $v = w + 1$, there is way to concatenate a finite collection ϕ_0, \dots, ϕ_m in Φ in such a way that $\phi_0(0) \sim x$ and $\phi_m(1) \sim y$.

We now provide an axiom that rules out example 4. It is stated in the weakest possible form, one that is easier to verify when, as in our application of section 3, we are trying to elicit utility or beliefs.

[†]The remainder of this proof appears on page 28

Axiom \mathcal{A} . If $x < y$, then there exists $\phi_0, \dots, \phi_m \in \Phi$ such that $\inf \phi_0 \lesssim x$, $y \lesssim \sup \phi_m$ and, for each $n < m$, $\inf \phi_{n+1} \lesssim \sup \phi_n$.

\mathcal{A} implies that every pair of prospects such that $x < y$ are connected by a finite chain of paths in Φ . In view of this, \mathcal{A} is an Archimedean condition. On the other hand, \mathcal{A} is clearly a completeness condition on the set of possible concatenations. That is to say, \mathcal{A} is also a richness condition on the set of paths in Φ . It goes without saying that, although this does place considerable structure on Φ relative to preferences, we are still far from requiring that X is a mixture set: there \mathcal{A} is satisfied with $m = 1$.

Lemma 5. [†] Let X be a partial mixture set and let preferences satisfy \mathcal{O} , \mathcal{C} , \mathcal{S} and \mathcal{A} . If $x < y$, then Φ generates a synchronising concatenation from x to y .

For the case where $x < y$, lemma 5 passes the main premise of lemma 4 to the axioms. The following simple example confirms that our axioms are still too weak to deliver a cardinal representation.

Example 5 (\mathcal{A} is too weak). Let X be the disjoint union of $\phi(I)$ and $\phi'(I)$, where $\phi, \phi' \in \Phi^<$ and suppose that, $\phi(1) \sim \phi'(0) =: x'$, so that \mathcal{A} also holds. Moreover, suppose that $U : X \rightarrow \mathbb{R}$ is a linear utility representation of preferences. Then, every element of $\phi'(I)$ other than x' strictly dominates every element of $\phi(I)$.

The trouble is that for any $0 < \mu < 1$, we may freely define a distinct mixture preserving concatenation f of ϕ and ϕ' such that $f(\mu) = x'$. This will not do for a cardinal representation, for each distinct pair f and g of such concatenations yields a pair of linear utility representations that are not related via a single positive affine transformation. Indeed, let $f(1/2) = x' = g(1/4)$, so that $x' < g(1/2)$. Now let $V := U \circ g \circ f^{-1}$. Then V is a well-defined linear utility by virtue of the fact that f is a bijection. Moreover, by construction, the image of V is the same as that of U . But, since $f^{-1}(x') = 1/2$, we have $(g \circ f^{-1})(x') = g(1/2)$, so that $V(x') = (U \circ g)(1/2)$. Since U is a utility representation, this number exceeds $U(x') = (U \circ g)(1/4)$.

Measuring utility at each $x \in X$ In the presence of the other axioms, \mathcal{A} ensures that the decision space is not too large relative to Φ . In particular, the preference ordering is spanned by a countable chain of paths in Φ : for every $x \in X$, $x \sim \phi(\lambda)$ for some $\phi \in \Phi$ and $\lambda \in I$. But example 5 demonstrates that this is not enough. We will now show that for cardinality, what is needed is that every $x \in X$ belongs to the relative interior of ϕ for some $\phi \in \Phi$. The timeless example of von Neumann and Morgenstern [29], upon which the following is based, helps to clarify this point.

Example 6. Suppose Val strictly prefers a glass of tea to a cup of coffee z in the afternoon. She also strictly prefers a cup of coffee to a plain glass of water x .

[†]See page 29 for proof.

Let ϕ denote the infusion over a five minute period of tea in a glass. Then von Neumann and Morgenstern's point is that, upto a reasonable approximation, we ought to be able to find a unique time $0 < \mu < 1$ such that $z \sim \phi(\mu)$. The same principle carries over to the setting where mixtures are only partially defined. If Val has a strong aversion to carbonated drinks, then it seems reasonable to suppose that any carbonation of the water (points on a path ϕ' from sparkling water to x) is strictly worse than x . But, if these are the only paths at our disposal, then example 5 tells us that we cannot measure utility on a single, cardinal scale. Equivalently, we cannot identify the strength of Val's preference for x relative to the worst prospect (sparkling water) and the best prospect (tea). The trouble is that preferences are such that x is not an interior point relative to any path of mixtures. In other words, we need to enrich the set of paths.

Examples 5 and 6 simply remind us that cardinality is purely an issue of measurement. For every prospect, there must be a suitable instrument with which to measure utility. Since our ultimate goal is to obtain a representation that is cardinal, the following axiom is unavoidable.

Axiom \mathcal{M} . If $x < z < y$, then $\inf \phi < z < \sup \phi$ for some $\phi \in \Phi$.

\mathcal{M} ensures that every interior point of the preference order is also in the relative interior of some path. \mathcal{M} is trivially satisfied when X is a mixture set, for fact that Φ is everywhere defined ensures that we can take $\phi = \phi_{xy}$. In view of this, the main innovation of \mathcal{M} is to formalise the fact that in practice we often have considerable flexibility in choosing the path with which to measure.

The following lemma improves on lemmas 1 to 5. It confirms that, together with the other axioms, \mathcal{M} allows us to measure the strength of preference of any given prospect z relative to any pair x and y such that $x < z < y$. The measuring instrument is a synchronising concatenation of paths in Φ .

Lemma 6. [†] Let X be a partial mixture set and let preferences satisfy \mathcal{O} , \mathcal{C} , \mathcal{S} , \mathcal{A} and \mathcal{M} . If $x < z < y$, then there is a unique $0 < \mu < 1$ such that $f(\mu) \sim z$ for every synchronising concatenation f from x to y that Φ generates.

The key contribution lemma 6 is the uniqueness of λ regardless of the choice of concatenation. This means that it does not matter how we “frame” the paths that form the synchronising concatenations, the strength of preference for z relative to x and y is the same. By virtue of its strength, Lemma 6 should provide a useful target for experimental testing in many settings. Any failure of preferences to satisfy its conclusion, can be investigated by testing the axioms individually to ascertain the source. This exercise is especially fruitful in light of the fact that the axioms are also necessary for the conclusion of the lemma. In fact, as our main theorem now shows if any of the axioms fails to hold, then there is no utility representation that is cardinal and linear.

[†]See page 29 for proof.

Theorem 1. [†] Let X be a partial mixture set. \mathcal{O} , \mathcal{C} , \mathcal{S} , \mathcal{A} and \mathcal{M} hold if and only if preferences have a cardinal and linear utility representation.

Unlike the case where X is a mixture set, the fact that \mathcal{A} and \mathcal{M} are necessary for such a representation is not immediately obvious. Examples 4 and 5 go some way towards demonstrating that they are. One final example completes the picture and also leads to a proof of the following novel characterisation of cardinality.

Corollary 1. [‡] Let X be a partial mixture set and let preferences have a linear utility representation U . Then U is cardinal if and only if both \mathcal{A} and \mathcal{M} hold.

3 Application: eliciting beliefs about stock prices

3.1 Basic facts on Wiener processes

Let $W_\bullet = \{W_t : t \in I\}$ be a standard Wiener process on an abstract probability space $(\Omega, \mathcal{F}, \mathbb{P})$. This means that $W_0 = a$ for some $a \in \mathbb{R}$ and W_\bullet is a family of random variables with continuous paths and stationary, independent increments $W_t - W_s$ that have distribution $N(0, t - s)$ for each $s \leq t$ in I .

By Bichteler [5, p.14], a standard Wiener process is associated with a certain *random path* on Ω . This is the unique path-valued random variable $\underline{W}_a : \Omega \rightarrow C_a(I; \mathbb{R})$ with Wiener measure as its probability distribution on $C_a(I; \mathbb{R})$. The latter is the space of continuous functions $w : I \rightarrow \mathbb{R}$ such that, for $t = 0$, $w_t = a$. We refer to a given realisation $w = W_a(\omega)$ as a *trajectory*. For each $t \in I$, \underline{W}_a is related to W_t via the *evaluation map* $w \mapsto w_t = \overline{W}_t(w)$. Explicitly, for each $t \in I$ and $\omega \in \Omega$,

$$W_t(\omega) = (\overline{W}_t \circ \underline{W}_a)(\omega).$$

The events on which Wiener measure is defined are of the form $E \subseteq C_a(I; \mathbb{R})$ such that $\underline{W}_a^{-1}(E) \in \mathcal{F}$. Basic events are those that are restricted only at finitely many times $t_1, \dots, t_m \in I$. They are known as *cylinder sets*:

$$E = \{w \in C_a(I; \mathbb{R}) : w_{t_n} \in A_n \text{ for some } A_1, \dots, A_m \subseteq \mathbb{R}\}.$$

The very simplest functions (also known as acts) that might be used to elicit beliefs in the Savage [25] and Anscombe and Aumann [2] framework are functions that take the value one E and zero on its complement. Even if the experimenter restricts attention only to events that are cylinder sets, we believe that even sophisticated subjects will struggle to compare two events E and E' by weighing up the factors that affect their probability: the number m versus the number m' of times each event is restricted; the distance from zero of times t_n and $t'_{n'}$; the relative measure in \mathbb{R} of each A_n and $A'_{n'}$; the relative centrality of the latter; and finally how to combine this information to make a judgement.

