A Foundation for Expected Utility in Decision Problems and Games*

Andrés Perea†
Maastricht University

This version: April 2021

Abstract

In a decision problem or game we typically fix the person's utilities but not his beliefs. What, then, do these utilities represent? To explore this question we assume that the decision maker holds a *conditional preference relation* – a mapping that assigns to every possible probabilistic belief a preference relation over his choices. We impose a list of axioms on such conditional preference relations, and show that they single out those conditional preference relations that admit an expected utility representation. If there are no weakly dominated choices, the key property is the *existence of a uniform preference increase,* which states that the decision maker should be able to uniformly increase the preference intensity for one of his choices without contradicting the conditional preference relation. In the presence of weakly dominated choices this condition is strengthened to the *existence of coherent uniform preference increases.* We also present a procedure that can be used to construct, for a given conditional preference relation satisfying the axioms, a utility function that represents it. If there are no weakly dominated choices, the existence of a uniform preference increase can be replaced by two easily verifiable conditions: *strong transitivity* and the *line property.*

JEL Classification: C72, D81

Keywords: Expected utility, decision problems, games, conditional preference relation, uniform preference increase, weak dominance.

*A special word of gratitude goes to Emiliano Catonini, Stephan Jagau and Peter Wakker for their very extensive and detailed feedback on this paper. I would also like to thank Geir Asheim, Christian Bach, Rubén Becerril, Giacomo Bonanno, Richard Bradley, Gabriel Frahm, Shuige Liu, Andrew Mackenzie, Martin Meier, Rineke Verbrugge and Marco Zaffalon for their useful suggestions and comments. Finally, I am grateful to the seminar audiences at the University of Glasgow, Maastricht University and the One World Mathematical Game Theory Seminar for their valuable feedback.

†*Address:* Dept. of Quantitative Economics and EpiCenter, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. *E-mail:* a.perea@maastrichtuniversity.nl *Web:* http://www.epicenter.name/Perea/

1 Introduction

What do the utilities in a decision problem or game represent? This is the key question we wish to explore in this paper. It is often argued that such utilities may be derived from Savage's (1954) framework. Indeed, from a particular player's point of view in a game, his opponents' choice combinations may be viewed as the set of states about which this player is uncertain, whereas his own choices correspond to acts that assign to every state some consequence. In that sense, a game can be embedded into Savage's model. The framework by Savage provides an axiomatic foundation for subjective expected utility maximization, by imposing axioms on the decision maker's preference relation over acts, and showing that the preference relations satisfying these axioms are precisely those that admit an expected utility representation. Such an expected utility representation consists of a subjective probabilistic belief about the set of states, together with a utility function assigning to every possible consequence some utility. One could thus argue that the utilities in a game, or a decision problem in general, may be derived from the player's preferences over acts, provided they satisfy the axioms as proposed by Savage.

In my view there are at least two problems with this approach. First, Savage assumes that the decision maker holds preferences over *all possible acts,* that is, over all possible functions from the set of states to the set of consequences. In a decision problem or game, however, many of these acts will not correspond to choices, and will therefore be unrelated to the decision problem or game at hand. It thus seems problematic to assume that the decision maker holds preferences even over these acts.

A second problem is that the axioms provided by Savage yield a *unique* subjective probabilistic belief for the decision maker about the set of states. In a game, therefore, these axioms lead, for a given player, to a *unique* probabilistic belief about the opponents' choices. At the same time, most game theory concepts select, for every player, *several* possible beliefs. Consider, for instance, the concepts of Nash equilibrium (Nash (1950, 1951)) and rationalizability (Bernheim (1984) and Pearce (1984)). In the spirit of Aumann and Brandenburger (1995), a mixed strategy in a Nash equilibrium may be interpreted as the belief that the other players have about this player's pure strategy. As a game typically has several Nash equilibria, the concept selects several possible beliefs for the same player. Analogously, a game typically has several rationalizable pure strategies for the same player. Therefore, also rationalizability typically selects multiple beliefs for a given player in the game. But also in one-person decision problems it may be natural to allow for several different beliefs. Consider, for instance, a decision problem where the consequence of a choice depends on the state of the weather. Then, we may naturally be interested in how the decision maker would rank his choices under several different weather forecasts.

