Context and Decision : Utility on a Union of Mixture Spaces

Patrick O'Callaghan

No 973

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF WARWICK

Context and Decision: Utility on a Union of Mixture Spaces

Patrick O'Callaghan*

University of Warwick

Abstract

Suppose a decision-maker is willing to make statements of the form: *"I prefer to choose alternative a when in context p, than to choose alternative b when in context q"*. Contexts p and q may refer to given probability distributions over a set of states, and b and c to alternatives such as: "turn left" or "turn right" at a junction. In such decision problems, the set of alternatives is discrete and there is a continuum of possible contexts. I assume there is a is a mixture operation on the space of contexts (eg. convex combinations of lotteries), and propose a model that defines preferences over a collection of mixture spaces indexed by a discrete set. The model yields a spectrum of possibilities: some decision-makers are well represented by a standard von Neumann–Morgenstern type of utility function; whilst for others, utility across some or all the mixture spaces is only ordinally comparable. An application to the decision problem of Karni and Safra (2000) leads to a generalization, and shows that state-dependence and comparability are distinct concepts. A final application provides a novel way of modeling incomplete preferences and explaining the Allais paradox.

*I thank Simon French, Peter Hammond, Andrea Isoni, Saul Jacka, Aron Toth and Horst Zank for their helpful suggestions and detailed feedback.

1 Introduction

When a decision-maker sees no benefit to fooling her opponent in a game, she presumably sees no reason to define her preferences over her own mixed strategies. Yet if the decision-maker is to be modeled using a von Neumann–Morgenstern (1944) [vNM] expected utility function, or one of the generalizations we discuss in more detail below, this is what she is required to do.

There are models, such as that of Gilboa and Schmeidler (2003) [GS] and O'Callaghan (2011) which address this concern by reducing the domain of preferences to be the set of alternatives (pure actions) that are available to the decision-maker. Then, preferences are indexed by the set of possible beliefs the decision-maker has regarding her opponent's move. Whilst this approach certainly addresses the issue concerning preferences over own mixed strategies, it also introduces new difficulties.

In particular, if preferences are to be represented by an expected utility function, they must satisfy a diversity condition. In [GS], this condition says that for every set of four alternatives available to the decision-maker, and each possible strict ordering of the four alternatives (there are $4! = 24$ of these), there exists a context such that preferences agree with that ranking. These diversity conditions are not only unnecessary for an expected utility representation, they are also strong enough to exclude the majority of possible decision-makers. Indeed, as is discussed in Ashkenazi and Lehrer (2001) and O'Callaghan (2011), for decision problems where there are only two states of nature and more than three alternatives, the diversity condition of [GS] is so strong as to exclude all possible preferences. Moreover, there appears to be no intuitive conditions on preferences that resolve this problem.

My purpose therefore is to identify a minimal domain that allows us to obtain an expected utility function without recourse to these extra conditions on preferences. I show that it suffices to define preferences over alternative-context pairs. The results are obtained in the more general setting of mixture preserving functions on mixture spaces in the spirit of classic paper of Herstein and Milnor (1953) [HM].

1.1 Two motivating examples

The first class of problems that is well suited to the model I present in this paper is where the decision-maker's "opponent" is nature. Moreover,

(i) she knows she *will* face a choice in the presence of uncertainty about the future state of nature;

(ii) she knows, that when the situation arises, she will have knowledge of the context she is in (for instance, she will know the likelihood of any given state of nature);

(iii) given a context, she knows her preference for one alternative over another;

(iv) given a particular course of action, she knows which context she would rather be in;

(v) building upon (iii) and (iv), she is willing to go further and make statements of the form "I prefer to choose alternative a when in context p, than to choose alternative b when in context q".[1]

Although (ii) of this list is arguably a strong assumption for the class of examples of this paper, I will simply take it as given and accept it as a topic for future research without further discussion.

[1]In this paper we take this statement to be equivalent to "I prefer to be in context p when choosing alternative a, than to be in context q when choosing alternative b".

Example. *Consider a planner who is devising a complete, contingent plan of how to respond to the future threat of a flood. There are two states (flood and no-flood) and there are two alternatives: do nothing or evacuate. Crucially, the contingencies are defined to be the set of possible probability distributions over states, not the states that may subsequently obtain. The idea being that, when the time comes, the action should be carried out without question.*

In this example, the planner should have no reason to consider mixtures over her set of alternatives: there is no obvious benefit to doing so. The situation is different from a game against a strategic opponent, such as rock-paper-scissors, where being predictable carries a cost. As Rubinstein (2000) and Gilboa and Schmeidler (2003) highlight, data on players' preferences over mixed strategies may be unreliable as it is not clear whether observable, pure actions of the player are part of a grand mixed strategy. If this hypothesis is true in games against other players, it is even more true in games that are "against nature".

On the other hand, there is good reason for the planner to consider each possible contingency (context) and, given this, say whether she would evacuate or not. It is not too much more to ask her to state her preferences over contingencies given a choice of alternative. This suggests that (iii) and (iv) are reasonable assumptions in the flood example.

This leaves (v). Now the planner needs to be willing to make statements of the form "I prefer to announce evacuate when the probability of flooding is $\frac{1}{2}$ than to make no announcement when the probability of flooding is $\frac{1}{4}$". Whilst there is no doubt that this is more demanding, it is worthwhile putting it into perspective by considering the following class of statements "I prefer to announce evacuate with probability $\frac{1}{2}$ when the probability of flooding is $\frac{1}{2}$ than to evacuate with probability $\frac{1}{4}$ when the probability of flooding is $\frac{1}{4}$". The latter type of statement is necessary if we wish to apply the benchmark [vNM] model of expected utility.

In fact the [vNM] model requires quite a bit more than this. It requires that the decision-maker is willing to make preference statements about any pair of probability distributions over the set of four outcomes defined by taking the product of the set of alternatives with the product of the set of states: (d,n)≡(do nothing, no flood), (d,f)≡(do nothing, flood), (e,n)≡(evacuate, no flood) and (e,f)≡(evacuate, flood).

This space of lotteries contains probability distributions where the (joint) probability of the outcome (evacuate, flood) is not equal to the product of the marginal probabilities. Let δ_x denote the probability measure assigning probability one to outcome x and consider the following lottery:

$$\tfrac{1}{2}\,\delta_{(d,n)} + \tfrac{1}{4}\,\delta_{(d,f)} + \tfrac{1}{4}\,\delta_{(e,n)}.$$

For outcome (d,n) to occur with probability $\frac{1}{2}$, it seems reasonable to assume, that in the absence of mischievous deities, there is a positive probability that the central planner chooses "do nothing" *and* a positive probability that "flood" occurs. Similarly, for (e,n) to occur with probability $\frac{1}{4}$, the planner ought to be choosing "evacuate" with positive probability. However, by taking the product of the marginal probabilities of evacuate and flood, outcome (e,f) occurs with positive probability. Now this can't be because the sum of $\frac{1}{2}$, $\frac{1}{4}$ and $\frac{1}{4}$ is 1.

Even so, suppose the planner was willing to entertain the possibility of her strategies somehow being correlated with those of nature's; would she be willing to say whether she preferred the above lottery to

$$\tfrac{7}{16}\,\delta_{(d,n)} + \tfrac{3}{16}\,\delta_{(d,f)} + \tfrac{1}{16}\,\delta_{(e,n)} + \tfrac{5}{16}\,\delta_{(e,f)}\,?$$

It seems that a simple but challenging decision problem has been turned into a complicated problem that is even more challenging. The difficulty is that to define preferences over the space of lotteries on $A \times S$, where A is the set of alternatives and S is the set of states, is to define preferences over the 3-dimensional unit simplex $\Delta(A \times S)$ in $\mathcal{R}^4$. This adds considerable complexity to the decision problem.

One solution to this latter difficulty is to use the multilinear expected utility model of Fishburn (1980). This model allows us to define preferences on the product of unit-simplices $\Delta(A) \times \Delta(S) = [0, 1] \times [0, 1]$. It therefore allows the decision-maker to avoid defining her preferences on the strange lotteries described above. However, it still requires that the decision-maker define her preferences over her own mixed strategies. As a result, the conditions on preferences that Fishburn imposes are more numerous and more complicated than the model of this paper.

Instead, I propose to define a simple model of preferences over the product $A \times \Delta(S)$. In the above example this amounts to two copies of the unit interval. Aside from the intuitive appeal in decision problems like the flood example above, the advantage to defining preferences on a smaller space, in terms of the complexity of eliciting a utility function may be significant in real-life complex decision problems such as those studied by computer scientists like Braziunas and Boutilier (2010).

In order to show that the present model is not restricted to choice under risk/ uncertainty, I now present a second example where the space of contexts need not be the set of probability distributions over states. It serves to motivate the concept of a mixture space which is defined formally below. For now it suffices to think of it as a suitable generalization of a convex space in which the operation of taking mixtures is defined.

Example 1.1. *Consider a customer in a restaurant choosing a meal with the assistance of a waiter. The menu defines a finite set of alternatives, and there are a continuum of possibilities for each dish: how well the flavors of the main ingredients combine, the degree to which the food will be cooked, the quantity of salt it will contain, etc. Suppose that the waiter is able to describe the context space as precisely as the customer wishes.*

Given the choice of a particular item on the menu, the customer may well be willing to state her preferences over the possible contexts, thus (iv) may be reasonable enough here. She may also be willing to make statements of the form: "I prefer the carrot soup with a table spoon and a half of cream, to the tomato salad with one quarter of an onion in it."

On the other hand, it is unclear that the decision-maker would be comfortable defining her preferences over mixtures of "soup" and "salad", even if by mixture we did mean probability mixtures. Presumably the only situation where she might even consider such alternatives is when she is unsure of what to do and flips a coin to break the tie. So Fishburn's multilinear model, which also holds for mixture spaces, is arguably inappropriate here. As above, to define preferences on a single mixture space, we must go even further: either assume the decision-maker considers lotteries over the product of the alternative space with the space of contexts; or allow for a continuum of portion sizes for all items, and expand the menu to include mixtures of all possible ingredients, cooking styles etc.

At this point it is natural to question the need for introducing any of the above strange considerations to the decision problem. The answer is that we are often interested in representing preferences with a utility function for which numerical values have meaning in same sense that temperature values do. That is given a choice of scale and origin, say, Celsius, we can speak of temperature as a number. Moreover, if someone else uses another scale and origin, such as Fahrenheit, their statements are meaningful to us because we know the affine transformation that converts one to the other.[2]

In game theory for instance the starting point is to specify a given player's "pay-offs" as numbers. [vNM] showed that this is a reasonable starting point provided preferences are defined on the set of probability distributions (or lotteries) over the outcomes of the game, and provided certain conditions on preferences apply. In this case, the player's payoffs are unique up to a common scale and origin (a positive affine transformation).

For the customer in the restaurant example, where lotteries may play no part, if we seek such a utility representation, we must appeal to the generalization of [vNM] by [HM]. They define preferences over a mixture space and impose very similar conditions on preferences to those of [vNM]. Their representation is also unique up to a positive affine transformation, and so utility units in their model are meaningful.

1.2 Outline

The main purpose of the present paper is to pin down the minimal conditions on preferences that extend [HM] (and hence [vNM]) to the setting where there is more than one mixture space and provide a precise characterization of the utility function.

The model of Karni and Safra (2000) (henceforth [KS]) comes closest to the present class of problems. There are two reasons why I choose to build a new model. The first is that they introduce additional structure on the

[2]The analogy is not so good if we consider Kelvin as this has a fixed origin.

space on which preferences are defined. This structure is not needed to make the extension of [vNM] and [HM] to the general setting where mixtures are not everywhere defined. I apply the model of the present paper to the space they define in section (4). Second, they impose conditions on preferences that are have no counterpart in the original models of [vNM] and [HM], and are specific to the space on which they define preferences. This means that the representation in this paper is therefore both simpler and more general than [KS] and the related paper of Karni (2009). Correspondingly, it is less straightforward to obtain, and the length of the derivation in section (3) testifies to this.[3]

What distinguishes the utility function of this paper from that of [vNM], [HM], Fishburn (1980), [KS], but also the "state-dependent" utility models, is the extent to which utility is numerically comparable over the domain of preferences. In "state-dependent" utility models (eg. Dréze (1961)), comparing utility numbers across states is meaningless, for there is an independent scale and origin to utility for each state. It is the polar opposite of [vNM] and the other models we have discussed.