[†]See proof on page 30.

[‡]See proof on page 32.




Figure 2: A particularly simple cylinder set E contains all continuous paths in \mathbb{R} that begin at $a = 0$ and pass through intervals $A_1 = [b, c]$, $A_2 = [b, a]$ and $A_3 = [b, b]$ at times t_1 , t_2 and t_3 respectively. Of the four plotted trajectories, only the lower two belong to E . The other two do not pass through A_2 and A_3 .

The Brownian bridge The random paths that we will work with are a simple transformation of a Wiener process. Since a standard Wiener process is nonrandom at $t = 0$, is said to be *pinned* at a . A *Brownian bridge* B_\bullet is a stochastic process that is pinned at both ends of the interval I . That is $B_0 = a$ and $B_1 = b$. B_\bullet is related to W_\bullet as follows:

$$B_t = (1 - t) \left(a + \int_0^t \frac{1}{1 - r} dW_r \right) + tb \quad (4)$$

for each $t \in I$. Following common practice, reference to $\omega \in \Omega$ in (4) has been suppressed. It is the stochastic Itô integral in this expression that makes B_t a random variable on \mathbb{R} . In the case where $t = 1$, clearly $B_t = b$. For each $t < 1$, distributing the constant $1 - t$ through the parentheses and the integral, yields a bounded integrand since $1 - t \leq 1 - r$ for each $0 \leq r \leq t < 1$. The continuity properties of the Itô integral then ensure that this Brownian bridge is associated with a path-valued random variable $\underline{W}_{ab} : \Omega \rightarrow C_{ab}(I; \mathbb{R})$.

A minor extension of this definition yields the paths that will form our partial mixture set. Let $x = (s, a)$, $y = (t, b)$ for some $0 \leq s < t \leq 1$ and $a, b \in \mathbb{R}$. Let \underline{W}_{xy} denote be the random variable that takes values in $C_{ab}([s, t]; \mathbb{R})$. Then the resulting Brownian bridge between x and y is described by the composition of the evaluation map \bar{W}_λ with \underline{W}_{xy} for each $\lambda \in I$

$$\bar{W}_\lambda \circ \underline{W}_{xy} = (1 - \lambda) \left(a + \int_s^{s+\lambda(t-s)} \frac{t-s}{t-r} dW_r \right) + \lambda b.$$

Note that since the Itô integral is a square-integrable, real-valued random variable on $(\Omega, \mathcal{F}, \mathbb{P})$, so is $\overline{W}_\lambda \circ \underline{W}_{xy}$ for each $\lambda \in I$. And since the expectation with respect to \mathbb{P} of an Itô integral is zero,

$$\mathbb{E}(\overline{W}_\lambda \circ \underline{W}_{xy}) = (1 - \lambda)a + \lambda b. \quad (5)$$

We now use these random paths to construct a partial mixture set.

A partial mixture set of random variables Let Y denote a nonempty set of pairs (x, y) that give rise to a Brownian bridge of the form we have just described. Let

$$X_1 \stackrel{\text{def}}{=} \bigcup \{ \overline{W}_\lambda \circ \underline{W}_{xy} : \lambda \in I \text{ and } (x, y) \in Y \} .^\dagger$$

We construct the set Φ_1 of paths that make X_1 a partial mixture as follows. For each $(x, y) \in Y$, let $\phi_{xy} : I \rightarrow X_1$ satisfy $\phi_{xy}(\lambda) \stackrel{\text{def}}{=} \overline{W}_\lambda \circ \underline{W}_{xy}$. Let Φ_Y denote this collection of paths. Every path in Φ_Y has deterministic endpoints in X .

We now define a path in Φ_1 to be a subpath (in the sense of proposition 1) of some path in Φ_Y . In particular, a path $\phi : I \rightarrow \overline{W}_t$ belongs to Φ_1 if and only if, for every $\lambda \in I$, $\phi(\lambda) = \phi_{xy}((1 - \lambda)\mu + \lambda\nu)$ for some $(x, y) \in Y$ and some $\mu, \nu \in I$. By construction, each ϕ satisfies proposition 1, and this ensures that X is a partial mixture set. Clearly, X is not a mixture set since there is no path between $\phi(\lambda)$ and $\gamma(\mu)$ for any $\phi \neq \gamma$ in Φ_Y such that $0 < \lambda, \mu < 1$. Moreover, note in the experimental setting we have in mind, Y would be some finite set of pairs.

The following proposition shows that the typical path in Φ_1 is not convex. This confirms the need for the more general, nonconvex framework that partial mixture sets permit.

Proposition 6. [‡] For any $\phi_{xy} \in \Phi_1$ such that $x \neq y$, if $\mu < \nu$ and $0 < \lambda < 1$, then

$$\phi_{xy}((1 - \lambda)\mu + \lambda\nu) \neq (1 - \lambda)\phi_{xy}(\mu) + \lambda\phi_{xy}(\nu).$$

The partial mixture set X_1 will allow us to elicit Val's subjective first moments of her beliefs. For subjective moments of power k for $k = 2, 3, \dots$, we define X_k as follows. For each $\phi \in \Phi_1$, let ϕ^k denote the k th power of ϕ , where the power is taken pointwise: for each $\lambda \in I$, $\phi^k(\lambda) \stackrel{\text{def}}{=} (\phi(\lambda))^k$. Plainly, any such ϕ^k is a continuous, random path that takes values in $\mathbb{R}^{[s, t]}$ for some $0 \leq s < t \leq 1$.

Now let $\Phi_k \stackrel{\text{def}}{=} \{ \phi^k : \phi \in \Phi \}$ and let X_k be the union of the image sets $\phi^k(I)$ such that $\phi^k \in \Phi_k$. Clearly each (X_k, Φ_k) is a partial mixture set. And whilst the disjoint union over k of the X_k would also form a partial mixture set, that would not be the appropriate structure for eliciting subjective moments. We only want to synchronise paths that are of the same power.

[†]Strictly speaking, a typical element of X_1 is a pair consisting of a time $s + \lambda(t - s)$ and a real-valued random variable on Ω . This is consistent with $x = (s, a)$ and $y = (t, b)$ being the deterministic endpoints of \underline{W}_{xy} . As with ω , when this index is superfluous, it is suppressed.

[‡]See page 32 for proof.

Preferences The above construction of X_1, \dots is explicitly known to our experimenter, Zal. In contrast, the subject Val of the experiment, whose preferences are of interest, only has a subjective view of the stochastic process W_\bullet . In the canonical case, Val has observed a “large” number of days of trading of the underlying AAPL stock which is a standard Wiener process. Large means enough for her beliefs to have converged by the time the experiment begins.

Zal needs to check whether Val’s subjective moments are compatible with Wiener measure. Val is presented with pairs in X_k for $k = 1, 2, \dots$ and required to choose her preferred prospect (indifference is allowed). Val is assumed to understand that $\phi \in \Phi_Y$ is generated by sampling a path from the stochastic process that generates AAPL and revealing to her only the values of that path that correspond to some pair of times s and t such that $s < t$. She also knows that, as the owner of one unit of AAPL stock, she will receive the value $\phi^k(\lambda, \omega)$ if she chooses $\phi^k(\lambda)$ over $\gamma^k(\lambda)$ for some $\gamma \in \Phi_Y$. She also knows that by choosing $\gamma^k(\lambda)$ over $\phi^k(\lambda)$, she will receive $\gamma^k(\lambda, \omega)$.

We assume that \mathcal{C} holds for $X = \bigcup_k X_k$. In particular, we assume that Val owns a unit of the AAPL stock. For this reason, Val prefers higher prices. Thus, for any $a, b \in \mathbb{R}$, (degenerate) random variables $x, y \in X$ such that $x = (a, s)$ and $y = (b, t)$, $x < y$ if and only if $a < b$. In the present setting, this is not farfetched since I corresponds to a day’s trading on the stock market. That is to say, this also implies that Val does not discount time over the interval I . Similarly, we assume that Val is risk neutral. These simplifying assumptions allow us to focus entirely on the issue of eliciting beliefs.

Whereas Wiener process has continuous trajectories, \mathcal{C} only requires that the random paths in Φ are continuous relative to the order topology on X . Therefore \mathcal{C} allows for more general stochastic processes. If Val’s beliefs are characterised by Wiener measure, then she ought to be certain the process has no jumps. That is, if $\lim_n \lambda_n = \lambda$, then Val should view $\{\omega : \lim_n \phi(\lambda_n, \omega) \neq \phi(\lambda, \omega)\}$ as a null event. This requirement should be relatively easy to verify in an experimental setting. We simply assume this in addition to \mathcal{C} and leave to future work the exploration of any relationship between these two types of continuity.

Note that \mathcal{A} and \mathcal{M} are satisfied provided the pairs in Y are suitably chosen. In view of the assumptions we have already placed on preferences, this is a simple matter.

As usual, \mathcal{S} is the crucial axiom. If Val satisfies \mathcal{S} on X_1 , then theorem 1 yields a cardinal and linear utility representation U on X_1 . In view of the fact that Val is risk neutral, this tells us that Val’s first moments are correct. Since this utility representation is cardinal, we can find a positive affine transformation of U that is the standard mathematical expectation. We denote this again by U . Then, for any pair $x, y \in X_1$ and $\lambda \in I$, $(U \circ \phi_{xy})(\lambda)$ satisfies (5) as required for Val’s first moments to be correct.