Despite the multiplicity of beliefs, decision problems and games typically view the utility function of a person as given. This, however, seems to be at odds with Savage's model, where the axioms on the decision maker's preferences over acts do not only lead to a utility function which is unique up to positive affine transformations, but also to a unique belief. What does it mean, then, that a decision problem or game specifies the person's utilities but not his belief?

As a possible answer to this question, this paper adopts a decision theoretic view on games which resembles Gilboa and Schmeidler's (2003), and which is fundamentally different from Savage (1954). Instead of assuming that a person holds preferences over all possible acts that can be derived from the decision problem or game, we suppose that the person's probabilistic belief is variable, and that he holds, for every possible belief, a preference relation over his own choices in the decision problem or game. The primitive object in our setup is thus a mapping which assigns to every probabilistic belief about the states a preference relation over his own choices. Such mappings are called *conditional preference relations,* and these are precisely the mappings used by Gilboa and Schmeidler (2003) for their foundation of expected utility in games. By adopting this approach we thus no longer fix the probabilistic belief of a decision maker, yet at the same time we make sure that the preferences of a decision maker only concern those acts that correspond to his actual choices in the decision problem or game.

We then ask: When does such a conditional preference relation have an expected utility representation? In other words, when can we find a utility function, assigning a utility index to every combination of a choice and a state, such that for every belief p and every two choices a and b, the decision maker prefers a to b exactly when the expected utility of a under p is larger than that of b under p ? We impose six axioms on conditional preference relations, and prove in Theorem 5.1 that the conditional preference relations satisfying the axioms are precisely those that admit an expected utility representation. Importantly, the proof of Theorem 5.1 is constructive and procedural: For a given conditional preference relation satisfying the axioms, we explictly show *how* to construct a utility function that represents it, by means of an easy and intuitive procedure.

On the road towards this theorem we zoom in on two special cases of conditional preference relations: The case of two choices and the case where there are no weakly dominated choices. For the first case we show in Theorem 3.1 that the basic regularity axioms of *completeness, transitivity, continuity*, *preservation of indifference* and *preservation of strict preference*, which also appear in Gilboa and Schmeidler (2003), are sufficient to characterize the conditional preference relations having an expected utility representation.

If there are no weakly dominated choices, but possibly more than two choices in total, then Theorem 4.1 shows that expected utility can be characterized by the regularity axioms together with a new axiom, the *existence of a uniform preference increase.* The latter axiom states that for at least one choice there must be an alternative conditional preference relation that uniformly increases the preference intensity for this choice by a constant degree.

For the general case, the conditional preference relations having an expected utility representation are precisely those that satisfy the regularity axioms together with a strengthening of the axiom above, called the *existence of coherent uniform preference increases.* This is the content of Theorem 5.1. Existence of coherent uniform preference increases states that for every choice there must be an alternative conditional preference relation that uniformly increases the preference intensity for this choice by a constant degree, and such that this collection of uniform preference increases induces at every belief a preference intensity between any two choices that never contradicts the original conditional preference relation.

An important feature of an expected utility representation in our setting is that the *same* utility function is used to represent the decision maker's preferences for *all* possible beliefs. This reflects the idea that the beliefs of the decision maker are often prone to change, due to reasoning or when new information is received, whereas the decision maker's tastes are generally viewed as more robust. When we write down a utility matrix, we thus assume that these various possible beliefs lead to preferences that are "consistent" with one another, in the sense that they are all governed by the same utilities. The axiom system in this paper reveals what is needed on behalf of the conditional preference relation to achieve such consistency.

Gilboa and Schmeidler (2003) also propose axioms on conditional preference relations. Contrary to the present paper, their axioms are sufficient but not necessary for guaranteeing an expected utility representation. More precisely, their axioms characterize those conditional preference relations that can be represented by a special type of utility function which they call *diversified.* The key axiom in their system is *diversity,* which states that for every strict ordering of at most four choices, there must be a belief that induces precisely that ordering. The diversity axiom, however, rules out many cases of interest, such as all scenarios with two states and at least three choices, all scenarios with three states and at least four choices, and all cases where there are weakly dominated choices. The axiom system we adopt allows for these scenarios, and does not impose any restriction on the utility function that can be used to represent the conditional preference relation at hand.

For the case of three or more choices, this paper shows that either the existence of one uniform preference increase, or the existence of a coherent collection of uniform preference increases, is required to make an expected utility representation possible. In fact, adding this condition to the regularity axioms is precisely what is needed. But what does this axiom entail intuitively?