Here however, there is a spectrum of possibilities. For some decision-makers, utility is numerically comparable across mixture spaces in the sense that the utility function is, to use the terminology of Karni (2009), unique up to a positive affine transformation that applies uniformly across the domain (i.e. utility is cardinally measurable and fully comparable across the domain (CFC)). For others, the utility function will only be CFC within certain sub-

[3]Moreover, the relevant result in that paper contains two mistakes. The first is their claim that the utility function is unique up to a common scale but not a common origin. As Karni (2009) points out and amends, it is in fact unique up to a positive affine transformation. The second mistake is also present in Karni (2009) and a counterexample to their claim is presented in section (4) below.

sets of its domain. If such a subset is maximal, in the sense that it is the largest subset for which numerical utility comparisons are possible, we will call it a quasi-component of preferences.[4] Across quasi-components, only ordinal utility comparisons are meaningful.

Such concepts may at first sight seem irrelevant to decision-making. However the essential idea is that when, regardless of context, one alternative is obviously better than another, it may be that the decision-maker has need for the high resolution measurement scale that [vNM] require. For other pairs of alternatives, the [vNM] model may be natural. The issue bears resemblance to the way that one does not typically need scales to decide whether a toddler is lighter than an adult, but periodically, we do need precise scales to measure whether our own weight has increased or decreased.

I argue that this property of preferences may well be justified in applications, and, unless there is good reason, it should not be assumed away. One such application is provided in section (5), the final section of this paper, where I will argue that, when preferences are incompletely defined on the space of lotteries over monetary outcomes, preferences that accord to the Allais paradox are straightforwardly captured by a mixture-preserving utility function that is not defined over the entire simplex of lotteries. This application also highlights the fact that in general the set A need not index alternatives, it may index the members of a union of subsets of a single simplex.

The next section presents the conditions we impose upon preferences, the mixture space as well as the space of alternatives. Following this, in section

[4]This terminology is related to the concept of a component in topology. These sets resemble components in some ways, but since they may contain limit points of other quasi-components, we have chosen this terminology.

(3) the two main representations are derived, the first generalizes the model of [HM], and the second the model of [vNM].[5]

The paper then concludes with two applications of the model. The first, in section (4) makes precise the differences between our model and that of [KS] and shows how state-dependent preference can hold despite utility being CFC across states. I then conclude in section (5) with the application to the Allais Paradox.

2 Conditions on preferences and the space of alternative-context pairs

To keep the notation as simple as possible, I assume that the space of contexts, $\mathcal{M}$, is the same for all the elements of A, which we will refer to simply as alternatives. Moreover $\mathcal{M}$ is a mixture space, which is defined in the following way:

Definition 2.1 (Mixture set). *A set $\mathcal{M}$ is said to be a* mixture set (or space) *if for any $x, y \in \mathcal{M}$ and any λ, we can associate another element, which we write as either $\lambda x + (1 - \lambda)y$ or $x\,\lambda\,y$, which is again in $\mathcal{M}$, and where*

(1) $x\lambda x = x$

(2) $x\lambda y = y(1 - \lambda)x$

(3) $(x\lambda y)\mu y = x(\lambda\mu)y$.

[5]Although great care has been taken to define every new concept and explain each step in the proofs, the derivation is rather technical. Hence, I suggest that non-specialists first read the following section, which makes precise the limitations of the model, and continue reading near the end of section (3) where the representations are to be found.

A mixture space is more general than a simplex. As Mongin (2001) shows, a mixture space needn't even be isomorphic to convex subset of a vector space. Even so, the following pair of examples shows that it is not general enough to include spaces such as $A \times \mathcal{M}$, where A is a general set of alternatives. The issue is of course that the mixture operation need not be defined on the whole of $A \times \mathcal{M}$.

Example. *Let $A = \{a, b\}$. Let (a, p) be an element of $\{a\} \times \mathcal{M}$ and $y = (b, p)$ an element of $\{b\} \times \mathcal{M}$. Clearly, there is no $\lambda \in [0, 1]$ other than 0 and 1 such that $x\lambda y$ lies in $A \times \mathcal{M}$.*

Example. *If A is the set of rational numbers the set $A \times \mathcal{M}$ is neither a mixture space, nor a product of mixture spaces. For if we are to take the interval $[0, 1]$, then $0 < \frac{1}{\sqrt{2}} < 1$ is not rational. Neither therefore is $(\frac{1}{\sqrt{2}}, p) \in A \times \mathcal{M}$ for any $p \in \mathcal{M}$.*

Note that the set we study can also be written as a union as follows

$$A \times \mathcal{M} \equiv \bigcup_{a \in A} \big(\{a\} \times \mathcal{M}\big).$$

Unlike [HM], I will assume that the mixture space $\mathcal{M}$ is endowed with a topology, and that, under this topology, it is compact. This, together with the conditions I impose on preferences, as I show below, is sufficient for the existence of a greatest lower bound (glb) and least upper bound (lub) of $\mathcal{M}$ under the order that preferences define for each alternative $b \in A$.

This is an important simplification that makes the proofs more straightforward. Further research is required to prove that this condition can be weakened so as to have a representation for a union of general mixture spaces.[6] Nonetheless, it still allows for a wide range of context spaces as the following example highlights.

[6] Although these results are not presented here, my current efforts appear to show that, for a topological mixture space at least, this should be possible.

Example 2.2. *For an arbitrary set S, $[0,1]^S$ is, by the Tychonov theorem, compact in the product topology. Now since $\Delta(S)$, is a closed subset of $[0,1]^S$, by theorem 26.2 of Munkres p181 $\Delta(S)$ is also compact, and as such, a valid space of contexts for the results that follow.*

Unless otherwise stated, the space of alternatives will be assumed to be finite. Where possible my proofs are written so that they would also apply to the countably infinite case. Preliminary research into the latter shows that it is somewhat more complicated, and that to progress we will need to make slight alterations to some of the concepts we introduce. For the case where A is uncountable, it is clear that further conditions on preferences will be needed. The reason being that there can only be countably many disjoint intervals (with nonempty interior) in $\mathcal{R}$, whereas in general preferences may lexicographically order elements of $A \times \mathcal{M}$ in the sense that for each $a, b \in A$, either (a, p) is strictly better than (b, q) for all $p, q \in \mathcal{M}$, or the reverse strict preference for all $p, q \in \mathcal{M}$, so that any representation of such preferences will need to map into the same number of disjoint intervals as the cardinality of A.

For any given $x, y \in A \times \mathcal{M}$, we take the statement "y is weakly preferred to x" to be equivalent to $x \precsim y$.[7] The conditions on the relation $\precsim$ that will be needed are defined as follows.

Definition (Complete pre-order (O)).
For all $x, y, z \in A \times \mathcal{M}$, both the following hold:
*(i) (*Completeness*) $x \precsim y$ or $y \precsim x$, and*
*(ii) (*Transitivity*) if $x \precsim y$ and $y \precsim z$, then $x \precsim z$.*

Definition (Continuity (C'ty)).

[7]This notation is used by Fishburn (1979) and Binmore (2009). I believe it to be easier to read in the present setting due to the fact that our proof often deals with "intervals" defined by preferences.

For all $x, y \in A \times \mathcal{M}$, the following sets are closed:

$$\{y \in A \times \mathcal{M} : y \gtrsim x\}, \quad \textit{and} \quad \{y \in A \times \mathcal{M} : x \gtrsim y\}$$

When the decision-maker makes a preference statement comparing different contexts $p, q \in \mathcal{M}$, *given that she is choosing a particular alternative $b \in A$*, I will use the shorthand $p \lesssim_b q$ and understand it to be the same statement as $(b, p) \lesssim (b, q)$.

Definition ([vNM]. *Independence on each mixture space (I)]*
For any $b \in A$, and any $p, q, r \in \mathcal{M}$ if $p \sim_b q$, then

$$p\tfrac{1}{2}r \sim_b q\tfrac{1}{2}r.$$

Condition (I) is stated in the form that [HM] introduced in their paper. The first part of that paper is dedicated to showing that this implies the more familiar form that [vNM] introduced, where the condition in (I) holds not only for $\lambda = \frac{1}{2}$, but for all $\lambda \in [0, 1]$. The following condition is to my knowledge new, and it is the condition that as I now show generalizes condition (I).

Definition (Congruent betweenness (CB))**.**
For any $b, c \in A$, $p, q, p', q' \in \mathcal{M}$, if both $(b, p) \sim (c, q)$ and $(b, p') \sim (c, q')$ then

$$(b, p\tfrac{1}{2}p') \sim (c, q\tfrac{1}{2}q').$$

Clearly if the cardinality of A is one, then condition (CB) is implied by condition (I) and transitivity; since (I) implies that whenever $p \sim_b q$ and $p' \sim_b q'$ we have

$$p\tfrac{1}{2}p' \sim_b q\tfrac{1}{2}p' \sim_b q\tfrac{1}{2}q',$$

and transitivity ensures that that indecisiveness propagates. In fact, by taking $p = q$ and $b = c$ in the definition of (CB), we see that, because $p = q$ implies $p \sim_b q$, (CB) implies (I). Thus, when the cardinality of A is one, the two conditions are equivalent in the presence of (O), but when A contains two or more elements, (CB) implies (I), but not vice versa.

This justifies my claim that (CB) is a natural generalization of [vNM] independence to unions of mixture spaces, or equivalently, spaces for which the mixture operation is not everywhere defined. Note that if for two mixture spaces each element of the first is strictly better than all elements of the second, then condition (CB) is silent for such comparisons, but it still has implications within each mixture space.

Note that [KS] provide an example that shows why their version of the (CB) is not implied by the combination of (O), (C'ty) and (I). The intuition is that whenever there are at least two distinct indifference sets, both of which contain elements from two particular mixture spaces, it is possible to find utility functions that are mixture preserving on each of the mixture spaces, together satisfy (O), (C'ty) and (I), but which fail to satisfy (CB).

If $|A| = 1$, then $A \times \mathcal{M}$ is in fact a mixture space, and if it satisfies conditions (O), (C'ty) and (CB), then we say that it is a vNM ordered space. More generally, we have the following definition.

Definition 2.3. *[Extended vNM ordered space]*
Let $\mathcal{M}$ be a mixture space and A a discrete set, and let $\precsim$ be a binary relation on $A \times \mathcal{M}$. Then $(A \times \mathcal{M}, \precsim)$ will be referred to as an extended vNM ordered space *if it satisfies conditions (O), (C'ty), and (CB).*

Any representation of preferences over a mixture set will involve a function that is, first and foremost, mixture preserving.

Definition 2.4 (A generalization of [HM] and Moulin (2001)[8])**.** *For any set*

[8] I thank Peter Hammond for recommending this form of the definition.

D, a function $f : D \to \mathcal{R}$ is said to be mixture preserving *(MP) if for all $p, q \in D$, and all $\lambda \in [0, 1]$ such that $p\lambda q \in D$*

$$f(p\lambda q) = \lambda f(p) + (1 - \lambda)f(q).$$

A special case of a mixture preserving function is of course an expected utility function such as that of [vNM].

3 Mixture preserving utility

The next lemma ensures that for each $a \in A$, the vNM ordered space $\{a\} \times \mathcal{M}$, has a mixture preserving representation.