For higher moments, the situation is not quite as straightforward. The issue is that the variance of a Brownian Bridge is nonlinear in λ and dependent on the size of the time increment. In particular, if $x = (s, a)$ and $y = (t, b)$ then

the true variance of $\phi_{xy}(\lambda)$ is

$$\text{Var}(\overline{W}_\lambda \circ \underline{W}_{xy}) = (\lambda - \lambda^2)(t - s).$$

To obtain the true expectation of $\phi_{xy}^2(\lambda)$, we simply add the square of $a + \lambda(b - a)$ to the variance. Assuming \mathcal{S} applies to paths in Φ_2 will not do. For taking $a = b$, implies $a^2 \sim b^2$ and, since $\phi_{xx}^2(\lambda) = a^2$ for every $\lambda \in I$. Then, \mathcal{O} and \mathcal{S} imply $\phi_{xy}^2(1/2) \sim a^2$, so variance plays no role whatsoever in determining preferences. In contrast, if Val's beliefs are correct, she would view the event $\{\omega : \phi_{xy}^2(1/2, \omega) > a^2\}$ as an event that occurs with probability one. The way to address the dependency on the time increment, is to constrain \mathcal{S} so that it applies only to paths $\phi_{xy}^2, \phi_{x'y'}^2 \in \Phi_2$ such that the corresponding time increments satisfy $t - s = t' - s'$.

Thus, suppose the certainty equivalent of $x, x' \in X_2$ is a and that of $y, y' \in Y_2$ is b . Then since $t - s = t' - s'$, Val's raw subjective second moments are correct provided

$$\phi_{xy}^2(1/2) \sim \phi_{x'y'}^2(1/2) \sim \frac{(t - s) + (a + b)^2}{4}$$

for every such x, x', y and y' .

Whilst this relaxation of \mathcal{S} allows for higher-order subjective moments of $\phi(\lambda)$ that are nonlinear in λ , it also appears to constitute a departure from the linear model that we have derived. We now demonstrate how to exploit the partial mixture set structure in order to extend the linear model to this case.

Consider the case where $a = b = 0$. The more general case is a straightforward extension. Then let

$$\xi \mapsto \lambda = \frac{1}{2} - \frac{\sqrt{1 - \xi}}{2},$$

The following extension of $\lambda \mapsto \phi(\lambda)$

Existence and uniqueness of beliefs Existence of a continuous cumulative distribution function on \mathbb{R} is characterised by the requirement that the k th moment m_k of any given real-valued random variable $x \in X_1$ satisfies

$$\begin{vmatrix} 1 & m_1 & \dots & m_k \\ m_1 & m_2 & \dots & m_{k+1} \\ \vdots & \vdots & \ddots & \vdots \\ m_k & m_{k+1} & \dots & m_{2k} \end{vmatrix} > 0 \quad \text{for } k = 0, 1, 2, \dots$$

Provided this additional condition is satisfied, for each nondegenerate $x \in X_1$, Val's subjective moments are characterised by a nonempty set of measures.

The collection of Val's subjective moments may not uniquely identify a single probability measure, be it Wiener or any other.

A sufficient condition for uniqueness is the Carleman condition $\sum_k m_{2k}^{-1/2k} = \infty$ (see Billingsley [6] for a related condition). In essence, this states that the

measure is unique if the rate of growth of the even moments is slow enough. A sufficient condition for nonuniqueness is the Krein condition. See the classic reference Akhiezer [1] and the more recent Stoyanov [28] for more in depth discussions of this topic, which is known as the *Hamburger moment problem*.

Entropy Of course in practice, there will be some finite k' such that the experimenter stops eliciting higher moments. Then clearly Val's beliefs will be underidentified. This is a practical concern, and it is worth noting that the situation would be the same if the experimenter chose to elicit beliefs using events instead. The hope is that more information may be elicited from the lower moments than from the functions on cylinder sets described above.

Having elicited the first k' moments for a representative sample of X_1 , the experimenter finds himself in a position that is rather common in physics and statistics. He has rich information about lower moments and no information about higher moments. If he adopts a conservative stance about Val's beliefs, he should model them with the probability distribution that *maximises the entropy* in the system.

We now explain how he might go about this. Let $F_x(k')$ denote the set of possible densities on \mathbb{R} that are compatible with his observations of Val's first k' moments for a given $x \in X_1$. For each $f \in F_x(k')$, let

$$H(f, 1) \stackrel{\text{def}}{=} \int_{\mathbb{R}} f (\ln f - \ln 1) dr.$$

The argument $\underline{f} \in F_x(k')$ that minimises $H(\cdot, 1)$ is the closest density to the uniform. This identifies the minimum relative entropy distribution, which coincides with the one that maximises entropy. The usual way to identify \underline{f} is to set up a Lagrangian using the function H and the constraints that characterise $F_x(k')$. (The latter are the subjective moment conditions of Val's preferences.)

For the case where $k' = 2$, the maximum entropy distribution is none other than the normal distribution with mean and variance satisfying Val's first two subjective moments. Thus, when Y is chosen to be a representative sample, the definition of Wiener measure is such that it is the unique maximiser.

More generally, when \underline{f} exists, it is the worst case model of Val's information that the data allow. However, because the uniform density on \mathbb{R} does not correspond to a probability distribution (its integral is infinite), there may be no solution \underline{f} to the maximum entropy problem. For some interesting examples and a more complete discussion of this issue see Rockinger and Jondeau [24]. Clearly, a sufficient condition is that $F_x(k')$ is compact, for H is continuous in f .

Whilst maximising the entropy is one of the most popular approaches to dealing with complex scenarios involving uncertainty and underidentification, it is not the only one. At the other extreme, the experimenter may be in a setting where it is better to give Val the benefit of the doubt. If, as in our example above, the experimenter knows the true distribution of each $x \in X_1$, then he may minimise relative to the true density f' . That is via the same constrained

minimisation procedure, with $H(\cdot, f')$ replacing $H(\cdot, 1)$, we obtain a density \bar{f} that represents the best case for Val's beliefs. \bar{f} represents the least conservative model. In contrast with \underline{f} , \bar{f} always exists and is equal to f' . Clearly, if f' is feasible, so that it belongs to $F_x(k')$, this means that the first k' subjective moments that Val revealed were correct.

Finally, for any $0 < \alpha < 1$, the convex combination $(1 - \alpha)\underline{f} + \alpha\bar{f}$ is an intermediate model of Val. This procedure for finding a distribution that represents Val's beliefs given her subjective moments, extends to the full collection of maximum and minimum entropy densities \underline{f}_x and \bar{f}_x such that $x \in X_1$.

4 Discussion

Related axioms in the literature We now describe the relationship between condition \mathcal{S} and the corresponding condition of Karni and Safra [18], Karni [17], and Grant et al. [11]. These papers take X to be of the form $A \times L$ for a compact set A and a mixture set L . (In Grant et al. [11], A is also a mixture set.) The axiom corresponding to \mathcal{S} is translated as follows

Suppose that $\phi, \gamma \in \Phi$ and $\phi(0) \sim \gamma(0)$. Then $\phi(1) \lesssim \gamma(1)$ if and only if $\phi(\lambda) \lesssim \gamma(\lambda)$ for every $\lambda \in (0, 1]$.

The fact that this axiom implies \mathcal{S} follows immediately if we take $\lambda = 1/2$. (Consider the fact that \sim is a subset of \lesssim and use the “only if” part of the statement.) Step 3 of the proof of lemma 4 shows that, together with \mathcal{O} and \mathcal{C} , \mathcal{S} is sufficient for this axiom.

The translation of axiom E2 of Fishburn [8, p.88] to the present notation is a slightly weaker version of the above: if $\phi, \gamma \in \Phi$, $\phi(0) \sim \gamma(0)$ and $\phi(1) < \gamma(1)$, then $\phi(\lambda) < \gamma(\lambda)$ for every $0 < \lambda < 1$. If either of these “independence” axioms is assumed in the place of \mathcal{S} , then theorem 1 holds when the following continuity axiom is assumed in the place of \mathcal{C}

If $\phi_{xy} \in \Phi$, $z \in X$ and $x < z < y$, then there are $0 < \mu, \nu < 1$ such that $\phi_{xy}(\mu) < z < \phi_{xy}(\nu)$.

Finally, \mathcal{A} and \mathcal{M} do not appear to have featured as axioms in the literature. The Archimedean axiom of Gilboa and Schmeidler [10, 9] is clearly similar to \mathcal{A} even if the structure of preferences is quite different. Karni [17] contains a background assumption which is closely related and indeed sufficient for \mathcal{A} and \mathcal{M} combined. Karni views this as a property of the set A of acts and refers to this property as “linked”. Similarly an appeal to a condition involving a finite sequence of overlapping mixture sets is made in proposition 6 of Grant et al. [11].

To our knowledge, nowhere is it shown that these are part of a necessary and sufficient group of axioms for a cardinal and linear utility representation. Moreover, as the initial discussion HM highlights, our set up allows for an infinite dimensional partial mixture set, where the linear operator that preferences define is unbounded and hence discontinuous. (See Fishburn [7], Inoue [15],

and Monteiro [20] for examples and further discussions of this point.) Only the multilinear utility model of Fishburn [8] is comparable in this latter sense.