Let us start with the existence of a (single) uniform preference increase. It requires that for at least one choice a, the decision maker should envisage a new, hypothetical conditional preference relation in which the attractiveness of a, compared to the other choices, is uniformly lifted by a constant degree. Then, this constant degree can be used as a "common scale" on which the preference intensities between a and every other choice, at each of the possible beliefs, can be expressed. For this common scale to be compatible with the original conditional preference relation, the beliefs at which the decision maker is indifferent between two choices b and c other than a must then be exactly the beliefs at which the induced preference intensity between a and b is the same as between a and c. Thus, the axiom intuitively states that it must be possible for the decision maker to uniformly lift the preference intensity for a particular choice by a constant amount, such that the induced preference intensities between the other choices never contradict the original conditional preference relation.

The existence of coherent uniform preference increases imposes more: First of all, it requires the decision maker to envisage, for *each* of the possible choices, a new, hypothetical conditional preference relation that uniformly lifts the preference intensity for this choice by a constant degree. We thus obtain a whole system of uniform preference increases. At a given belief, and for a given pair of choices b and c, the preference intensity between b and c can typically be derived from different chains of uniform preference increases in this system. For this system of uniform preference increases to be coherent, the preference intensities induced by all these chains must never contradict the original conditional preference relation. This is the intuitive content of the stronger axiom "existence of coherent uniform preference increases".

The case in which there are no weakly dominated choices plays a prominent role in this paper. First, it may be viewed as a canonical case for a rational decision maker. To see this, note that a rational decision maker may reasonably be expected not to go for any of the weakly dominated choices. Indeed, for a weakly dominated choice there will always be another choice that is weakly preferred at all states, and strictly preferred at some state. But if we eliminate all weakly dominated choices from the problem, then we will be left with a collection of choices that do not weakly dominate one another.

We have seen above that in this case, the regularity axioms together with the existence of a (single) uniform preference increase are both necessary and sufficient for guaranteeing an expected utility representation. That is, the stronger axiom "existence of coherent uniform preference increases" is not needed. If, in addition, the sets of beliefs for which the decision maker is indifferent between two choices do not all coincide, then Theorem 4.2 states that the utility differences are uniquely determined up to a positive multiplicative constant. In that case, the utility differences between two choices a and b may be viewed as expressing the decision maker's "preference intensity" between a and b. That is, the utility difference between two choices does not only indicate which choice is preferred by the decision maker, but also by "how much". This is similar to the approaches by Anscombe and Aumann (1963) and Wakker (1989), where the axioms of state independence and state independent preference intensity, respectively, guarantee that the utility difference between two consequences is the same at every state, and may be viewed as expressing the "preference intensity" between these consequences.

Moreover, in the absence of weakly dominated choices we are able to provide an alternative characterization of expected utility, in Theorem 6.1, by relying on two conditions, *strong transitivity* and the *line property.* Strong transitivity states that for every three choices a, b and c, the linear extensions of the sets of beliefs where the decision maker is indifferent between a and b, between b and c, and between a and c, respectively, must have a common intersection, possibly outside the belief simplex. The second condition states that for four choices a, b, c and d and a line L, if we know for which beliefs on the line L the decision maker is indifferent between e and f, for any $\{e, f\} \neq \{a, b\}$, then we also know for which belief on the line he will be indifferent between a and b. The advantage of these conditions is that they do not require us to construct alternative conditional preference relations, but purely relate to the original conditional preference relation, making them directly and easily verifiable. The drawback is that the intuitive content of these properties is not as obvious, and it is not clear at this moment how these properties could be extended to allow for weakly dominated choices.

This paper is organized as follows. In Section 2 we present the necessary mathematical definitions. In Section 3 we introduce the notion of a conditional preference relation, present the regularity axioms, and show that for the case of two choices these are necessary and sufficient for an expected utility representation. In Section 4 we discuss the case where there are no weakly dominated choices, whereas Section 5 treats the general case. In Section 6 we zoom in on strong transitivity and the line property. We conclude with a discussion in Section 7. All the proofs can be found in the appendix and the online appendix.