Lemma 3.1. *Let $|A| = 1$. Then $(A \times \mathcal{M}, \precsim) \equiv (\{a\} \times \mathcal{M}, \precsim)$ is a vNM ordered space if and only if there exists a mixture preserving function $U : \{a\} \times \mathcal{M} \to \mathcal{R}$ such that for every $p, q \in \mathcal{M}$,*

$$(a, p) \precsim (a, q) \quad \Leftrightarrow \quad U(a, p) \leqslant U(a, q) \tag{1}$$

Proof. The proof follows from the fact that when $A \times \mathcal{M}$ is a vNM ordered space it satisfies the conditions of [HM]. The only part that is not immediate, is the observation that condition (C'ty) is stronger than the continuity condition of [HM]. Their condition is stated as follows.
For any $p, q, r \in \Delta$, the following sets are closed:

$$\{\lambda \in [0, 1] : p\lambda q \precsim_a r\} \quad \text{and} \quad \{\lambda \in [0, 1] : r \precsim_a p\lambda r\}.$$

We prove the contrapositive. That is, we prove that if [HM]'s condition fails to hold, then so does (C'ty). Suppose there exists a sequence $\{\lambda_n : n \in \mathbb{N}\}$ in $[0, 1]$ that converges to λ' with the property that, for all n, $p\lambda_n q \precsim_a r$, whilst at λ' we have $r \prec_a p\lambda' q$. In the presence of complete preferences, this is the only possibility. Now, this implies that (C'ty) indeed fails to hold. $\square$

Lemma (3.1) also provides justification for the following identity

$$(\{a\} \times \mathcal{M}, \precsim) \equiv (\mathcal{M}, \precsim_a)$$

and we will use the latter as shorthand for the former, thus $(a, p) \sim (a, q)$ if and only if $p \sim_a q$ for example.

The first result that arises from the assumption of compactness is the following.

Lemma 3.2. *Let $[\alpha, \gamma]$ be a closed interval in $\mathcal{R}$ such that $\alpha < \gamma$ and let $(\mathcal{M}, \precsim)$ be a vNM ordered space with $\prec \neq \varnothing$. Then given any mixture preserving representation $\widetilde{U}$ of $(\mathcal{M}, \precsim)$, there exists unique $\theta, \kappa \in \mathcal{R}$ with $\theta > 0$ such that the function*

$$U : \mathcal{C} \to [\alpha, \gamma], \quad p \mapsto \theta \widetilde{U}(p) + \kappa,$$

satisfies

$$p \precsim_{.} q \quad \Leftrightarrow \quad U(p) \leqslant U(q)$$

for each $p, q \in \mathcal{M}$.

Proof of lemma (3.2). Since $(\mathcal{M}, \precsim)$ is a vNM ordered space, by lemma (3.1) there exists a mixture preserving function $\widetilde{U} : \mathcal{M} \to \mathcal{R}$ that represents preferences on $\mathcal{M}$. The existence of mixture preserving representation is sufficient for condition (C'ty), and so $\widetilde{U}$ is continuous. Then since $\mathcal{M}$ is compact, by theorem 26.4 of Munkres p182, the image of $\mathcal{M}$ under $\widetilde{U}$ is compact.

By the extreme value theorem (theorem 27.4 of Munkres 190), this implies that there exists a greatest lower bound g and a least upper bound l, both in $\mathcal{M}$ such that

$$\widetilde{U}(g) \leqslant \widetilde{U}(p) \leqslant \widetilde{U}(l)$$

for every $p \in \mathcal{M}$. Let $\widetilde{\alpha} = \widetilde{U}(g)$ and $\widetilde{\gamma} = \widetilde{U}(l)$. By the fact that $\widetilde{U}$ is a representation, we know that for all $p \in \mathcal{M}$, $g \precsim p \precsim l$.

By condition (2) of the definition of mixture spaces, for every $\lambda \in [0,1]$ $g\lambda l$ is an element of $\mathcal{M}$. Then since $\widetilde{U}$ is mixture preserving, for all $\lambda \in [0,1]$

$$\widetilde{U}(g\lambda l) = \lambda\widetilde{\alpha} + (1-\lambda)\widetilde{\gamma},$$

so that the image of $\widetilde{U}$ is equal to the interval $[\widetilde{\alpha}, \widetilde{\gamma}]$. The fact that $\prec$ is nonempty, and implies that $l \prec g$ and hence $\widetilde{\alpha} < \widetilde{\gamma}$.

Now let θ satisfy the equation $\theta(\widetilde{\gamma} - \widetilde{\alpha}) = \gamma - \alpha$. Then $\theta > 0$ and it is uniquely identified. Next, let κ satisfy $\theta\widetilde{\alpha} + \kappa = \alpha$; it too is uniquely identified. Then since

$$\theta\widetilde{\gamma} + \kappa = \theta\widetilde{\gamma} + \alpha - \theta\widetilde{\alpha} = \gamma,$$

we see that $U := \theta\widetilde{U} + \kappa$ is a candidate for the required function.

The fact that U is mixture preserving follows readily from the fact that $\widetilde{U}$ is mixture preserving and the fact that for each $p, q \in \mathcal{M}$, $0 \leqslant \lambda \leqslant 1$

$$\begin{aligned} U(p\lambda q) :=& \theta\,\widetilde{U}(p\lambda q) + \kappa \\ =& \lambda\,(\theta\,\widetilde{U}(p) + \kappa) + (1-\lambda)(\theta\,\widetilde{U}(q) + \kappa) \\ =& \lambda\,U(p) + (1-\lambda)\,U(q); \end{aligned}$$

whilst the fact that it is order preserving is similarly easy to show. □

Remark 3.3. *If $\prec$ is empty, then for all $p, q \in \mathcal{M}$, $p \sim q$ and so $\widetilde{U}(p) = \widetilde{\alpha}$ for all $p \in \mathcal{M}$. Clearly if α is any other element of $\mathcal{R}$, then there exists an infinity of solutions to the equation $\theta\widetilde{\alpha} + \kappa = \alpha$.*

Lemma 3.4. *Suppose that there exist $a \neq b$ in A and $p, q, p' \in \mathcal{M}$ such that preferences satisfy $(a,p) \precsim (b,q) \precsim (a,p')$ and $(a,p) \prec (a,p')$. Then there exists a unique $\lambda \in [0,1]$ such that $(a, p\lambda p') \sim (b,q)$.*

Proof. Existence: This argument parallels the proof of Theorem 1 of [HM]. Consider the the set

$$T := \{\lambda \in [0,1] : (b,q) \precsim (a, p\lambda p')\}.$$

By condition (C'ty) and the proof of lemma (3.1), T is a closed subset of $[0,1]$. Since $(b,q) \precsim (a,p')$, $1 \in T$; so T is nonempty. By the same argument, $W := \{\lambda \in [0,1] : (a, p\lambda p') \precsim q\}$ is also closed in $[0,1]$; it is nonempty as $0 \in W$. Now since $\mathcal{M}$ is a completely preordered mixture set, $T \cup W = [0,1]$; the fact that $[0,1]$ is a connected set, implies that no pair of its subsets define a separation thereof, and so $T \cap W$ is nonempty. Let $\lambda_0 \in T \cap W$; by construction of these sets I have $(a, p\lambda_0 p') \precsim (b,q) \precsim (a, p\lambda_0 p')$. So that by asymmetry of $\prec$ I have $(a, p\lambda_0 p') \sim (b,q)$.

Uniqueness: Clearly if $(a,p) \prec (b,q) \prec (a,p')$, then $0 < \lambda_0 < 1$. Now take any $r \in \mathcal{M}$ such that $(a,r) \sim (b,q)$. Transitivity of $\sim$ implies that $(a,r) \sim (a, p\lambda_0 p')$. Moreover, theorem 6 of [HM] together with the fact that I have assumed $(a,p) \prec (a,p')$ implies that λ_0 is unique. $\square$

Definition 3.5. *Let x and y be elements of $A \times \mathcal{M}$. Then the pair (x,y) is called a* gap *if both the following conditions hold:*

(i) $x \prec y$; and

(ii) the set $]x,y[:= \{z : x \prec z \prec y\}$ is empty.

Remark 3.6. *Note that more generally the set $]x,y[:= \{z : x \prec z \prec y\}$ can, by de Morgan's laws and the fact that $\precsim$ is complete and transitive, be rewritten as $\{z : \neg(z \precsim x)\} \cup \{z : \neg(y \precsim z)\}$. Then condition (C'ty) implies that each of the sets in this union is open; therefore $]x,y[$ is open.*

Definition 3.7. *Let $x = (a,p)$ and $y = (b,q)$ be elements of $A \times \mathcal{M}$. Then the pair (x,y) is called a* quasi-gap *(or alternatively a* q-gap*) if it is either a gap or if both the following conditions hold:*

(i) $x \sim y$*; and*

(ii) the sets

$$D_x = \{c \in A : \textit{for some } r \in \mathcal{M}, (c,r) \prec x\};\ \textit{and}$$
$$E_y = \{c \in A : \textit{for some } r \in \mathcal{M}, y \prec (c,r)\}$$

are disjoint with $a \in D_x$ *and* $b \in E_y$*.*

Definition 3.8. *In what follows I will denote any element that is a greatest lower bound for a set* $\{c\} \times \mathcal{M}$ *by* g_c*. Similarly,* l_c *will refer the the least upper bound.*

Lemma 3.9. *If* (x, y) *is a q-gap, then for all* $c \in A$*, either* $(c, r) \precsim x$ *for all* $r \in \mathcal{M}$*, or* $y \precsim (c, r)$ *for all* $r \in \mathcal{M}$*.*

Proof. If $x \sim y$, so that (x, y) is a q-gap but not a gap, then in property (ii) of the definition of a q-gap, the fact that the sets D_x and E_y are disjoint is sufficient for the conclusion of this lemma to hold.

If $x \prec y$, then the conclusion of this lemma is equivalent to the following:

$$\forall c \in A \quad \neg(x \prec l_c \text{ and } g_c \prec y)$$

which in turn, by the fact that $]x, y[= \varnothing$, is equivalent to:

$$\forall c \in A \quad \neg(g_c \precsim x \prec y \precsim l_c).$$

So by way of contradiction, suppose not. If $g_c \sim x \prec y \sim l_c$, then the fact that $\{c\} \times \mathcal{M}$ is a vNM ordered set together with theorem 2b of [HM] implies that there exists $z \in \{c\} \times \mathcal{M}$ such that $g_c \prec z \prec l_c$. Transitivity then implies that $x \prec z \prec y$, so that the pair (x, y) cannot be a gap; therefore, either $g_c \prec x$ or $y \prec l_c$ must hold.

Suppose that $g_c \prec x$, then the fact that $g_c \prec x \prec l_c$ together with by lemma (3.4) there exists $v \in \{c\} \times \mathcal{M}$ such that $v \sim x$. Repeating the argument of the previous paragraph leads us to the same contradiction. The case where $y \prec l_c$ is identical. $\square$

Lemma 3.10. *Every q-gap (x, y) of $(A \times \mathcal{M}, \precsim)$ has the property that $x \sim l_a$ and $y \sim g_b$ for some $a \neq b$ in A.*

Proof. Suppose that (x, y) is a gap, and suppose that $x \not\sim l_c$ for all $c \in A$. Then there exists $d \in A$, $p, q \in \mathcal{M}$ such that $x = (d, p) \prec l_d = (d, q)$. Now since $(\mathcal{M}, \precsim_d)$ is a vNM ordered space, by theorem 2b of [HM] I have $p \prec_d p\frac{1}{2}q \prec_d q$, so that because $]x, y[= \varnothing$, and $\precsim_d$ is complete, $y \precsim \big(d, p\frac{1}{2}q\big)$. Now the same theorem implies that $y \precsim \big(d, p\frac{1}{2}(p\frac{1}{2}q)\big)$ which, by conditions (2) and (3) of the definition of a mixture space

$$p\tfrac{1}{2}(p\tfrac{1}{2}q) = p\tfrac{1}{2}(q\tfrac{1}{2}p) = (q\tfrac{1}{2}p)\tfrac{1}{2}p = q\tfrac{1}{4}p = p(1 - \tfrac{1}{4})q.$$

Thus $y \precsim \big(d, p(1-\frac{1}{4})q\big)$. Indeed by the induction hypothesis, and an identical argument I see that for all $j \in \mathbb{N}$ $y \precsim \big(d, p(1 - 2^{-j})q\big)$, so that by condition (C'ty) and the fact that $(d, p(1 - 2^{-j})q)$ converges to $x = (d, p)$, we see that $x \sim y$. This implies that (x, y) cannot be a gap; this contradiction implies that there exists $a \in A$ such that $x \sim l_a$. An identical argument shows that $y \sim g_b$ for some $b \in A$ whenever (x, y) is a gap. Moreover, by the definition of a gap, $a \neq b$.