4.1 Related experiments in literature

Experiments that have conducted related exercises include Stecher, Shields, and Dickhaut [27]. There, agents are shown realisations of processes that do not have bounded moments. Effectively they consider processes that lie beyond the method we propose. There is an active debate in the finance literature as to whether actual stock prices have moments of high order, but there seems to be little doubt about the lower order moments. In particular first two moments : the expected return and volatility of return (see Barndorff-Nielsen and Shephard [3] for a detailed discussion and further references to the literature). In econometric models, even conservative models require the first four moments (see for instance Hong, Linton, and Zhang [14]).

In any case, it seems plausible that an agent would trade any random variable that an experimenter might offer for a finite amount of money. The fact that we consider random bridges between deterministic points makes this case even more convincing. Finally, whilst our example is financial, the main motivation remains the same in other settings where the dimension of the state space is large and events are difficult to describe in an intuitive way. In such settings, we suggest that moments offer a plausible, alternative approach to eliciting beliefs.

A Proofs

PROOF OF PROPOSITION 1 (FROM PAGE 5). It suffices to show that for every $x', y' \in \phi(I)$ the path $\phi_{x'y'}$ belongs to Φ . For every such x' and y' , there exists $\mu, \nu \in I$ such that $\phi_{xy}(\mu) = x'$ and $\phi_{xy}(\nu) = y'$. By \mathcal{P}_3 , $\phi_{xx'}, \phi_{xy'} \in \Phi$ and $\phi_{xx'}(\lambda) = \phi_{xy}(\lambda\mu)$ and $\phi_{xy'}(\lambda) = \phi_{xy}(\lambda\nu)$ for every $\lambda \in I$. W.l.o.g., suppose that $\mu \leq \nu$.

If $\mu = \nu$, then $x' = y'$. Moreover, by \mathcal{P}_2 , $\phi_{x'x}(\lambda) = \phi_{xx'}(1 - \lambda)$ and by \mathcal{P}_1 , $x' = \phi_{x'x}(0)$. Then one further application of \mathcal{P}_3 yields $\phi_{x'y'}(\lambda) = \phi_{x'x}(\lambda 0) = x'$ for every $\lambda \in I$. In this case, clearly, (1) holds with $\mu = \nu$ for every $\lambda \in I$.

If $\mu < \nu$, then, for some $\lambda' < 1$, $\mu = \lambda'\nu$. Then $\phi_{xy'}(\lambda') = x'$ and since $\phi_{y'x}(1 - \lambda') = \phi_{xy'}(\lambda')$, we see that $x' = \phi_{y'x}(1 - \lambda')$. A final application of \mathcal{P}_3 yields $\phi_{y'x'}(\lambda) = \phi_{y'x}(\lambda(1 - \lambda'))$ for each $\lambda \in I$. Next note that $\lambda' = \mu/\nu$, so that a substitution for λ' and straightforward simplification yields $\phi_{y'x'}(\lambda) = \phi_{y'x}(\lambda(\nu - \mu)/\nu)$ for each $\lambda \in I$. Similarly, we have

$$\begin{aligned} \phi_{y'x}(\lambda(\nu - \mu)/\nu) &= \phi_{xy'}(1 - \lambda(\nu - \mu)/\nu) && \text{by } \mathcal{P}_2 \\ &= \phi_{xy}(\nu - \lambda(\nu - \mu)) && \text{by } \mathcal{P}_3. \end{aligned}$$

Let $\lambda \mapsto \kappa = 1 - \lambda$. Substituting for λ we have $\phi_{y'x'}(1 - \kappa) = \phi_{xy}(\mu + \kappa(\nu - \mu))$. One final application of \mathcal{P}_2 to the left-hand-side of the latter equality yields both $\phi_{x'y'}$ and (1) \square

PROOF OF PROPOSITION 2 (FROM PAGE 7). By assumption, $X \subset \mathbb{R}^3$ is parameterised by equations for a standard sphere. That is $x_1 = r \cdot \cos(2\pi\nu) \cdot \sin(\pi\kappa)$, $x_2 = r \cdot \sin(2\pi\nu) \cdot \sin(\pi\kappa)$ and $x_3 = r \cdot \cos(\pi\kappa)$, where r is the radius of the sphere, $\nu, \kappa \in I$, $\pi = 3.14\dots$, and \cdot denotes standard scalar multiplication. This parametrisation centers the sphere at the origin $(0, 0, 0)$.

For every $x, y \in X$, Let ϕ_{xy} be the geodesic (the shortest path from x to y on X). Conditions \mathcal{P}_1 and \mathcal{P}_2 of the definition of a mixture set are immediately satisfied. We now show that condition \mathcal{P}_3 holds for any path ϕ_{xy} that belongs to a typical “great circle” on X . Such circles are parametrised by an injective function $f : [0, 1) \rightarrow X$. As a canonical example, let $\kappa = 1/2$, so that $x_3 = 0$. Then the set of points $(x_1, x_2) = f(\nu) := r \cdot (\cos(2\pi\nu), \sin(2\pi\nu))$ such that $\nu \in [0, 1)$, define a standard circle in \mathbb{R}^2 as well as a “great circle” on the earth’s surface.

Let $x \neq y$ on this great circle. Then for some $\nu_x, \nu_y \in [0, 1)$, $x = f(\nu_x)$ and $y = f(\nu_y)$. Without loss of generality let $\nu_x < \nu_y$. Note that $\phi_{xy}(\lambda) = f((1-\lambda)\nu_x + \lambda\nu_y)$ for each $\lambda \in I$. That is, f provides a parametrisation of ϕ_{xy} . Suppose that $\phi_{xy}(\mu) = z$ for some $\mu \in I$. Then $f(\nu_z) = z$ where $\nu_z = (1-\mu)\nu_x + \mu\nu_y$. Recall that for condition \mathcal{P}_3 of mixture sets, we need to show $\phi_{xz}(\lambda) = \phi_{xy}(\lambda\mu)$. This follows by substituting for ν_z in $f((1-\lambda)\nu_x + \lambda\nu_z)$ and exploiting the convexity of $[0, 1)$. \square

PROOF OF LEMMA 1 (FROM PAGE 9). Consider the the set $L = \{\lambda : \phi_{xy}(\lambda) \lesssim z\}$. By condition \mathcal{C} , L is a closed subset of I . It is nonempty since $\phi_{xy}(0) = x < z$. Similarly, the set $U = \{\lambda : z \lesssim \phi_{xy}(\lambda)\}$ is closed and nonempty since $z < \phi_{xy}(1) = y$. By \mathcal{O} , I is the union of L and U . If $L \cap U = \emptyset$, then I is the union of two disjoint, nonempty and closed subsets. Since this would imply that I is disconnected, this intersection is necessarily nonempty. That is, there exists $\mu \in I$ such that $\phi_{xy}(\mu) \sim z$. We have already seen that $\mu \neq 0, 1$. \square

PROOF OF LEMMA 2 (FROM PAGE 10). Suppose that $x < z < y$ for some $z \in X$. Then by lemma 1, there exists at least one μ satisfying the required condition. Suppose that there exists $\mu' < \mu$ such that $\phi_{xy}(\mu') \sim z$. Then by \mathcal{O} , we have $\phi_{xy}(\mu') \sim \phi_{xy}(\mu)$. But since proposition 1 ensures $\phi_{xy}(I)$ is a mixture set and $x < y$, theorem 4 of HM implies that $\phi_{xy}(\mu') < \phi_{xy}(\mu)$ if and only if $\mu' < \mu$. Note that the latter theorem only applies because \mathcal{G} now holds.

Now suppose that $z \lesssim x < y$ and $z \sim \phi_{xy}(\mu)$ for some unique $0 < \mu < 1$. Let $z' = \phi_{xy}(\mu)$. Then \mathcal{P}_3 ensures that $\phi_{xz'}(\lambda) = \phi_{xy}(\lambda\mu)$ for every $\lambda \in I$. Indeed, proposition 1 ensures that $\phi_{z'y}(\lambda) = \phi_{xy}((1-\lambda)\mu + \lambda 1)$. \mathcal{O} implies $z' \lesssim x < y$ and the first part of this proof ensures that $\phi_{z'y}(\nu) \sim x$ for some unique $0 \leq \nu < 1$. Let $x' = \phi_{z'y}(\nu)$. Then $x' = \phi_{z'y}(\nu) = \phi_{xy}((1-\nu)\mu + \nu)$. Let $\nu^* = (1-\nu)\mu + \nu$. Then $\phi_{xy}(\lambda) \sim x$ for both $\lambda = 0$ and $\lambda = \nu^*$. Clearly $0 < \mu$ implies $0 < \nu^*$: another contradiction of theorem 4 of HM. \square

PROOF OF PROPOSITION 4. Let $x \sim x'$ and $y \sim y'$. Since X is a mixture set, the paths between pairs of elements are defined. \mathcal{G} implies both $\phi_{xy}(1/2) \sim \phi_{xy'}(1/2)$ and $\phi_{y'x}(1/2) \sim \phi_{y'x'}(1/2)$. Since $\lambda = 1 - \lambda = 1/2$, property \mathcal{P}_2 mixture

sets implies both $\phi_{xy'}(1/2) = \phi_{y'x}(1/2)$ and $\phi_{x'y'}(1/2) = \phi_{y'x'}(1/2)$. Finally, Θ (in particular transitivity) ensures that $\phi_{xy}(1/2) \sim \phi_{x'y'}(1/2)$. \square