2 Mathematical Definitions

In this section we introduce the mathematical definitions and notation needed for this paper, mainly from linear algebra. For a finite set X, we denote by $\mathbf{R}^X$ the set of all functions $v : X \to \mathbf{R}$. Scalar multiplication and addition on $\mathbf{R}^X$ are defined in the usual way: For a function $v \in \mathbf{R}^X$ and a number $\lambda \in \mathbf{R}$, the function $\lambda \cdot v$ is given by $(\lambda \cdot v)(x) = \lambda \cdot v(x)$ for all $x \in X$. Similarly, for functions $v, w \in \mathbf{R}^X$, the sum $v + w$ is given by $(v + w)(x) = v(x) + w(x)$ for all $x \in X$. The set $\mathbf{R}^X$ together with these two operations constitutes a *linear space,* and elements in $\mathbf{R}^X$ are called *vectors.* By $\underline{0}$ we denote the vector in $\mathbf{R}^X$ where $\underline{0}(x) = 0$ for all $x \in X$. For two subsets $V, W \subseteq \mathbf{R}^X$ and numbers $\alpha, \beta \in \mathbf{R}$, we define the set

$$\alpha V + \beta W := \{\alpha v + \beta w \mid v \in V \text{ and } w \in W\}.$$

For every two vectors $v, w \in \mathbf{R}^X$, the *vector product* is given by $v \cdot w := \sum_{x \in X} v(x)w(x)$.

A subset $V \subseteq \mathbf{R}^X$ is called a *linear subspace* of $\mathbf{R}^X$ if for every $v, w \in V$ and every $\alpha, \beta \in \mathbf{R}$, we have that $\alpha v + \beta w \in V$. For a subset $V \subseteq \mathbf{R}^X$, we denote by

$$\langle V \rangle := \{\sum_{k=1}^{K} \alpha_k v_k \mid K \geq 1,\ \alpha_k \in \mathbf{R} \text{ and } v_k \in V \text{ for all } k \in \{1, ..., K\}\}$$

the set of all (finite) *linear combinations* of elements in V, and call it the (*linear) span* of V. Here, $\sum_{k=1}^{K} \alpha_k v_k$ is called a *linear combination* of the vectors $v_1, ..., v_K$. The span $\langle V \rangle$ is always a linear subspace, and if V itself is a linear subspace then $\langle V \rangle = V$. Vectors $v_1, ..., v_K \in \mathbf{R}^X$ are called *linearly independent* if none of the vectors is a linear combination of the other vectors. Consider a linear subspace V of $\mathbf{R}^X$, and vectors $v_1, ..., v_K \in V$. The set of vectors $\{v_1, ..., v_K\}$ is a *basis* for V if $v_1, ..., v_K$ are linearly independent, and $\langle \{v_1, ..., v_K\} \rangle = V$. Every basis for V has the same number of vectors, and this number is called the *dimension* of V, denoted by $\dim(V)$. If $V = \{\underline{0}\}$, then $\dim(V) = 0$.

For a subset $V \subseteq \mathbf{R}^X$, we denote by

$$conv(V) := \{\sum_{k=1}^{K} \alpha_k v_k \mid K \geq 1,\ \alpha_k \in \mathbf{R},\ \alpha_k \geq 0,\ v_k \in V \text{ for all } k \in \{1, ..., K\}$$
$$\text{and } \sum_{k=1}^{K} \alpha_k = 1\}$$

the set of all (finite) *convex combinations* of elements in V, and call it the *convex hull* of V.

A *hyperplane* is a set of the form $H = \{v \in \mathbf{R}^X \mid v \cdot w = c\}$, where $w \in \mathbf{R}^X \backslash \{\underline{0}\}$ and $c \in \mathbf{R}$. If $c = 0$ then H is a linear subspace of dimension $|X| - 1$, where $|X|$ denotes the number of elements in X. Two hyperplanes H and H' are *parallel* if there is some $v \in \mathbf{R}^X$ such that $H' = H + \{v\}$. In that case, there is some $w \in \mathbf{R}^X \backslash \{\underline{0}\}$ and $c, c' \in \mathbf{R}$ such that $H = \{v \in \mathbf{R}^X \mid v \cdot w = c\}$ and $H' = \{v \in \mathbf{R}^X \mid v \cdot w = c'\}$.

A mapping $f : \mathbf{R}^X \to \mathbf{R}$ is *linear* if for every $v, w \in \mathbf{R}^X$ and every $\alpha, \beta \in \mathbf{R}$ it holds that $f(\alpha v + \beta w) = \alpha f(v) + \beta f(w)$. A mapping $f : \mathbf{R}^X \to \mathbf{R}$ is *affine* if for every $v, w \in \mathbf{R}^X$ and every $\alpha \in \mathbf{R}$ it holds that $f(\alpha v + (1 - \alpha)w) = \alpha f(v) + (1 - \alpha) f(w)$.