Now suppose that $x \sim y$, that is (x, y) is a q-gap, but not a gap. By property (ii) of q-gaps, there exists $d \in E_y \backslash D_x$ with $y = (d, p)$ for some $p \in \mathcal{M}$. The fact that d lies outside D_x implies, by condition (O), that $x \precsim (d, r)$ for all $r \in \mathcal{M}$, so that $y \sim g_d$. Finally, by the same property of q-gaps, there exists $(c', p') \in A \times \mathcal{M}$ such that $(c', p') = x$ and $c' \in D_x \backslash E_y$. Now the fact that c lies outside E_y implies that for all $r \in \mathcal{M}$, $(c, r) \precsim y$, so that by $(c, r) \precsim (c, p')$ for all $r \in \mathcal{M}$. This implies that $x \sim l_c$; moreover, I can see that $d \neq c$. $\square$

I will now use the quasi-gaps of $(A \times \mathcal{M}, \precsim)$ to define quasi-components. These are subsets of $A \times \mathcal{M}$ upon which my eventual representation will have a common scale and origin. That is, on a given quasi-component, in the language of measurability and comparability, preferences are cardinally measurable and fully comparable.

Definition 3.11 (Quasi-component)**.** *Let (x, y) be a quasi-gap in $(A \times \mathcal{M}, \precsim)$.*

(i) If for all other q-gaps (w, z), $x \precsim w$, then the (nonempty) set

$$[\leftarrow, x] := \{t \in A \times \mathcal{M} : x \precsim t\}$$

is called a quasi-component or q-component of $(A \times \mathcal{M}, \precsim)$.

(ii) If for all other q-gaps (w, z), $z \precsim y$, then the (nonempty) set

$$[y, \rightarrow] := \{t \in A \times \mathcal{M} : y \precsim (a, p)\}$$

is called a quasi-component of $(A \times \mathcal{M}, \precsim)$.

(iii) If (u, v) is another q-gap, distinct from (x, y) with $y \prec v$, and for all other q-gaps (w, z), $z \precsim y$ or $u \precsim w$, then the (nonempty) set

$$[y, u] := \{t \in A \times \mathcal{M} : y \precsim t \precsim u\}$$

is called a quasi-component or q-component of $(A \times \mathcal{M}, \precsim)$.

(iv) If there are no q-gaps in $(A \times \mathcal{M}, \precsim)$, then the set $A \times \mathcal{M}$ is itself a quasi-component.

Definition 3.12 (Component)**.** *If the quasi-gap(s) that identify a quasi-component $[u, v]$ are gap(s), then $[u, v]$ is also called a component.*

Definition 3.13. *A quasi-gap (w, z) is said to be distinct from another quasi-gap (x, y) if either $w \prec x$ or $y \prec z$. One quasi-component is said to be distinct from another if at least one of the quasi-gaps that define it is distinct from each of the quasi-gaps of the other. A collection of distinct quasi-components is such that every pair in the collection is mutually distinct.*

The following lemma together with the fact that for every q-gap (x, y) the set $]x, y[$ is empty shows that the definition of q-components is such that for all $x \in A \times \mathcal{M}$, x belongs to some q-component.

Lemma 3.14. *For every q-gap (x, y) in $(A \times \mathcal{M}, \precsim)$, there exists $w, z \in A \times \mathcal{M}$ such that $[w, x]$ and $[y, z]$ are q-components.*

Proof. If for every q-gap (w', w) in $(A \times \mathcal{M}, \precsim)$, $x \precsim w'$, then by (i) of the definition of q-components, $[\leftarrow, x]$ is a q-component. The fact that $[\leftarrow, x] = [w, x]$ for some $w \in A \times \mathcal{M}$ follows from the fact that A is finite and $\mathcal{M}$ is compact. The same is true of $[y, \rightarrow]$ if $w \precsim y$ for every q-gap (w', w).

Now suppose that there exists a q-gap (v', v) with $v' \prec x$. I will first show that there exists a q-gap (w', w) with the property that every q-gap (u', u) with $v' \precsim u' \precsim x$ satisfies $u' \precsim w'$.

Suppose not. That is, for every q-gap (w', w) such that $w' \prec x$, there exists another (u', u) with $w' \prec u' \prec x$. Starting with v', I may, using the first element in the pair that defines each of these q-gaps, construct a sequence $\{v'_i : i \in \mathbb{N}\}$ with the property that $v'_1 = v'$ and for all $i \geqslant 2$, $v'_{i-1} \prec v'_i \prec x$. By assumption this sequence is infinite. However, by lemma (3.10) I know that for all i, $v'_i \sim l_{a(i)}$ for some $a(i) \in A$; now the fact that $v'_{i-1} \prec v'_i$ for each i implies that A is an infinite set, a contradiction of my assumption that it is finite.

Thus there exists a q-gap (w', w) such that every other q-gap lies outside the set $[w, x]$; part (iii) of definition (3.11) then states that this is a q-component. $\square$

Lemma 3.15. *For all $a \in A$, $\{a\} \times \mathcal{M}$ is a subset of some q-component*

$[u, v]$. *Moreover, if* $\prec_a \neq \varnothing$, *then* $\{a\} \times \mathcal{M}$ *belongs to at most one (distinct) q-component, otherwise it belongs to at most two.*

Proof. Lemma (3.14) implies that $x = (a, p) \in A \times \mathcal{M}$ lies in some q-component $[u, v]$. In turn, lemma (3.9) implies that the q-gaps that define $[u, v]$ must lie outside $\{a\} \times \mathcal{M}$, and this is sufficient for the first part of lemma.

It is also sufficient for the the fact that $\{a\} \times \mathcal{M}$ belongs to $[u, v]$ alone whenever $\prec_a$ is nonempty. For in this case, even if for instance $g_a \sim u$, there exists $q \in \mathcal{M}$ such that $g_a \prec (a, q)$, so that (a, q) cannot belong to any q-component in the set $[\leftarrow, u]$. In the same way, we know that $l_a \precsim v$, so that $(a, q) \notin [v, \rightarrow]$.

To see that when $\prec_a = \varnothing$, the set $\{a\} \times \mathcal{M}$ can belong to more than one q-component, note that if $g_a \sim u$ and $[y, u]$ is also a q-component for some $y \in A \times \mathcal{M}$, then the fact that $l_a \sim u$ implies that $\{a\} \times \mathcal{M}$ belongs to both q-components. For any other q-component $[w, z]$ that is distinct from both $[y, u]$ and $[u, v]$, either $w \prec y$ or $v \prec z$, in either of these cases, one cannot have $u \in [w, z]$. □

Remark 3.16. *By this last lemma, denote any q-component* $[u, v]$ *as* $B \times \mathcal{M}$ *for some* $B \subset A$ *such that there are no q-gaps in* $B \times \mathcal{M}$. *Moreover, also by this lemma, for some* $Q \in \mathbb{N}$ *I may write*

$$A \times \mathcal{M} = \bigcup_{i=1}^{Q} B_i \times \mathcal{M},$$

where for each i, $B_i \times \mathcal{M}$ *is a quasi-component.*

Lemma 3.17. *For any q-component* $[u, v] := B \times \mathcal{M}$, *if* $u \prec v$, *then there exists* $b \in B$ *such that* $\prec_b \neq \varnothing$ *and* $u \sim g_b$. *If on the other hand* $u \sim v$, *then* $[u, v]$ *is a component.*

Proof. Suppose that despite the fact that $u \prec v$, it holds that for all $b \in B$, the set $\prec_b$ is empty. By this assumption, together with lemma (3.10), we know that $u \sim l_a \sim g_a$ and $v \sim g_c \sim l_c$ for some $a, c \in A$. Now since $[u, v]$ contains no q-gaps, there exists $d \in B \backslash \{a, c\}$ such that $g_a \prec g_d \sim l_d \prec g_c$. In turn, there exists $e \in B \backslash \{a, c, d\}$ such that $g_a \prec g_e \sim l_e \prec g_d$. A finite iteration of this argument exhausts the elements of the finite set B, so that we contradict the fact that there are no gaps in $[u, v]$; I conclude that there exists $b \in B$ with $\prec_b \neq \varnothing$.

An identical argument to the preceding paragraph shows that it is not the case that for all $b \in B$ with $\prec_b \neq \varnothing$ we have $u \prec g_b$; thus $u \sim g_b$ for some such b.

Suppose that $u \sim v$ and $[u, v]$ is a q-component but not a component, that is the q-gaps that identify $[u, v]$ are not gaps. Then let (x, u) (v, y) be quasi-gaps with $x \sim u$ and $v \sim y$. Transitivity, via condition (O), then implies that these two quasi-gaps are not distinct and so $[u, v]$ is not a q-component. □

Lemma 3.18. *The number Q of (distinct) quasi-components in $(A \times \mathcal{M}, \precsim)$ is less than or equal to the cardinality of A. The number of (distinct) quasi-gaps is one less than the number of quasi-components.*

Proof. By lemma (3.15) we know that every q-component $[u, v] = B \times \mathcal{M}$ contains at least one set $\{b\} \times \mathcal{M}$. If if $u \prec v$ then by lemma (3.17) $g_b \prec l_b$ for some $b \in B$ and by lemma (3.15) therefore, $[u, v]$ is the only q-component to which the set belongs. If $u \sim v$, then $g_b \sim l_b$ for all $b \in B$, and $[u, v,]$ is a component, and in this case every b in the nonempty set B satisfies the property that $[u, v]$ is the only q-component to which $\{b\} \times \mathcal{M}$ belongs.

This in itself is sufficient for the proof of the first part of this lemma. For the second part, I proceed by induction. First consider the "lowest" q-component in the order $\precsim$ which is denoted by $[\leftarrow, v_1]$. Either there are no q-gaps in $(A \times \mathcal{M}, \precsim)$, so that $A \times \mathcal{M}$ is the only q-component, or there exists y_1 such that (v_1, y_1) is a q-gap. In the former case the statement of this lemma is true, and in the latter there are once more two possibilities. By lemma (3.14), there exists v_2 such that $[y_1, v_2]$ is a q-component, either $[y_1, v_2] = [y_1, \rightarrow]$ is a q-component, or there exists y_2 such that (v_2, y_2) is a q-gap. Once again, in the former case, the statement of the lemma is true, and in the latter there are two similar possibilities. It is clear that for each $2 \leqslant j \leqslant |A|$ the inductive step is identical to the above and is therefore omitted. $\square$

Lemma 3.19. *If $[u, v] = B \times \mathcal{M}$ is a q-component, then either $u \sim v$ or there exists a minimal sequence $\{a_1, \ldots a_h\}$ in A such that for all $j = 1, \ldots, h$, $\{a_j\} \times \mathcal{M}$ is a subset of $[u, v]$, and*

$$g_1 \prec g_2 \prec l_1 \precsim g_3 \prec l_2 \precsim g_4 \prec \cdots \prec g_{h-1} \prec l_{h-2} \precsim g_h \prec l_{h-1} \prec l_h.$$

Definition 3.20. *The set $\{a_1, \ldots, a_h\} \times \mathcal{M}$ is called a* strict cover *of $[u, v]$.*

Proof. Let $u \prec v$, then by lemma (3.17) we may suppose that $g_b \prec l_b$ for some $b \in B$ and take g_b to be the glb of $[u, v]$. By the fact that B is finite together with condition (O), choose b such that for every c with $g_c \sim u$, $l_c \precsim l_b$. Consider the set

$$L_b := \{a \in A : g_a \prec l_b \prec l_a\}.$$

First observe that $L_b \subset B$, for I know that $l_b \precsim v$. If L_b is empty, then for every $a \in A \backslash L_b$, either $l_b \precsim g_a$ or $l_a \precsim l_b$. The former of these two relationships implies that if for some $p \in \mathcal{M}$, $l_b \prec (a, p)$ then $l_b \precsim g_a$, whilst together they imply for all $a \in B$, $l_a \precsim l_b$. Thus whenever L_b is empty $[u, v] = [g_b, l_b]$ and the sequence I seek is simply $\{g_b, l_b\}$ with $h = 1$.