PROOF OF LEMMA 3 (FROM PAGE 12). Let $\phi = \phi_{xy}$. Lemma 2 guarantees the existence of a candidate $0 < \lambda < 1$ such that $\phi(\lambda) \sim z$. Take γ to be any other path in Φ satisfying $\gamma(0) \sim x$ and $\gamma(1) \sim y$. \mathcal{S} ensures that $\phi(1/2) \sim \gamma(1/2)$. Condition \mathcal{P}_3 then ensures the existence of a subpath $\phi_{0\frac{1}{2}}$ from x to $\phi(1/2)$ such that $\phi_{0\frac{1}{2}}(\lambda) = \phi(\lambda/2)$ for every $\lambda \in I$. A similar path $\gamma_{0\frac{1}{2}}$ exists from x' to $\gamma(1/2)$. An application of \mathcal{S} yields $\phi_{0\frac{1}{2}}(1/2) \sim \gamma_{0\frac{1}{2}}(1/2)$. This implies that $\phi(1/4) \sim \gamma(1/4)$. An application of condition \mathcal{P}_2 of partial mixture sets and a similar argument applied to the paths $\phi_{\frac{1}{2}1}$ and $\gamma_{\frac{1}{2}1}$ yields $\phi(3/4) \sim \gamma(3/4)$. (Using proposition 1 to translate indifferences on subpaths to indifferences between ϕ and γ .) In this way, the above argument yields $\phi(\rho) \sim \gamma(\rho)$ for every dyadic rational $\rho \in I$. Then, since the dyadic rationals are dense in I , there exists a sequence $\lim_n \rho_n = \lambda$, where recall λ is our candidate for the proof. W.l.o.g., we may take the sequence to be increasing. Then by the proof of lemma 2, $\phi(\rho_n) \preceq \phi(\lambda)$ for each n . \mathcal{C} then ensures that λ belongs to $\{\lambda' : \gamma(\lambda') \preceq \phi(\lambda)\}$. Repeating the argument with the roles of ϕ and γ reversed yields the reverse weak preference, so that $\phi(\lambda) \sim \gamma(\lambda)$, as required. \square

PROOF OF LEMMA 4 (FROM PAGE 13). If $x \sim y$, then the fact that Φ generates a concatenation from x to y ensures that $\phi(0) \sim x$ for some $\phi \in \Phi$. Proposition 1 ensures that $\phi_{x'x'} \in \Phi$ for some x' such that $x' \sim x \sim y$. Then by the discussion immediately following \mathcal{G} , $\phi_{x'x'}(\lambda) \sim x$ for every $\lambda \in I$. Thus, $\phi_{x'x'}$ is a synchronising concatenation from x to y .

The remainder of the proof accounts for the case $x < y$. To this end, let $\Phi^< = \{\phi : \phi(0) < \phi(1)\}$.

STEP 1 (*There exists a minimal, strictly increasing concatenation f from x to y*). Recall from proposition 1 that every $\phi \in \Phi$ has the property that its image $\phi(I)$ is a mixture set. Then theorem 4 and 5 of HM ensure that every $\phi \in \Phi$ is of one of the following three types: (i) $\phi(\lambda) \sim \phi(\lambda')$ for every $\lambda, \lambda' \in I$; (ii) $\phi(\lambda) < \phi(\lambda')$ if and only if $\lambda < \lambda'$; or (iii) $\phi(\lambda') < \phi(\lambda)$ if and only if $\lambda < \lambda'$. Note that in the presence of the axioms assumed, ϕ is a type (ii) path if and only if $\phi \in \Phi^<$.

By assumption, there is a concatenation f' from x to y . Let f' be a concatenation of m' paths. Since the nonnegative integers are well-ordered there exists a concatenation f from x to y of the smallest possible number $m \leq m'$ of paths in Φ . Clearly f is concatenation only of paths in $\Phi^<$, for otherwise, we could exclude a path and obtain a suitable concatenation with even fewer paths. This completes the proof of step 1.

Let $\phi_1, \dots, \phi_m \in \Phi^<$ and $0 = \mu_0 < \dots < \mu_{m+1} = 1$ be the sequences that characterise f .

STEP 2 (*f is synchronised with ϕ_1, \dots, ϕ_m*). Take any $n = 0, \dots, m$. Recall the transformation T that was introduced in the derivation of 3. Since $\mu_n < \mu_{n+1}$, the inverse $T^{-1}(\nu) = (\nu - \mu_n)/(\mu_{n+1} - \mu_n)$ is defined for each $\nu \in [\mu_n, \mu_{n+1})$.

Then, for each $\lambda < 1$, $T(\lambda) = \lambda(\mu_{n+1} - \mu_n) + \mu_n$ so that $\phi_n(\lambda) = (f \circ T)(\lambda)$ as required. Finally, for $\lambda = 1$, $\phi_n(\lambda) \sim \phi_{n+1}(0) = f(\mu_{n+1})$.

STEP 3 (ϕ_n is synchronised with $\phi \in \Phi^<$ if $\inf \phi_n \lesssim \inf \phi$ and $\sup \phi \lesssim \sup \phi_n$). If $\phi(0) \sim f(\mu_n)$ and $\phi(1) \sim f(\mu_{n+1})$, then the fact that ϕ is synchronised with ϕ_n follows directly by lemma 3. Now suppose that $f(\mu_n) \lesssim \phi(0)$ and $\phi(1) \lesssim f(\mu_{n+1})$, with at least one relation holding strictly. Lemma 2 implies $\phi_n(\mu) \sim \phi(0)$ and $\phi(1) \sim \phi_n(\mu')$ some unique $\mu, \mu' \in I$ such that either $0 < \mu$ or $\mu' < 1$.

Since $\phi \in \Phi^<$, $\mu < \mu'$ follows by theorem 4 of HM and the fact that $\phi_n \in \Phi^<$. Write ϕ_n as a concatenation of (at most) three subpaths $\phi', \phi'', \phi''' \in \Phi^<$ such that in particular $\phi_n(\nu) = \phi''((\nu - \mu)/(\mu' - \mu))$ for each $\mu \leq \nu \leq \mu'$. (If $\mu = 0$ or $\mu' = 1$, there are just two paths.) Then $\phi''(\lambda) = \phi_n((1 - \lambda)\mu + \lambda\mu')$ for every $\lambda \in I$. Then $\phi''(0) \sim \phi(0)$ and $\phi''(1) \sim \phi(1)$, and lemma 3 ensures $\phi''(\lambda) \sim \phi(\lambda)$ for every $\lambda \in I$. In turn, we see that ϕ is synchronised with ϕ_n .

STEP 4 (f is synchronised with ϕ of step 3). Recall that $\phi_n(\xi) \sim f((1 - \xi)\mu_n + \xi\mu_{n+1})$ for every $\xi \in I$. For each $\lambda \in I$, let $\xi = (1 - \lambda)\mu + \lambda\mu'$, where μ and μ' were defined in step 3. Then a substitution and straightforward rearrangement yields $\phi(\lambda) \sim f((1 - \lambda)\mu_* + \lambda\mu^*)$ where $\mu_* = (1 - \mu)\mu_n + \mu\mu_{n+1}$ and $\mu^* = (1 - \mu')\mu_n + \mu'\mu_{n+1}$. Since $\mu_n < \mu_{n+1}$ and $\mu < \mu'$, it is clear that $\mu_* < \mu^*$ as required for f to be synchronising.

STEP 5 (f is synchronised with ϕ if $\phi(0) < f(\mu_n) < \phi(1)$ for some n). By step 1, m is minimal and this ensures that $f(\mu_{n-1}) \lesssim \phi(0)$ and $\phi(1) \lesssim f(\mu_{n+1})$ (with at least one relation holding strictly). The difficulty here is that the μ_n may be incorrectly specified, so that f travels at a different rate on distinct intervals $[\mu_{n-1}, \mu_n)$ and $[\mu_n, \mu_{n+1})$. Example 5 (from page 15) demonstrates the degree of freedom we have in choosing μ_n . We now show that we can always respecify the μ_n and obtain a new concatenation that is synchronised with ϕ . Since the resulting concatenation g is composed of the same paths as f , it passes step 2 and step 3 of this proof.

Let $1 \leq n \leq m$ be the smallest number such that $\phi(0) < f(\mu_n) < \phi(1)$ for some ϕ . By lemma 2, there is a unique $0 < \lambda_n < 1$ such that $\phi(\lambda_n) \sim f(\mu_n) = \phi_{n+1}(0)$. Let f_n denote the subconcatenation of f such that $f_n(\lambda) = f(\lambda\mu_n)$ for each $\lambda < 1$ and $f_n(1) = \phi_n(1)$. That is, f_n is a concatenation of ϕ_0, \dots, ϕ_n . Then lemma 2 ensures that $f_n(\kappa_*) \sim \phi(0)$ and $\phi_n(\kappa^*) \sim \phi(1)$ for some unique $0 \leq \kappa_* < 1$ and $0 < \kappa^* \leq 1$. (Note that $\kappa_* = 0$ if and only if $\phi(0) \sim x$ and $\kappa^* = 1$ if and only if $\phi(1) \sim f(\mu_{n+1})$ and, since m is minimal, both do not hold simultaneously.)

For every $\lambda \in [\mu_{n+1}, 1]$, take g to satisfy $g(\lambda) = f(\lambda)$. On the interval $[0, \mu_{n+1})$, g will be the concatenation of f_n, ϕ_{n+1} , but unless f is synchronised with ϕ to begin with, $g \neq f$.