If L_b is nonempty, then let $a_1 := b$, $g_1 := g_b$ and $L_1 := L_b$. Then take

$a_2 \in L_b$ and $l_2 := l_{a_2}$ to satisfy $l_c \precsim l_2$ for all $c \in L_1$. Such an element exists because of condition (O) and the fact that L_1 is nonempty and finite. Now consider the set

$$L_2 := \{a \in A : g_a \prec l_2 \prec l_a\}.$$

By the same argument as for L_1 above, I see that $L_2 \subset B \backslash \{a_1\}$. If it is empty I have found a strict cover:

$$u \sim g_1 \prec g_2 \prec l_1 \prec l_2.$$

If not, then let a_3 be the element of L_2 satisfying $l_a \precsim l_3$ for all $a \in L_2$. If L_3, defined recursively as for L_1 and L_2 is empty then I claim my sequence is

$$g_1 \prec g_2 \prec l_1 \precsim g_3 \prec l_2 \prec l_3.$$

The only relationship that needs further explanation is $l_1 \precsim g_3$. This holds because otherwise $a_3 \in L_1$ and the fact that $l_2 \prec l_3$ would contradict the fact that l_2 was maximal. The general case follows by induction, the argument of which is identical to the one just given and is hence omitted.

The fact that the strict cover is minimal follows immediately from the construction. $\square$

Lemma 3.21. *If $(a, p) \sim (b, q)$ and $(a, p') \sim (b, q')$, then, for all $0 < \lambda < 1$, $(a, p\lambda p') \sim (b, q\lambda q')$.*

Proof. Without loss of generality, let $(a, p) \prec (a, p')$, so that, by transitivity, $(b, q) \prec (b, q')$. By condition (CB), $(a, p\frac{1}{2}p') \sim (b, q\frac{1}{2}q')$. Successive applications of this condition show that for all dyadic rational numbers, $0 < \pi = \sum_{i=1}^{n(\pi)} \zeta_i/2^i < 1$, where $\zeta_i = 0$ or 1, we have $(a, p\pi p') \sim (b, q\pi q')$. For the remainder of this proof, π will refer to the binary expansion of some dyadic rational number. Recall the fact that the set of such numbers is dense in the real numbers and so every $0 < \lambda < 1$ is the limit of some such sequence

$\{\pi_j : j \in \mathbb{N}\}$.

For $0 < \lambda < 1$ consider the set $T_b := \{x \in A \times \mathcal{M} : x \precsim (b, q\lambda q')\}$. If $\{\pi_j\}$ is such that $\pi_j < \lambda$ for each j and $\lim_j \pi_j = \lambda$, then $(b, q\pi_j q') \prec (b, q\lambda q')$ for all j by theorem 4 of [HM]. By condition (BC) and transitivity, $(a, p\pi_j p') \prec (b, q\lambda q')$ for each j and so $\{(a, p\pi_j p')\} \subset T_b$. By condition (C'ty) T_b is closed, and because $\lim_j p\pi_j p' = p\lambda p'$ this implies that $(a, p\lambda p') \precsim (b, q\lambda q')$.

By the same argument $T_a := \{x \in A \times \mathcal{M} : x \precsim (a, p\lambda p')\}$ is closed and contains the set $\{(b, q\pi_j q')\}$ which converges to $(b, q\lambda q')$. Thus $(b, q\lambda q') \precsim (a, p\lambda p')$, so that by asymmetry of $\prec$ the proof is complete. □

Theorem 3.22 (Mixture preserving representation on q-components)**.** *$B \times \mathcal{M}$ is a quasi-component of the* extended *vNM ordered space $(A \times \mathcal{M}, \precsim)$ if and only if both the following conditions hold.*

(i) There exists a mixture preserving function $U : B \times \mathcal{M} \to \mathcal{R}$ such that for all $x, y \in B \times \mathcal{M}$

$$x \precsim y \quad \Leftrightarrow \quad U(x) \leqslant U(y).$$

(ii) If V is any other function with the same properties as U, then, for some $\theta, \kappa \in \mathcal{R}$ with $\theta > 0$, $V = \theta U + \kappa$.

Proof. If for all $x, y \in B \times \mathcal{M}$, $x \sim y$, then take U to satisfy $U(\cdot) \equiv 1$ and remark (3.3) ensures that every other representation V is the form $V = \theta U + \kappa$ for a one-dimensional set of suitable θ, κ combinations. In the opposite direction, suppose that $B \times \mathcal{M}$ is not a quasi-component. Then by lemma (3.15) $B \times \mathcal{M}$ is a union of q-components and so it contains at least one q-gap which, by lemma (3.10), which I denote by (l_a, g_b) for some $a, b \in B$. Now by the definition of q-gap, a and b belong to different q-gaps, indeed there exists $p, q \in \mathcal{M}$ such that $(a, p) \prec (b, q)$. Of course this contradicts the assumption that on $B \times \mathcal{M}$, $\prec$ is empty.

To prove the theorem, suppose that the set of pairs of elements in $B \times \mathcal{M}$ for which strict preference holds is nonempty. By lemma (3.19) there exists a subset B_h of B with $|B_h| = h \leqslant |A|$ and an enumeration of its elements $\{a_1, \dots, a_h\}$ such that if $g_j := g_{a_j}$ and $l_j := l_{a_j}$, then

$$g_1 \prec g_2 \prec l_1 \precsim g_3 \prec l_2 \precsim g_4 \prec \cdots \prec g_{h-1} \prec l_{h-2} \precsim g_h \prec l_{h-1} \prec l_h.$$

I first show that the conditions that constitute a vNM space imply part (i) of the present theorem. By lemma (3.1) the exists a mixture preserving function $U_1 : \{a_1\} \times \mathcal{M} \to \mathcal{R}$ such that for all $p, q \in \mathcal{M}$

$$(a_1, p) \precsim (a_1, q) \quad \Leftrightarrow \quad U_1(a_1, p) \leqslant U_1(a_1, q)$$

By compactness of $\mathcal{M}$, for some $\beta_1, \gamma_1 \in \mathcal{R}$ I have $U_1(a_1, \mathcal{M}) = [\beta_1, \gamma_1]$ with $U_1(g_1) = U_1(a_1, \underline{p})$ and $U_1(l_1) = U_1(a_1, \overline{p})$ for some $\underline{p}, \overline{p} \in \mathcal{M}$.

Since $g_1 \prec g_2 \prec l_1$, lemma (3.4) implies that there exists unique $0 < \lambda < 1$, and $p(\lambda) := \underline{p}\lambda\overline{p}$ such that $(a_1, p(\lambda)) \sim g_2$. Similarly, let $g_2 = (a_2, \underline{q})$ and $l_2 = (a_2, \overline{q})$, then there exists unique $0 < \nu < 1$ such that $l_1 \sim (a_2, q(\nu))$, where $q(\nu) := \underline{q}\nu\overline{q}$.

By lemma (3.1), the vNM ordered set defined by projecting preferences onto the set $\{a_2\} \times \mathcal{M}$, has a mixture preserving representation $U_2' : \{a_2\} \times \mathcal{M} \to \mathcal{R}$. Thus, with a view to leaving the image of U_1 unchanged in my construction of a mixture preserving representation on the projection of preferences onto $\{a_1, a_2\} \times \mathcal{M}$, I recall lemma (3.2) states that: for any interval $[\beta_2, \gamma_2]$ in $\mathcal{R}$, there exists unique $\theta_2, \kappa_2 \in \mathcal{R}$, $\theta_2 > 0$ such that

$$[\theta_2 U_2'(g_2) + \kappa_2,\ \theta_2 U_2'(l_2) + \kappa_2] = [\beta_2, \gamma_2].$$

The content of two preceding paragraphs suggests that I should choose β_2 and γ_2 to satisfy $\beta_2 = \lambda\beta_1 + (1-\lambda)\gamma_1$ and $\gamma_1 = \nu\beta_2 + (1-\nu)\gamma_2$. Substituting for β_2 in the second of these two equations, I obtain $\gamma_2 = \frac{1}{1-\nu}(\gamma_1 - \nu\beta_2)$. Let $U_2 : \{a_1, a_2\} \times \mathcal{M} \to \mathcal{R}$ be defined as

$$U_2(x) = \begin{cases} U_1(x) & \text{if } x \in \{a_1\} \times \mathcal{M} \\ \theta_2 U_2'(x) + \kappa_2 & \text{if } x \in \{a_2\} \times \mathcal{M} \end{cases}$$

U_2 is clearly mixture preserving and its image is the interval $[\beta_1, \gamma_2]$. In order to show that it is also a representation, it suffices to check pairs x and y where $x := (a_1, p')$ and $y := (a_2, q')$ for some $p', q' \in \mathcal{M}$.

First consider the more straightforward cases, that is where $x \prec g_2 \precsim y$ and $x \precsim l_1 \prec y$. For the former, since

$$U_2(g_2) = \beta_2 = \lambda\beta_1 + (1-\lambda)\gamma_1 = U_2(a_1, p(\lambda)),$$

the fact that $x = (a_1, p') \prec g_2 \sim (a_1, p(\lambda))$ together with theorem 4 and 6 of [HM] imply that there is a unique $\lambda_x < \lambda$ such that $x \sim (a_1, p(\lambda_x))$. In this case,

$$U_2(x) = U_2(a_1, p(\lambda_x)) = \lambda_x\beta_1 + (1-\lambda_x)\gamma_1 < \beta_2 \leqslant U_2(y),$$

and so I conclude that whenever $x \prec g_2 \precsim y$

$$x \precsim y \quad \Leftrightarrow \quad U_2(x) \leqslant U_2(y),$$

as required. (Note that in making this statement I have made use of the fact that in this case there are no pairs x, y such that $y \precsim x$.) For the set of pairs x, y where $x \precsim l_1 \prec y$ the proof that U_2 is a representation on such pairs is the same.

Now consider the remaining case where $x = (a_1, p')$ and $y = (a_2, q')$ for some $p', q' \in \mathcal{M}$ and $g_2 \precsim x, y \precsim l_1$. I recall that

$$(a_2, \underline{q}) = g_2 \sim (a_1, p(\lambda)) \quad \text{and} \quad (a_1, \overline{p}) = l_1 \sim (a_2, q(\nu)),$$

and once again by theorem 6 of [HM], there exists a unique $0 \leqslant \mu \leqslant 1$ with $p' \sim_{a_1} p(\lambda)\mu\overline{p}$ such that $x \sim \big(a_1, p(\lambda)\mu\overline{p}\big)$. By lemma (3.21) therefore

$$x \sim \big(a_1, p(\lambda)\mu\overline{p}\big) \sim \big(a_2, \underline{q}\mu q(\nu)\big).$$

Now suppose that $x \sim y$. In this case, transitivity implies that $y \sim \big(a_2, \underline{q}\mu q(\nu)\big)$. Then since $U_2(a_1, p(\lambda)) = \beta_2 = U_2(a_2, \underline{q})$ and $U_2(1, \overline{p}) = \gamma_1 = U_2(a_2, q(\nu))$ and U_2 is mixture preserving (whenever the mixture operation is defined) I have

$$U_2(a_1, p') = U_2(a_1, p(\lambda)\mu\overline{p}) = \mu\beta_2 + (1-\mu)\gamma_1$$

and

$$U_2(a_2, q') = U_2(a_1, \underline{q}\mu q(\nu)) = \mu\beta_2 + (1-\mu)\gamma_1$$

as required.

The remaining cases, where $x \prec y$ and $y \prec x$ follow by virtue of the following facts: by theorem 6 of [HM], I may find two unique values $0 < \mu_x, \mu_y < 1$ such that $p' \sim_{a_1} p(\lambda)\mu_x\overline{p}$ and $q' \sim_{a_2} \underline{q}\mu_y q(\nu)$; by lemma (3.21) I know that $(a, p(\lambda)\mu\overline{p}) \sim (a_2, \underline{q}\mu q(\nu))$ for all $0 \leqslant \mu \leqslant 1$; and by theorem 4 of [HM] I have $\mu_x < \mu_y$ if and only if both

$$\underline{q}\mu_x q(\nu) \prec_{a_2} \underline{q}\mu_y q(\nu) \quad \text{and} \quad p(\lambda)\mu_x\overline{p} \prec_{a_1} p(\lambda)\mu_y\overline{p}.$$

Thus far we've seen that U_2 is a mixture preserving representation of the projection of preferences $\precsim$ onto the subset $\{a_1, a_2\} \times \mathcal{M}$ of the strict cover $B_h \times \mathcal{M}$. In order to extend U_2 to the rest of this strict cover, I proceed by induction and the proof follows by precisely the same argument as above. The resulting representation is a function $U_h : B_h \times \mathcal{M} \to [\beta_1, \gamma_h]$ that is a standard mixture preserving representation on each of the sets $\{a_j\} \times \mathcal{M}$ and an extended mixture preserving representation on the whole set.

In order to extend U_h to the rest of the quasi-component $B \times \mathcal{M}$, note that for all $b \in B \backslash B_h$, $g_1 \precsim g_b \precsim l_b \precsim l_h$. In fact, because a strict cover is minimal, I conclude that $g_j \precsim g_b \precsim l_b \precsim l_{j+1}$ for some $j \in \{1, \ldots, j-1\}$. (See the proof of lemma (3.19) for the construction of a strict cover.) Thus for some $x, y \in \{a_j, a_{j+1}\} \times \mathcal{M}$ I have $x \sim g_b$ and $y \sim l_b$.

Now by the proof that U_2 is a representation of the projection of preferences onto $\{a_1, a_2\} \times \mathcal{M}$, and the fact that my construction of U_h is recursive in j, I know that for each j, U_h restricted to $\{a_j, a_{j+1}\} \times \mathcal{M}$ will have image $[\beta_j, \gamma_{j+1}]$ where $\beta_j < \gamma_{j+1}$. If I choose $\beta_b = U_h(x)$ and $\gamma_b = U_h(y)$ then by lemma (3.2) or remark (3.3) given any representation U_b of the projection of preferences onto $\{b\} \times \mathcal{M}$ there exists a positive affine transformation $\theta_b > 0$, $\kappa_b \in \mathcal{R}$ that is geometrically unique and satisfies

$$[\beta_b, \gamma_b] = [\theta_b U_b(g_b) + \kappa_b, \theta_b U_b(l_b) + \kappa_b].$$

Indeed because this is true for all $b \in B \backslash B_h$, the proof that

$$U_B(x) := \begin{cases} U_h(x) & \text{if } x \in B_h \times \mathcal{M} \\ \theta_b U_b(x) + \kappa_b & \text{if } x \in \{b\} \times \mathcal{M},\ b \in B \backslash B_h \end{cases}$$

is a representation for the projection of preferences onto the quasi-component $B \times \mathcal{M}$ follows by the same techniques that I have used in showing that U_2 is a representation.

It remains to be shown that the fact that $(A \times \mathcal{M}, \precsim)$ is a vNM ordered space implies part (ii) of the present theorem. That is, if V is any other mixture preserving representation of there exists a single positive affine transformation $\theta > 0$, $\kappa \in \mathcal{R}$ such that $V = \theta U + \kappa$. If $V(B \times \mathcal{M}) = [\pi, \rho]$ for some $\pi < \rho$ in $\mathcal{R}$, then by the proof of lemma (3.2), I let θ satisfy $\theta(\gamma_h - \beta_1) = \rho$ and κ satisfy $\theta\beta_1 + \kappa = \pi$. I recall that in the construction of the image of U_h, the representation of the strict cover of $B \times \mathcal{M}$, only β_1 and γ_1 were free variables. The degrees of freedom associated with every other representation $\widetilde{U}_j$ of $\{a_j\} \times \mathcal{M}$ were used to obtain the unique transformation $\theta_j \widetilde{U} + \kappa_j$ with image $[\beta_j, \gamma_j]$, where for each j, $\beta_j = \lambda_j \beta_{j-1} + (1 - \lambda_j)\gamma_{j-1}$ and $\gamma_j = \frac{1}{1-\nu_j}(\gamma_{j-1} - \nu_j \beta_j)$ and where the λ_j and ν_j are uniquely determined by preferences. As such it suffices to check that U and V agree on $\{a_1\} \times \mathcal{M}$. This of course directly follows by lemma (3.2).

This complete the proof that the conditions that constitute an extended vNM ordered space imply (i) and (ii) of the present theorem. In the opposite direction, conditions (O), (C'ty) and (I) are standard and therefore omitted. The necessity of (CB) is seen by noting that if it fails, that is there exists $a, b \in B$, $p, q, p', q' \in \mathcal{M}$ with $(a, p) \sim (b, q)$ and $(a, p') \sim (b, q')$ but

$$(a, p\tfrac{1}{2}p') \nsim (b, q\tfrac{1}{2}q').$$

Without loss of generality, suppose that $(a, p) \prec (b, q)$, then for any mixture preserving representation U of such preferences $\beta = U(a, p) = U(b, q)$ and $\gamma = U(a, p') = U(b, q')$ for some $\beta < \gamma$ in $\mathcal{R}$. However, the fact that U is mixture preserving on each of $\{a\} \times \mathcal{M}$ and $\{b\} \times \mathcal{M}$ implies

$$\begin{aligned} U(p\mu p') &= \mu\beta + (1 - \mu)\gamma \\ &= U(b, q\mu q') \end{aligned}$$

for all $0 < \mu < 1$; which clearly implies the desired contradiction.

The necessity of the assumption that $B \times \mathcal{M}$ is a quasi-component of $(A \times \mathcal{M}, \precsim)$ follows from the fact that if it is not then there exists at least two quasi-components $B' \times \mathcal{M}$ and $B'' \times \mathcal{M}$ whose union is $B \times \mathcal{M}$. I will obtain a contradiction for the case where there are only two q-components. Let g', g'', l' and l'' be the respective glbs and lubs of B' and B''. It will suffice to consider the cases where they are components, that is $g' \prec l' \prec g'' \prec l''$, and where they are quasi-components but not components, that is where $g' \prec l' \sim g'' \prec l''$.

First consider the case where $l' \prec g''$. Let U be a mixture preserving representation of $B \times \mathcal{M}$ with $U(g') = \beta'$, $U(l') = \gamma'$, $U(g'') = \beta''$ and $U(l'') = \gamma''$. Now let $[\pi', \rho']$ and $[\pi'', \rho'']$ be any pair of intervals such that $\pi' < \rho' < \pi'' < \rho''$. It is clear that in general I will need two different positive affine transformations $\theta', \theta'' > 0$ and $\kappa', \kappa'' \in \mathcal{R}$ are needed to shift and rescale the intervals $[\beta', \gamma']$ and $[\beta'', \gamma'']$ so that $[\theta'\beta' + \kappa', \theta'\gamma' + \kappa'] = [\pi', \rho']$ and $[\theta''\beta'' + \kappa'', \theta''\gamma'' + \kappa''] = [\pi'', \rho'']$. However it is also clear that any mixture preserving function that maps $B' \times \mathcal{M}$ into $[\pi', \rho']$ and $B'' \times \mathcal{M}$ into $[\pi'', \rho'']$ and which is representation on each of these quasi-components is also a representation on $B \times \mathcal{M}$. This implies that (ii) is violated. So that in this case, the fact that $B \times \mathcal{M}$ is a quasi-component is necessary.

The case where $g' \prec l' \sim g'' \prec l''$ follows by an identical argument and is therefore omitted. This completes the proof of the theorem. $\square$

Theorem 3.23 (Mixture preserving representation)**.**
Let $\mathcal{M}$ be a compact mixture space and A a finite set. Then $(A \times \mathcal{M}, \precsim)$ is an extended *vNM ordered space if and only if both the following conditions hold.*

(1) There exists a mixture preserving function $U : A \times \mathcal{M} \to \mathcal{R}$ such that for all $x, y \in A \times \mathcal{M}$

$$x \precsim y \quad \Leftrightarrow \quad U(x) \leqslant U(y).$$

(2) the number Q of quasi-components[9] has the property that if V is any

[9]For those who have not read the construction of quasi-components above, these should be interpreted as corresponding to maximal intervals in the image of the utility function

other function satisfying (1), then for each $a \in A$

$$V(a,\cdot) = \theta_i U(a,\cdot) + \kappa_i$$

for some $i = 1,\ldots,Q$, $\kappa_i \in \mathcal{R}$ and $\theta_i > 0$.

Proof. ($\Rightarrow$ (1)) Since $A \times \mathcal{M}$ can be written as the union of an arbitrary enumeration of its q-components $A_i \times \mathcal{M}$, $i = 1,\ldots,Q \leqslant |A|$. (Recall that we are implicitly referring to distinct q-components.) Each q-component, by theorem (3.22), has a mixture preserving representation $\widetilde{U}^i : \{a_i\} \times \mathcal{M} \to [\beta_i', \gamma_i']$ for some $\beta_i' \leqslant \gamma_i'$ in $\mathcal{R}$. I proceed by rearranging the intervals $[\beta_i', \gamma_i']$, to match the order $\precsim$.[10]

The proof then follows by lemma (3.2) as we are free to take the necessary positive affine transformations of the representations $\widetilde{U}^i$ such that the transformed representation U^i has image equal to the desired interval $[\beta_i, \gamma_i]$. My mixture preserving representation will be the function that coincides with each of the functions U^i for each q-component $A_i \times \mathcal{M}$.

Let $[\beta_1, \gamma_1] = [0, 1]$ if $\beta_1' < \gamma_1'$ and $\{0\}$ otherwise. By induction, for the i_{th} interval in the enumeration, consider the following cases:

- $l_i \prec g_m$ for all $m \leqslant i-1$: if so let $[\beta_i, \gamma_i] = [-2^i - 1, -2^i]$ if $\beta_i' < \gamma_i'$ and $\{-2^i\}$ otherwise;
- $l_m \prec g_i$ for all $m \leqslant i-1$: if so let $[\beta_i, \gamma_i] = [2^i, 2^i + 1]$ if $\beta_i' < \gamma_i'$ and $\{2^i\}$ otherwise;

in $\mathcal{R}$ that are non-overlapping except, possibly, at the endpoints. It is this slightly messy property that leads the construction to be involved even though the concept is straightforward.

[10] This method of proof allows for a countable set A, although my construction of q-components precludes the theorem from applying to that case.

- $l_h \prec g_i \precsim l_i \prec g_k$ for some $1 \leqslant h, k \leqslant i-1$ and for every $m \leqslant i-1$ I have either $l_m \precsim l_h$ or $g_k \precsim g_m$.

If the last of these is true and $\beta_i' < \gamma_i'$, then let $\beta_i = \frac{1}{2}(\gamma_h + \beta_k - \frac{\beta_k - \gamma_h}{2^i})$ and $\gamma_i = \frac{1}{2}(\gamma_h + \beta_k + \frac{\beta_k - \gamma_h}{2^i})$ and let them both equal $\frac{1}{2}(\gamma_h + \beta_k)$ otherwise.

In the first of these cases, for each $i \geqslant 2$, then $\gamma_i = -2^i < -2^{i-1} - 1$ where $-2^{i-1} - 1$ is a lower bound for β_m for every $m \leqslant i-1$. The second case is similar. For the third case, it is also clear that for all $i \geqslant 2$ $\gamma_h < \beta_i \leqslant \gamma_i < \beta_k$.

Now consider the remaining cases where the q-component $[g_i, l_i]$ shares an indifference set with either one or two q-components with respect to which it is distinct. Recall, that by lemma (3.17) in this case, we cannot have $g_i \sim l_i$.