Recall that $\phi(\lambda_n) \sim f(\mu_n)$. We seek g such that $g(\nu_n) \sim \phi(\lambda_n)$ where

$$\nu_n = (1 - \lambda_n)\mu_* + \lambda_n\mu^* \tag{6}$$

for some unique $\mu_* < \mu^*$ such that $g(\mu_*) \sim \phi(0)$ and $g(\mu^*) \sim \phi(1)$. Since g is to

be concatenation of f_n on $[0, \mu_n)$, we also require $f_n(\lambda) = g(\lambda\nu_n)$ for every $\lambda < 1$. The latter equality together with the indifferences $g(\mu_*) \sim \phi(0) \sim f_n(\kappa_*)$ yields the equation $\mu_* = \kappa_*\nu_n$. Similarly, from the concatenation relation between f and ϕ_{n+1} we obtain the equation $\mu^* = (1 - \kappa^*)\nu_n + \kappa^*\mu_{n+1}$. Substituting for μ_* and μ^* in (6) and solving for ν_n we find

$$\nu_n = \frac{\lambda_n \kappa^* \mu_{n+1}}{(1 - \lambda_n)(1 - \kappa_*) + \lambda_n \kappa^*}.$$

Now since $\kappa_* < 1$, $0 < \lambda_n < 1$ and $0 < \kappa^*, \mu_{n+1}$, a suitable ν_n exists and uniquely so. By construction, g is synchronised with ϕ .

STEP 6 (g is synchronised with every remaining $\phi \in \Phi$). By the preceding steps in this proof, g is synchronised with every ϕ' such that $g(0) \lesssim \phi'(0)$ and $\phi'(1) \lesssim g(\mu_{n+1})$. Note that the definition of g is such that its rate is, in general, different from f on the entire interval $[0, \mu_{n+1})$. In particular, for each $n' \leq n$, $\phi_{n'}(0) \sim g(\nu_{n'})$, where $\nu_{n'} = (\nu_n / \mu_n) \cdot \mu_{n'}$.

A similar argument applies to μ_{n+1} if $\phi'(0) < g(\mu_{n+1}) < \phi'(1)$ for some ϕ' . Since m is finite, we exhaust the set of μ_n such that $1 \leq n \leq m$ in a finite number of such steps. For $n = 0$ and $n = m + 1$, where μ_n equals 0 and 1 respectively, a very similar argument can also be implemented to ensure the resulting concatenation is synchronised with every $\phi \in \Phi^<$ such that $x < \phi(\lambda) < y$ for some $0 < \lambda < 1$. Let g denote this latter concatenation.

For every other $\phi \in \Phi^<$, $\phi(1) \lesssim x$ or $y \lesssim \phi(0)$. We adopt the convention that g is trivially synchronised with $\phi \in \Phi^<$ if there is at most one indifference set that is common to them both.

If $\phi(0) \sim \phi(1)$, then, since $\phi(I)$ is a mixture set, theorem 5 of HM ensures that ϕ is of type (i). Then $\mu = \mu'$ and the proof of this case is immediate. Finally, condition \mathcal{P}_2 accounts for every ϕ such that $\phi(1) < \phi(0)$, so that ϕ is of type (iii). Since these three cases exhaust Φ , our proof is complete. \square

PROOF OF PROPOSITION 5 (FROM PAGE 14) CONTINUED. Let $Z = \{x_0, x_1, \dots\}$ be any infinite subset of \mathbb{U}_+ such that $x_0 = 0 \times_w 0$ for some $w \in \mathbb{W}$ and $x_n < x_{n+1}$ for each n in \mathbb{Z}_+ . Suppose that, for some $w \in \mathbb{W}_+$, $x_{n+1} \in L_w$, then, since $x_n < x_{n+1}$, there exists $x'_n \in L_w$ such that $x'_n \sim x_n$, so that the path ϕ_n from x'_n to x_{n+1} belongs to $\Phi^<$. Then by lemma 4, there exists a synchronising concatenation satisfying $f(\nu) = \phi_n((\nu - \mu_n) / (\mu_{n+1} - \mu_n))$ for each n . Note that f is a concatenation, but not a path since $f : [0, 1) \rightarrow X$ whenever Z is infinite. The properties of I and fact that $\mu_n < \mu_{n+1}$ for each n implies that such a concatenation is only possible when Z is countable. Let \mathcal{Z} denote the collection of all such countable and increasing sequences in \mathbb{U}_+ and let f_Z denote the synchronising concatenation corresponding to each $Z \in \mathcal{Z}$.

The proof of ?? shows that \mathbb{U}_+ is order isomorphic to the long line. As such, every countable set has an upper bound in \mathbb{U}_+ (see for instance Munkres [21, Theorem 10.3]). Thus for each $Z \in \mathcal{Z}$, there exists $x \in X$ such that $z \lesssim x$ for every $z \in Z$. W.l.o.g., suppose that $x \in L_w$. Then, since L_w is a linear

continuum and the image of f_Z is a nonempty subset, $\sup f_Z$ is a well-defined element of L_w . For any pair $Z, Z' \in \mathcal{Z}$, suppose $\sup f_Z \lesssim \sup f_{Z'}$. The fact that the concatenations are synchronising ensures that, for every $x_n \in Z$, there exists $\mu, \mu' \in I$ such that $\phi_n(\lambda) \sim f_{Z'}((1-\lambda)\mu + \lambda\mu')$ for every $\lambda \in I$. When this holds, we say that the set $Z'' = Z \cup Z'$ is *reducible* to Z' .

For any subcollection $\mathcal{Z}' \subset \mathcal{Z}$, we say that \mathcal{Z}' is reducible to $Z' \in \mathcal{Z}$ if and only if every $Z \in \mathcal{Z}'$ is reducible to Z' . Let V be a well ordering of the elements of \mathcal{Z} . We proceed by induction. Since each pair Z_v and Z_{v+1} is reducible to the set with the greatest supremum, the initial step and the inductive step for successor ordinals is complete. If v is a limit ordinal, then the induction hypothesis ensures that, for each $v' < v$, the subcollection $\mathcal{Z}_{v'} = \bigcup \{Z_{v''} : v'' \leq v'\}$ is reducible to one of its component sets, say Z_u . Clearly $Z_u \cup Z_v$ is reducible to Z_u or Z_v depending on which has the greatest supremum. This shows that \mathcal{Z} is reducible to some countable set Z . That is for every $\phi \in \Phi$, there exists $\mu, \mu' \in I$ such that $\phi(\lambda) \sim f_Z((1-\lambda)\mu + \lambda\mu')$ for every $\lambda \in I$.

Since Z is countable, $\sup Z \in \mathbb{U}_+$ and this ensures the existence of $w \in \mathbb{W}_+$ and $y \in L_w$ such that $\sup f_Z < y$. But then there is no path $\phi_{xy} \in \Phi$ such that $x = 0 \times_w 0$, for otherwise $f_Z(\nu) \sim y$ for some $\nu \in I$. This implies a contradiction of \mathcal{O} or the fact that L_w is a mixture set. \square

PROOF OF LEMMA 5. Fix $x < y$.

CASE 1 (*the implication of \mathcal{A} holds with indifference*). By lemma 2, either $\inf \phi_0 = \phi_0(0)$ or $\sup \phi_0 = \phi_0(0)$. If it is the latter, appeal to condition \mathcal{P}_2 to obtain ϕ'_0 such that $\phi'_0(0) \sim x$. If it is the former, simply relabel ϕ_0 as ϕ'_0 and proceed to ϕ_1 . Similarly, if $\sup \phi_0 = \phi_0(0)$, then once again appeal to condition \mathcal{P}_2 to obtain ϕ'_1 such that $\phi'_1(0) = \inf \phi'_1$. Since m is finite, we exhaust the sequence by this recursive procedure.

Let f be a concatenation of ϕ'_0, \dots, ϕ'_m . Then lemma 4 ensures the existence of a synchronising concatenation g from x to y .

CASE 2 (*not case 1*). By the argument of case 1, w.l.o.g. suppose that $\inf \phi_n = \phi_n(0)$ for each n . Suppose that $\phi_0(0) < x$. W.l.o.g., we may assume that m is the minimal length of a sequence satisfying \mathcal{A} . Then $x \sim \phi_0(\mu)$ for some $\mu < 1$, for otherwise, we can exclude ϕ_0 and obtain a shorter sequence. By \mathcal{O} and lemma 2, μ is unique. Using the partial mixture set conditions (see discussion before definition 2), take $\phi'_0 \in \Phi$ such that $\phi'_0(\lambda) = \phi_0((1-\lambda)\mu + \lambda)$ for each $\lambda \in I$. If $\phi_1(0) < \phi'_0(1)$, then, by the same method that was used to obtain ϕ'_0 , we can find ϕ'_1 such that $\phi'_1(0) \sim \phi'_0(1)$, and $\phi'_0(1) < \phi'_1(\lambda)$ if $0 < \lambda$. Once again, since m is finite, in this way we can define a suitable sequence ϕ'_0, \dots, ϕ'_m such that the conditions for case 1 of the present proof to apply. \square

PROOF OF LEMMA 6 (FROM PAGE 16). Fix $x < z < y$. By lemma 5, Φ generates a synchronising concatenation from x to y . Let ϕ_0, \dots, ϕ_m and $0 = \mu_0 < \dots < \mu_{m+1} = 1$ be the sequences that define f . Take g to be any other synchronising concatenation from x to y that is characterised by $\gamma_0, \dots, \gamma_k$ and $0 = \nu_0 < \dots < \nu_{k+1} = 1$.