- $l_i \sim g_h$, where $g_h \precsim g_m$ for all $m \leqslant i-1$: if so let $[\beta_i, \gamma_i] = [\beta_h - 1, \beta_h]$;
- $l_k \sim g_i$, where $l_m \precsim l_k$ for all $m \leqslant i-1$: if so let $[\beta_i, \gamma_i] = [\gamma_k, \gamma_k + 1]$;
- $l_h \precsim g_i \prec l_i \precsim g_k$ for some $1 \leqslant h, k \leqslant i-1$ and for every $m \leqslant i-1$ we have either $l_m \precsim l_h$ or $g_k \precsim g_m$.

If the last of these is true, then there are three sub-cases to consider. If $l_h \sim g_i \prec l_i \sim g_k$, then I simply set $[\beta_i, \gamma_i] = [\gamma_h, \beta_k]$. If $l_h \prec g_i \prec l_i \sim g_k$, then let $\gamma_i = \beta_k$ and $\beta_i = \beta_k - \frac{\beta_k - \gamma_h}{2^i}$. Similarly, if $l_h \sim g_i \prec l_i \prec g_k$, then I let $\beta_i = \gamma_h$ and $\gamma_i = \gamma_h + \frac{\beta_k - \gamma_h}{2^i}$. $\square$

For the special case, where Δ is the set of probability distributions over fixed finite set, I obtain the following representation.

Theorem 3.24 (vNM for $A \times \Delta$)**.**
Let A and S be finite sets and let Δ be the set of probability measures on S. Then $(A \times \Delta, \precsim)$ is an extended *vNM ordered space if and only if both the following conditions hold.*

(1) There exists a function $U : A\times\Delta \to \mathcal{R}$, $(a,p) \mapsto U(a,p) = \sum_{s\in S} p_s u(a,s)$ such that for all $x, y \in A \times \Delta$

$$x \precsim y \quad \Leftrightarrow \quad U(x) \leqslant U(y).$$

(2) The number Q of quasi-components[11] *satisfies the property that if V is any other function satisfying (1), then for each $a \in A$*

$$V(a,\cdot) = \theta_i U(a,\cdot) + \kappa_i$$

for some $i = 1,\ldots,Q$, $\kappa_i \in \mathcal{R}$ and $\theta_i > 0$.

To conclude this section, the following remark describes the relationship between condition (CB) and the corresponding condition of [KS].

Remark 3.25. *The following is the constrained independence condition of [KS]. I emphasize that this is not assumed anywhere in this paper, and is only presented so as to show that it is somewhat stronger than (CB).*

Definition 3.26 (Constrained independence)**.** *For any $b, c \in A$, $p, q, p', q' \in \mathcal{M}$, and $\lambda \in [0,1]$ if $(b,p) \sim (c,q)$ then $(b,p') \precsim (c,q')$ if and only if*

$$(b, p\lambda p') \precsim (c, q\lambda q').$$

The fact that constrained independence implies condition (CB) follows immediately if we take $\lambda = \frac{1}{2}$, consider the fact that $\sim\subset\precsim$ and use the "only if" part of the statement. Moreover, it is clear that in the absence of condition (O), (CB) does not imply the "if" part of constrained independence, for (b,p') and (c,q') may be incomparable. We claim without proof, that in the presence of (O) and (C'ty) the conditions are equivalent

[11] We recall that these should be interpreted as corresponding to maximal intervals in the image of the utility function in $\mathcal{R}$ that are non-overlapping except, possibly, at the endpoints.

The remaining conditions of [KS] serve only to ensure that the representation is state-independent and is particular to the relationship between the the set A and the mixture space $\mathcal{M}$ that they consider. That is to say, their paper is a special case, with additional conditions on preferences and structure on the space on which preferences are defined. Thus to my knowledge, this paper constitutes the most general in its class, and is the natural generalization of vNM and [HM] to the setting where there are is more than one mixture space, or where the mixture operation is not everywhere defined.

4 Separating the issue of comparability from state-dependence

Consider the following change in interpretation of the space $A \times \mathcal{M}$ of the above theorems. Take A to be the product space

$$\prod_{s \in S} A_s$$

for some finite set of states S and, for each $s \in S$ a finite set of state-outcomes A_s. Then take $\mathcal{M}$ to be the set of probability distributions $\Delta(S)$ on S. Now as an element of $A \times \Delta(S)$ is of the form $(a, p) := \big((a_1, \ldots, a_{|S|}), (p_1, \ldots, p_{|S|})\big)$ which, in the case where the elements of A are also vectors in $\mathcal{R}^{|S|}$, may be rewritten as an inner product

$$\langle a, p \rangle = \sum_{s \in S} p_s \, a_s.$$

Or, in the usual lottery form where $\delta(a_s)$ is the lottery assigning probability one to the state-outcome a_s

$$\sum_{s \in S} \delta(a_s) \, p_s.$$

Now for each a in the finite set A, $a \times \Delta(S)$ is a mixture space, and so by theorem (3.24) preferences satisfy (O), (C'ty) and (CB), if and only if there exists a function $U : A \times \Delta(S) \to \mathcal{R}$, $(a, p) \mapsto U(a, p) = \sum_{s \in S} p_s u_s(a_s)$ such that for all $(b, q), (c, r) \in A \times \Delta$

$$(b, q) \precsim (c, r) \quad \Leftrightarrow \quad \sum_{s \in S} u_s(b_s)\, q_s \leqslant \sum_{s \in S} u_s(c_s)\, r_s\,.$$

Note that the means to which we obtain this may be easier to see if we write $u_s(a_s) = u(a_s, \delta_s)$, but it is standard in the literature to make this small abuse of notation. Now this is clearly a state dependent utility function. That is, for any pair of states $s, t \in S$, it may well be that $b_s = b_t = b_{.}$, $c_s = c_t = c_{.}$, $u_s(b_{.}) < u_s(c_{.})$ and $u_t(c_{.}) < u_t(b_{.})$ simultaneously hold. By contrast, theorem (3.24) implies that the uniqueness properties of the representation depend entirely upon preferences. Suppose that the extended vNM ordered set $(A \times \Delta(S), \precsim)$ defines a single component. In this case, any other representation of preferences can be rewritten as a positive affine transformation of U. Thus, in the presence of condition (CB), full comparability across states and state-dependence of preferences is possible.

Indeed, this is the property that ensures the representation of [KS] and Karni (2009) has the uniqueness form that it does. They obtain this by assuming a condition called "coordinate essentiality". This states that, for each state s, there exists $b, c \in A$ such that $(b, \delta_s) \prec (c, \delta_s)$. Whilst this condition is sufficient for the representation to be unique in the sense I have just described, contrary to what [KS] and Karni (2009) claim, it is certainly not necessary. The following example highlights this fact.

Example 4.1. *Suppose there are two states. Take A_1 to be a singleton.*[12] *Then preferences do not satisfy the condition in question. Now take A_2 to be of cardinality 2 and note that the cardinality of A is therefore also two. Now suppose that $(b, \delta_2) \prec (c, \delta_2)$ and that $(b, \delta_1) \equiv (c, \delta_1) \prec (b, \delta_2)$. Then by lemma (3.4) we know that there exists a unique $0 < \lambda < 1$ such that*

[12] We could also suppose that $(b, \delta_1) \sim (c, \delta_1)$ for all $b, c \in A$.

$(b, \delta_2) \sim (c, \delta_1 \lambda \delta_2)$. *We conclude from the fact that there are two distinct indifference sets each of which contains elements of both the mixture spaces* $\{b\} \times [0,1]$ *and* $\{c\} \times [0,1]$, *and so given a particular utility representation* U *of preferences, any other can be written as a positive affine transformation of* U.

The "certainty principle" of [KS] states that for all states s, when the decision-maker is certain of being in state s, what might happen in other states doesn't matter, that is, she reduces the comparison of (a, δ_s) and (b, δ_s) to the comparison of a_s and b_s. Whilst this is not necessary for the uniqueness discussion above, it is necessary and sufficient for there to be a state-independent utility representation of the form

$$U : A \times \Delta(S) \to \mathcal{R}, \quad (a, p) \mapsto U(a, p) = \sum_{s \in S} p_s u(a_s),$$

where now, since $u(a_s, \delta_s) = u(a_s)$ for all s, we see that the function u depends only on the state through the state-outcome a_s.

5 Incompletely defined preferences and the Allais Paradox

By preferences that are incompletely defined we mean that there exist entire sets of lotteries, say, that the decision-maker has not even considered, but wherever preferences are defined they are complete. This is a special case of incomplete preferences which holds when there exists at least one pair of alternatives for which preferences have nothing to say, perhaps as a result of the decision-maker genuinely being unable to state preference in either direction. For instance, for incomplete preferences in general, it may be the case that x is worse than both y and z, whilst y and z are incomparable as in some ways y dominates z and in others z dominates y. The following diagram presents such a scenario for a standard (element by element) ordering of vectors in R^2.

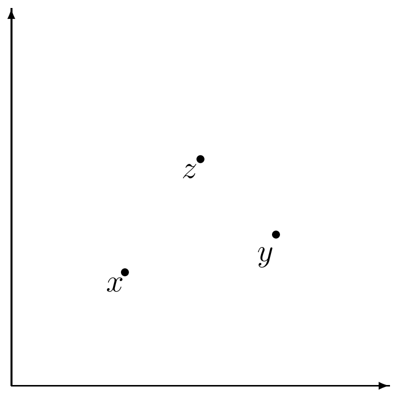

Instead, by incompletely defined preferences we suppose that comparability is transitive. When two lotteries, x and y, are comparable, we mean that either x is preferred to y or y is preferred to x. By transitivity of the relation of of being comparable that if x is comparable to both y and z, then y is comparable with z. If we denote "comparable" by "$\leftrightarrow$" then this condition is summarized thus: for all x, y and z

$$\text{if} \quad x \leftrightarrow y \quad \text{and} \quad y \leftrightarrow z \quad \text{then} \quad x \leftrightarrow z.$$

This assumption ensures that the sets where preferences are complete are distinct from those where they are complete. By itself it is stronger than the usual transitivity of preference condition, but it is not strong enough to partition X into two sets, one containing comparable elements and the other elements for which preferences are not defined. There may be two or more sets within each of which all elements are comparable, but across which elements are not, or at least not comparable in the same way.

The concept at play is closely related to the approach of Schmeidler (1989). On p.576 he concedes that completeness is the most restrictive and imposing assumption of expected utility theory, but notes that "*One can view the weakening of the completeness assumption as a main contribution of all the* *other axioms.*" Indeed he then goes on to weaken the independence condition so that it applies only for a particular subclass of lotteries. The present application pertains to the same viewpoint.

Specifically, suppose that, for a fixed set of prizes such as $\{\$0, \$10, \$50\}$, a given decision-maker has been offered the choice between two pairs of lotteries over the prizes such as the pair $p = 0.01 \cdot \delta_0 + 0.99 \cdot \delta_{10}$ (win nothing with probability 0.01 and win \$10 with probability 0.99) and

$$q = 0.17 \cdot \delta_0 + 0.83 \cdot \delta_{50},$$

and the pair

$$r = 0.9 \cdot \delta_0 + 0.1 \cdot \delta_{50} \quad \text{and} \quad u = 0.11 \cdot \delta_0 + 0.89 \cdot \delta_{10},$$

where, once again, δ_x is the measure assigning probability one to outcome x. Now the chances are she will not have given much thought to the vast number of other possible lotteries over the set of prizes. However, she would almost certainly agree that both of p and q are strictly better than either r or u. In fact this may be so apparent that, when making such comparisons, the decision-maker may have no need for the high resolution scale of measurement the vNM model implies. After all, is it not the case that a person who is asked to state which of two rather different weights is heavier would have no need to ask if they could use scales before providing an answer?

By contrast, her decisions between p and q, on the one hand, and r and u on the other, are likely to require a good deal more consideration. Moreover, whilst for a given pair, such as p and q, mixtures nearby or in between may give rise to similar judgements, and be approximated by the vNM conditions, this need not hold globally. Unless more lotteries are placed before the decision-maker, thus allowing her to explore her own attitude to risk, her preferences may not even be defined globally. Near r and u, for instance, there may be another region where the model of vNM is a good approximation, but across the two regions, other than ordinal statements of the form "*anything in the region near p and q dominates anything in the region near r and u*, the decision-maker may be agnostic.

It seems reasonable that the utility function that characterizes such a decision-maker's preferences should reflect this reflect this asymmetry in the decision-making task across and within regions of the simplex. Furthermore, without further probing by the experimenter, say, the decision-maker may have no cause whatsoever to compare mixtures of elements in one region with those of the other, and if preferences are not defined there, then why should the utility function be?

In the example we have been considering we have so far left the preferences of the decision-maker unspecified. If however they took the form we see in figure (5) (see final page) we obtain an example of preferences that satisfy the Allais paradox, and which are also well described by the discussion above. If, as in the shaded regions of the diagram, we assume that each of these sets are convex, then they are mixture spaces and the model of this paper provides a simple way of representing preferences with a family of mixture preserving utility functions (a simple generalization of an expected utility function that is defined below). Each member of this family being defined on one of the regions for which preferences are well defined. Moreover, these functions combine to define a single mixture preserving utility function on the union of the given regions.

On the complement of this union, preferences and hence the above utility function is simply not defined. However, if we extend each of the mixture preserving utility functions from their domain of definition to the whole simplex, each will take the form of a vNM expected utility function (also defined below). Together, over the entire simplex, these combine to form a multi-expected utility function that resembles that of Dubra, Maccheroni and Ok (2001).

The resulting representation may well fail to characterize preferences in the sense that, for some pair of lotteries p, q each of the vNM expected utility functions may happen to agree that p assigns greater utility than q, even though they both lie outside the regions where preferences are defined. Nonetheless, further research into finding appropriate conditions may provide a way of completing preferences for such pairs.

Another important difference between the representations is in the uniqueness properties the multi-expected utility functions possess. In the present model preferences may be such that each member of the family of utilities is numerically fully comparable with one or more of the others. That is, there may even be a single (multi-)utility scale. This may well be useful in applications.

References

[1] Ernest W. Adams. Elements of a theory of inexact measurement. *Philosophy of Science*, 32(3/4):205–228, 1965.

[2] S. Barberá, P.J. Hammond, and C. Seidl. *Handbook of Utility Theory: Extensions*. Handbook of Utility Theory. Kluwer Academic Publishers, 2004.

[3] Truman F. Bewley. Knightian decision theory: Part I. *Cowles foundation working paper.* No. 807, Yale University, 1986.

[4] K.G. Binmore. *Rational decisions.* Gorman lectures in economics. Princeton University Press, 2009.

[5] C Blackorby, D Donaldson, and J.A Weymark. Social choice with interpersonal utility comparisons: a diagrammatic introduction. *International Economic Review*, 25(2):327–356, 1984.

[6] N. Bourbaki. *Elements of Mathematics: Algebra 1.* Number v. 2. Springer, 1998.

[7] Darius Braziunas and Craig Boutilier. Assessing regret-based preference elicitation with the utpref recommendation system. *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC'10)*, pages pp.219–228, 2010.

[8] S. H. Chew. Axiomatic utility theories with the betweenness property. *Annals of Operations Research*, 19(1):273–298, 1989.

[9] C. d'Aspremont and L. Gevers. Social welfare functionals and interpersonal comparability. *Handbook of social choice and welfare*, 1:459–541, 2002.

[10] M.H.A. Davis. *Markov models and optimization.* Monographs on statistics and applied probability. Chapman & Hall, 1993.

[11] G Debreu. Representation of a preference ordering by a numerical function. *Decision processes*, pages 159–165, 1954.

[12] G Debreu. Topological methods in cardinal utility theory. *Cowles Foundation Discussion Papers*, 1959.

[13] G. Debreu. Continuity properties of paretian utility. *International Economic Review*, 5:285–293, 1964.

[14] E. Dekel. An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom* 1. *Journal of Economic Theory*, 40(2):304–318, 1986.

[15] J. H. Dreze. Les fondemonts logiques des l'utilite cardinale et de la probabilité subjective. *La Décision.* Paris: Colloques Internationaux du CNRS, 1961.

[16] Juan Dubra, Fabio Maccheroni, and Efe A. Ok. Expected utility theory without the completeness axiom. *Cowles Foundation Discussion Papers 1294,* Cowles Foundation for Research in Economics, Yale University, 2001.

[17] S. Evans. Lecture notes from the 2010 probability at warwick workshop. Available at `www2.warwick.ac.uk/fac/sci/statistics/research/paw/paw2010`.

[18] P.C. Fishburn. *Utility theory for decision making.* Publications in operations research. R. E. Krieger Pub. Co., 1979.

[19] P.C. Fishburn. Multilinear expected utility. *Mathematics of Operations Research*, 5(4):pp. 502–509, 1980.

[20] P.C. Fishburn. An axiomatic characterization of skew-symmetric bilinear functionals, with applications to utility theory. *Economics Letters*, 8(4):311–313, 1981.

[21] I. Gilboa. *Theory of decision under uncertainty.* Econometric Society monographs. Cambridge University Press, 2009.

[22] I. Gilboa and D. Schmeidler. *A theory of case-based decisions.* Cambridge University Press, 2001.

[23] I. Gilboa and D. Schmeidler. A derivation of expected utility maximization in the context of a game. *Games and Economic Behavior*, 44(1):172 – 182, 2003.

[24] Itzhak Gilboa and David Schmeidler. Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639, 1995.

[25] V. Guillemin and A. Pollack. *Differential topology.* Mathematics Series. Prentice-Hall, 1974.

[26] Paul R. Halmos. *Finite-dimensional vector spaces.* Springer-Verlag, New York, 1974.

[27] Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988. 10.1007/BF00129168.

[28] Peter J. Hammond. Consistent plans, consequentialism, and expected utility. *Econometrica*, 57(6):1445–1449, 1989.

[29] P.J Hammond. Ex-ante and ex-post welfare optimality under uncertainty. *Economica*, 48(191):235–250, 1981.

[30] I. N. Herstein and J. Milnor. An axiomatic approach to measurable utility. *Econometrica*, 21(2):pp. 291–297, 1953.

[31] W. Hurewicz and H. Wallman. *Dimension theory.* Princeton mathematical series. Princeton University Press, revised 1948 edition, 1941.

[32] E. Karni and Z. Safra. An extension of a theorem of von Neumann and Morgenstern with an application to social choice theory. *Journal of Mathematical Economics*, 34(3):315–327, 2000.

[33] Edi Karni. A theory of medical decision making under uncertainty. *Journal of Risk and Uncertainty*, 39(1):1–16, 2009.

[34] Veronika Kobberling and Peter P. Wakker. Preference foundations for nonexpected utility: A generalized and simplified technique. *Mathematics of Operations Research*, 28(3):395–423, 2003.

[35] D.M. Kreps. *Notes on the theory of choice.* Underground classics in economics. Westview Press, 1988.

[36] K. Kuratowski. *Topology.* Number v. 1 in Topology. Academic Press, 1968.

[37] K. Kuratowski. *Topology.* Number v. 2 in Topology. Academic Press, 1968.

[38] G. Loomes and R. Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824, 1982.

[39] A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic theory.* Oxford University Press, 1995.

[40] Philippe Mongin. "notes and comments: A note on mixture sets in decision theory". *Decisions in Economics and Finance*, 24(1):59–69, 2001.

[41] S Morris and T Ui. Best response equivalence. *Games and Economic Behavior*, 49(2):260–287, 2004.

[42] J.R. Munkres. *Topology.* Prentice Hall, 2nd edition, 2000.

[43] Patrick H.D. O'Callaghan. Context and decision: a geometric approach. *University of Warwick Thesis,* 2nd chapter, 2011.

[44] P. Orlik and Conference Board of the Mathematical Sciences. *Introduction to arrangements*. Regional conference series in mathematics. Published for the Conference Board of the Mathematical Sciences by the American Mathematical Society, 1989.

[45] Kevin W. S. Roberts. Interpersonal comparability and social choice theory. *The Review of Economic Studies*, 47(2):421–439, 1980.

[46] A. Rubinstein. *Economics and Language*. Cambridge Univ. Press, 2000.

[47] L.J. Savage. *The foundations of statistics*. Dover Publications, 1972.

[48] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, 57(3):571–587, 1989.

[49] E. Shafir, I. Simonson, and A. Tversky. Reason-based choice* 1. *Cognition*, 49(1-2):11–36, 1993.

[50] Lynn A. Steen and J. Arthur Seebach. *Counterexamples in Topology*. Hot, Rinehart and Winston, Inc., USA, 1970.

[51] G. Strang. *Linear algebra and its applications*. Harcourt, Brace, Jovanovich, Publishers, 1988.

[52] K. Vind. Independent preferences. *Journal of Mathematical Economics*, 20(1):119–135, 1991.

[53] Karl Vind and Birgit Grodal. *Independence, Additivity, Uncertainty*. Springer, http://springer.de, 2003.

[54] J. von Neumman and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, Princeton and Oxford, 1944.

[55] R. Webster. *Convexity*. Oxford science publications. Oxford University Press, 1994.

[56] J. Zabczyk. *Chance and decision: stochastic control in discrete time*. Scuola Normale Superiore, 1996.

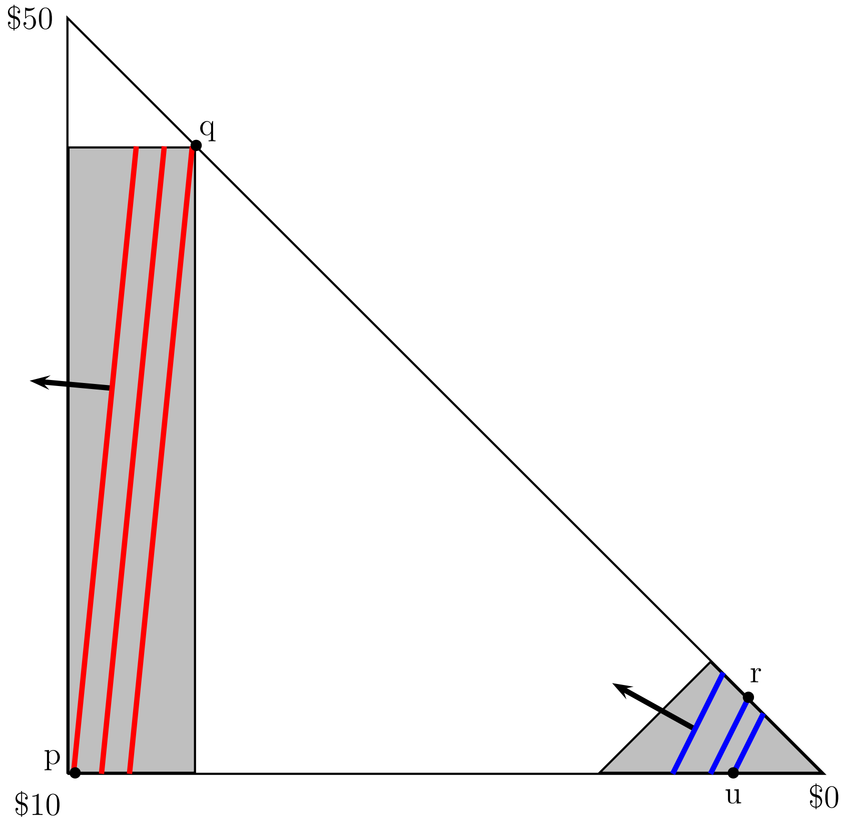

Figure 1: The Allais paradox and incomplete preferences. The decision-maker's preferences are such that the slope of the indifference curves in the triangular region containing δ_0 are shallower than those in the rectangular region containing p and q. If preferences satisfy the conditions for the main theorem of this paper over the union of these two regions, then the decision maker has a mixture preserving utility representation on this union. The representation is the restriction of two distinct expected utility functions to the respective regions. Since every element in the rectangle dominates every element in the triangle, preferences are lexicographic across regions and so comparing the utility value of an element in the rectangle with that of an element in the triangle is meaningless. Only ordinal statements are meaningful. By contrast, within each region, the representation is numerically meaningful in the sense of vNM and HM53.