CASE 3 (for every $1 \leq n \leq m$ and every $1 \leq j \leq k$, $f(\mu_n) \not\sim g(\nu_j)$). Note that in this case, we do not need to appeal to \mathcal{M} : we can use the paths that make f and g . W.l.o.g., suppose that $f(\mu_1) < g(\nu_1) < f(\mu_2)$. Then, by lemma 2, $\gamma_0(1) \sim \phi_1(\mu')$ for some unique $0 < \mu'$. (The fact that $\phi_0(1) \sim \phi_1(0) < \gamma_0(1)$ implies that $0 < \mu'$.) Then, for some ν' , $f(\nu') = \phi_1(\mu')$ and, since f is synchronising, $\gamma_0(\xi) \sim f(\xi\nu')$ for every $\xi \in I$. That is, $f(\xi\nu') \sim g(\xi\nu_1)$ for every $\xi \in I$. If $g(\nu_1) \sim y$, then $\nu_1 = 1$ and, since $f(\xi) \sim y$ if and only if $\xi = 1$, we have $\nu' = \nu_1$ as required.

If $g(\nu_2) < f(\mu_2)$, then repeat the argument of the last paragraph: using γ_1 instead of γ_0 to show that $f(\xi\nu'') \sim g(\xi\nu_2)$ for every $\xi \in I$ and some $0 < \nu'' \leq 1$. If on the other hand, suppose that $f(\mu_2) < g(\nu_2) < f(\mu_3)$, then the argument of the previous paragraph applies, with the exception that now we use the fact ϕ_1 is synchronised with g . In all cases, the fact that m and k are finite, $\mu_0 = \nu_0 = 0$ and $\mu_{m+1} = \nu_{k+1} = 1$ means that we eventually establish that $f(\xi) \sim g(\xi)$ for every $\xi \in I$ as required.

CASE 4 (for some $1 \leq n \leq m$ and $1 \leq j \leq k$, $f(\mu_n) \sim g(\nu_j)$). Suppose that $x < f(\mu_1) \sim g(\nu_1) < y$. Together, \mathcal{M} , lemma 2 and condition \mathcal{P}_2 ensure the existence of ϕ such that $\psi(0) < f(\mu_1) < \psi(1)$ and $\psi(\lambda') \sim f(\mu_1)$ for some unique $\lambda' \in I$. W.l.o.g. (by passing to a subpath where necessary) suppose that $x \lesssim \psi(0)$ and $\psi(1) < g(\nu_2) \lesssim f(\mu_2)$.

Since $g(\nu_2) \lesssim f(\mu_2)$, there exists ν'' such that $f(\nu'') \sim g(\nu_2)$. Let $\phi^* \in \Phi$ be the subpath of ϕ_1 such that $\phi^*(\xi) = f((1-\xi)\mu_1 + \xi\nu'')$ for every $\xi \in I$. Then, lemma 3, $\phi^*(\xi) \sim \gamma_1(\xi)$ for every $\xi \in I$. Similarly, since $f(\mu_1) \sim g(\nu_1)$, $\phi_0(\xi) \sim \gamma_0(\xi)$ for every $\xi \in I$.

Then there exist unique $\kappa_* < 1$ and $0 < \kappa^*$ such that $\phi_0(\kappa_*) \sim \psi(0)$ and $\phi^*(\kappa^*) \sim \psi(1)$. Then, by the argument of step 5 in the proof of lemma 4 and the fact that f is synchronising to begin with,

$$\mu_1 = \frac{\lambda' \kappa^* \nu''}{(1-\lambda')(1-\kappa_*) + \lambda' \kappa^*} \quad (7)$$

But, since we also have $\gamma_0(\kappa_*) \sim \psi(0)$ and $\gamma_1(\kappa^*) \sim \psi(1)$, we may derive a similar equation to (7), with ν_1 and ν_2 substituting for μ_1 and ν'' respectively. If $\nu_2 = 1$, then, as in case 3, $\nu' = 1$ and the proof is complete. Otherwise, we may minor variations of the present argument or that of case 3, to extend and show that ν'' depends in a similar way on μ_2 and ν_2 on some μ'' such that $g(\mu'') \sim f(\mu_2)$. Since m and k are finite, we repeat until we reach $\mu_{m+1} = \nu_{k+1} = 1$, at which point $f(\xi) \sim g(\xi)$ for every $\xi \in I$ and the proof is complete. \square

PROOF OF THEOREM 1 (FROM PAGE 17). If $\leq = \emptyset$, then by \mathcal{O} , $x \sim y$ for every $x, y \in X$. In this case, every utility representation is both linear and cardinal. Conversely, both \mathcal{A} and \mathcal{M} hold vacuously when $\leq = \emptyset$. This ensures that the axioms are necessary and sufficient in this case.

Henceforth, suppose that $\leq \neq \emptyset$.

STEP 7 (*Necessity of the axioms*). The proof that \mathcal{O} , \mathcal{C} and \mathcal{S} are necessary for a cardinal, linear representation is clear and therefore omitted.

We now use counterexamples to prove that \mathcal{A} and \mathcal{M} are necessary for such a representation. First note that example 5 is a representative counterexample for the case where \mathcal{A} holds and \mathcal{M} does not. For the case where there exists a linear utility representation and \mathcal{M} holds but \mathcal{A} does not, see example 7 in the proof of corollary 1 below. A similar accounts for the case where both \mathcal{M} and \mathcal{A} fail to hold.

In each of these cases, any linear utility representation fails to be cardinal. The only remaining case is where \mathcal{A} and \mathcal{M} both hold. In this case, a representation of the required form exists by the following argument.

STEP 8 (*Sufficiency of the axioms*). Consider the quotient set $X_{/\sim}$. Each element of $X_{/\sim}$ consists of an indifference class generated by preferences on X . $X_{/\sim}$ is well-defined because \mathcal{O} ensures that the indifference classes partition X . Let $p : X \rightarrow X_{/\sim}$ be the natural projection $x \mapsto \{y : y \sim x\}$. Let $F_{/\sim}$ be the set of paths f' in $X_{/\sim}$ such that $f' = p \circ f$ for some synchronising concatenation f that is generated by Φ in X . The arguments of the next two paragraphs demonstrate that $(X_{/\sim}, F_{/\sim})$ is a mixture set whenever $< \neq \emptyset$. To be clear, all remaining arguments in this proof (and indeed the paper) proceed in X , not $X_{/\sim}$.

If $x < y$, then lemmas 5 and 6 guarantee the existence and (upto indifference) uniqueness of a synchronising concatenation f from x to y . A repeated application of condition \mathcal{P}_2 yields the concatenation g from y back to x exists and satisfies $g(\lambda) \sim f(1 - \lambda)$ for every $\lambda \in I$. For condition \mathcal{P}_3 , suppose that $z \sim f(\mu)$ for some $z \in X$ and $\mu \in I$. Since f is generated by Φ , take $g(\nu) = f(\nu\mu)$ for each $\nu \in I$. Then $g(\nu) = \phi_n((\nu\mu - \mu_n)/(\mu_{n+1} - \mu_n))$ for each n such that $\mu_n < \mu$ and each $\nu \in I$. If $\mu > 0$, then simple division of the numerator and denominator in each ϕ_n shows that g is a well-defined concatenation. If $\mu = 0$, then $g(\nu) = \phi_0(0)$ for every $\nu \in I$ and by condition \mathcal{P}_3 , there exists $\phi \in \Phi$ such that $g = \phi$.

If $x \sim y$, then first suppose $x \lesssim x'$ for every $x' \in X$. Since $< \neq \emptyset$, there exists y' such that $x < y'$. In this case, lemma 5 ensures the existence of a synchronising concatenation f from x to y' . Then $f(\mu) \sim x$ for $\mu = 0$ and the preceding paragraph completes the proof. The case where $x' \lesssim x$ for every $x' \in X$ is similar, so we proceed to the case where $x' < x < y'$ for some $x', y' \in X$. In this case, \mathcal{M} and lemma 2 ensure that $x \sim z' = \phi(\mu)$ for some $\phi \in \Phi$ and $0 < \mu < 1$. By proposition 1, $\phi(I)$ is a mixture set and $\phi_{z'z'} \in \Phi$. Finally, since $\phi_{z'z'}(\lambda) \sim x$ for every $\lambda \in I$, $\phi_{z'z'}$ is synchronising concatenation from x to y . Finally, since condition \mathcal{P}_1 holds in every case, we have shown that, upto indifference, X a mixture set.

For the axioms, recall from our discussion following definition 2, that, by virtue of \mathcal{C} , concatenations inherit continuity from the continuity of members of Φ . Moreover, lemma 6 is clearly sufficient for the property $f^{(1/2)} \sim g^{(1/2)}$ for any pair of synchronising concatenations such that $f(0) \sim g(0) < f(1) \sim g(1)$. That is to say, synchronising concatenations satisfy an independence axiom akin

to \mathcal{S} . We may therefore apply the main theorem of HM to obtain a cardinal and linear utility representation of preferences. In summary, we have completed the proof that the axioms are sufficient for the required utility representation. \square

PROOF OF COROLLARY 1 (FROM PAGE 17). The fact that the axioms are sufficient for cardinality when preferences have a linear representation follows from theorem 1: \mathcal{O} , \mathcal{C} and \mathcal{S} hold whenever preferences have a linear representation.

For the necessity of the axioms, note that, since example 4 fails to give rise to a utility representation, it is of no use for the present proof. Example 5 accounts for the case where \mathcal{A} holds, \mathcal{M} does not. The following example accounts for the case where \mathcal{M} holds, but \mathcal{A} does not.

Example 7. Let X be the disjoint union of $\{\phi_n(I), \gamma_n(I) \in \Phi^< : n \in \mathbb{Z}_+\}$. Suppose that for every $n \in \mathbb{Z}_+$: $\phi_n(0) \sim \phi_{n+1}(0)$ and $\phi_n(1) < \phi_{n+1}(1)$. Thus every ϕ_n has the same infimum and the supremum is strictly increasing in n . Moreover, suppose that, for every $m, n \in \mathbb{Z}_+$, $\phi_m(1) < \gamma_n(0)$, so that any element of any $\gamma_n(I)$ is strictly better than any element of any $\phi_m(I)$. Finally, in contrast, let every γ_n have the same supremum and let the infimum be decreasing in n . That is, $\gamma_n(1) \sim \gamma_{n+1}(1)$ and $\gamma_n(0) < \gamma_{n+1}(0)$ for every $n \in \mathbb{Z}_+$.

In this case, the fact that \mathbb{Z}_+ is unbounded ensures that \mathcal{M} holds. Take any $x \in \phi_m(I)$ and $y \in \gamma_n(I)$. If $z \in \phi_{m+1}(I)$, then ϕ_{m+2} is a path such that $\phi_{m+2}(0) < z < \phi_{m+2}(1)$. However, this example clearly violates \mathcal{A} , for Φ does not generate a concatenation from x to y .

Suppose preference have a linear representation U . We now show that U is not cardinal. Let $[r, s]$ and $[t, u]$ be the respective image sets of $\{U \circ \phi_m\}$ and $\{U \circ \gamma_n\}$. Since U is a utility function, $s \leq t$. Define $V : X \rightarrow \mathbb{R}$ such that $V(z) = U(z)$ for every $z \in \bigcup_m \phi_m(I)$. For every $z \in \bigcup_n \gamma_n(I)$, let $V(z) = \kappa + \theta U(z)$, where $\theta > 1$ and $\kappa + \theta t = t$. Such a choice is always possible since $t - \theta t < 0$ and we have complete freedom choose $\kappa < 0$.

It is straightforward to confirm that V is a linear utility representation. But since it not a positive affine transformation of U , the latter is not cardinal.

This accounts for all but the case where both \mathcal{A} and \mathcal{M} fail to hold. The latter is accounted for by a simplification of example 7 to the case where X is the union of just two paths ϕ and γ . Since $\phi(\lambda) < \gamma(\mu)$ for every $\lambda, \mu \in I$, \mathcal{A} fails to hold, just as in example 7. The fact that \mathcal{M} fails to hold follows by considering that $\phi(0) < \phi(1) < \gamma(1)$: there is no path in Φ that satisfies \mathcal{M} for the point $\phi(1)$.

Since the linear representations in all these examples fails to be cardinal, we have the implication: not $[\mathcal{A}$ and $\mathcal{M}]$ implies not cardinal as required. \square

PROOF OF PROPOSITION 6 (FROM PAGE 19). Let $\xi = (1 - \lambda)\mu + \lambda\nu$. We only prove the case for $\phi = \phi_{xy} \in \Phi_Y$. The difference between this and the more general case is only a matter of notation. Moreover, since $(1 - \lambda)[a + \mu(b - a)] + \lambda[(a + \nu(b - a))]$ is equal to $a + \xi(b - a)$, which is convex combination, the remainder of the argument shows that the nonconvexity arises from the product

of $(1 - t)$ and the integral term in (4). Since $1 - \xi = (1 - \lambda)(1 - \mu) + \lambda(1 - \nu)$, we have

$$\begin{aligned} (1 - \lambda)(1 - \mu) \int_0^\mu \frac{1}{1 - r} dW_r + \lambda(1 - \nu) \int_0^\nu \frac{1}{1 - r} dW_r \\ = (1 - \xi) \int_0^\xi \frac{1}{1 - r} dW_r + \text{error}, \end{aligned}$$

where

$$\text{error} = - (1 - \lambda)(1 - \mu) \int_\mu^\xi \frac{1}{1 - r} dW_r + \lambda(1 - \nu) \int_\xi^\nu \frac{1}{1 - r} dW_r.$$

The fact that the latter two integrals are defined on nondegenerate and nonoverlapping intervals ensures that the two stochastic integrals are independent random variables that are nonzero with probability one. With the same probability, the error is therefore nonzero and ϕ is nonconvex. \square

References

- [1] N.I. Akhieser. *The Classical Moment Problem*. Trans. by N. Kemmer. Edinburgh and London: Oliver and Boyd Ltd., 1965.
- [2] Francis J Anscombe and Robert J Aumann. “A definition of subjective probability”. In: *Annals of mathematical statistics* (1963), pp. 199–205.
- [3] Ole E Barndorff-Nielsen, Neil Shephard, et al. “Basics of Lvy processes”. In: *Draft Chapter from a book by the authors on Lvy Driven Volatility Models* (2012).
- [4] Daniel Bernoulli. “Specimen Theoriae novae de Mensura Sortis”. In: *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5 (1738).
- [5] Klaus Bichteler. *Stochastic integration with jumps*. Vol. 89. Cambridge University Press, 2002.
- [6] Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.
- [7] PC Fishburn. “Utility functions on ordered convex sets”. In: *Journal of Mathematical Economics* 12.3 (1983), pp. 221–232.
- [8] Peter C. Fishburn. *The foundations of expected utility*. Theory & Decision Library V.31. Springer Science and Business Media, BV, 1982.
- [9] Itzhak Gilboa and David Schmeidler. *A theory of case-based decisions*. Cambridge University Press, 2001. ISBN: 9780521003117.
- [10] Itzhak Gilboa and David Schmeidler. “Inductive Inference: An Axiomatic Approach”. In: *Econometrica* 71.1 (2003), pp. 1–26.
- [11] Simon Grant et al. “Generalized utilitarianism and Harsanyi’s impartial observer theorem”. In: *Econometrica* 78.6 (2010), pp. 1939–1971.

- [12] Aviad Heifetz, Martin Meier, and Burkhard C Schipper. “Interactive unawareness”. In: *Journal of Economic Theory* 130.1 (2006), pp. 78–94.
- [13] Israel N. Herstein and John Milnor. “An Axiomatic Approach to Measurable Utility”. English. In: *Econometrica* 21.2 (1953), pp. 291–297. ISSN: 00129682.
- [14] Seok Young Hong, Oliver Linton, and Hui Jun Zhang. “An investigation into multivariate variance ratio statistics and their application to stock market predictability”. In: (2015).
- [15] Tomoki Inoue. “A utility representation theorem with weaker continuity condition”. In: *Journal of Mathematical Economics* 46.1 (2010), pp. 122–127. ISSN: 0304-4068. DOI: <http://dx.doi.org/10.1016/j.jmateco.2009.09.002>.
- [16] Daniel Kahneman and Amos Tversky. “Prospect theory: and analysis of decision under risk”. In: *Econometrica* 47(2) (1979), pp. 263–291.
- [17] E. Karni. “A theory of medical decision making under uncertainty”. In: *Journal of Risk and Uncertainty* 39.1 (2009), pp. 1–16.
- [18] E. Karni and Z. Safra. “An extension of a theorem of von Neumann and Morgenstern with an application to social choice theory”. In: *Journal of Mathematical Economics* 34.3 (2000), pp. 315–327.
- [19] D. H. Krantz et al. *Foundations of measurement, Vol.I: Additive and polynomial representations*. New York: Academic Press, 1971.
- [20] Paulo Klinger Monteiro. “Some results on the existence of utility functions on path connected spaces”. In: *Journal of Mathematical Economics* 16.2 (1987), pp. 147–156.
- [21] James R. Munkres. *Topology*. 2nd. Prentice Hall, 2000. ISBN: 9780131816299.
- [22] John Quiggin. “A theory of anticipated utility”. In: *Journal of Economic Behavior & Organization* 3.4 (1982), pp. 323–343.
- [23] Michael Richter and Ariel Rubinstein. “Back to Fundamentals: Equilibrium in Abstract Economies”. In: *American Economic Review* 105.8 (2015), pp. 2570–94. DOI: 10.1257/aer.20140270.
- [24] Michael Rockinger and Eric Jondeau. “Entropy densities with an application to autoregressive conditional skewness and kurtosis”. In: *Journal of Econometrics* 106.1 (2002), pp. 119–142.
- [25] Leonard J. Savage. *The foundations of statistics*. Dover Publications, 1972. ISBN: 9780486623498.
- [26] David Schmeidler. “Subjective probability and expected utility without additivity”. In: *Econometrica* 57.3 (1989), pp. 571–587.
- [27] Jack Stecher, Timothy Shields, and John Dickhaut. “Generating ambiguity in the laboratory”. In: *Management Science* 57.4 (2011), pp. 705–712.

- [28] Jordan Stoyanov. “Krein condition in probabilistic moment problems”. In: *Bernoulli* 6.5 (Oct. 2000), pp. 939–949.
- [29] John von Neumann and Oscar Morgenstern. *Theory of games and economic behavior*. Sixtieth anniversary. Princeton and Oxford: Princeton University Press, 1944.