A *probability distribution* on X is a vector $p \in \mathbf{R}^X$ such that $\sum_{x \in X} p(x) = 1$ and $p(x) \geq 0$ for all $x \in X$. The set of probability distributions on X is denoted by $\Delta(X)$. For a given element $x \in X$, we denote by $[x]$ the probability distribution in $\Delta(X)$ where $[x](x) = 1$ and $[x](y) = 0$ for all $y \in X \backslash \{x\}$. A probability distribution p has *full support* if $p(x) > 0$ for all $x \in X$.

A (*directed*) *graph* is a pair (V, E) where V is a set of vertices, and $E \subseteq V \times V$ the set of (directed) edges. We assume that $(v, v) \notin E$ for every $v \in V$. A *path* from $v \in V$ to $w \in V$ is a sequence $(v_1, v_2, ..., v_K)$ where $v_1 = v$, $v_K = w$ and $(v_k, v_{k+1}) \in E$ for all $k \in \{1, ..., K - 1\}$. Similarly, $(v_1, v_2, ..., v_K)$ is an *undirected path* from v to w if for every $k \in \{1, ..., K - 1\}$ either $(v_k, v_{k+1}) \in E$ or $(v_{k+1}, v_k) \in E$. A *tree* is a graph (V, E) with a vertex $r \in V$, the root, such that for every $v \in V \backslash \{r\}$ there is a unique path from r to v. Within a tree T with root r, and for a given vertes v, we denote by *depth*(v) the number of edges on the unique path from r to v. By $depth(T)$ we denote the maximal depth within T. For a given vertex $v \neq r$, the *predecessor* of v is the vertex that comes before v in the unique path from r to v. A graph (V, E') is a *spanning tree* for (V, E) if (V, E') is a tree and $E' \subseteq E$. For a graph (V, E), a subgraph (V', E') with $V' \subseteq V$ and $E' \subseteq E$ is an *undirected connected component* if (i) $E' = E \cap (V' \times V')$, (ii) for every $v, w \in V'$ there is an undirected path in (V, E) from v to w, and (iii) for every $v \in V', w \in V \backslash V'$ there is no undirected path in (V, E) from v to w.

3 Case of Two Choices

In this section we formally introduce a conditional preference relation as the primitive notion of our model. Subsequently, we impose some regularity axioms on conditional preference relations, and show that for the case of two choices these suffice to single out the conditional preference relations that admit an expected utility representation.

3.1 Conditional Preference Relations

In line with Gilboa and Schmeidler (2003), the primitive object in this paper is that of a *conditional preference relation* – a mapping that assigns to every probabilistic belief over the states a preference relation over the available choices. In accordance with the literature on decision making under uncertainty, we refer to such choices as *acts.* Consider a decision maker (DM) who must choose from a finite set of acts A. The final outcome depends not only on the act $a \in A$, but also on the realization of a state $s \in S$ from a finite set of states S. We assume that the decision maker first forms a probabilistic belief p on S, which then induces a preference relation $\succsim_p$ on A. Formally, a *preference relation* $\succsim_p$ on A is a binary relation $\succsim_p \subseteq A \times A$. If $(a, b) \in \succsim_p$ we write $a \succsim_p b$, and the interpretation is that the DM weakly prefers act a to act b if his belief is p.

Definition 3.1 (Conditional preference relation) *Consider a finite set of acts A and a finite set of states S. A conditional preference relation on (A, S) is a mapping $\succsim$ that assigns to every probabilistic belief $p \in \Delta(S)$ a preference relation $\succsim_p$ on A.*

In a game, the DM would be one of the players, A would be his set of actions in the game, and S the set of opponents' action combinations. For two acts a and b, we write that $a \sim_p b$ if $a \succsim_p b$ and $b \succsim_p a$. The interpretation is that the DM is indifferent between a and b while having the belief p. Similarly, we write $a \succ_p b$ if $a \succsim_p b$ but not $b \succsim_p a$, representing a case where the DM strictly prefers a to b. For two acts $a, b \in A$ we define the sets of beliefs $P_{a\sim b} := \{p \in \Delta(S) \mid a \sim_p b\}$, $P_{a\succ b} := \{p \in \Delta(S) \mid a \succ_p b\}$ and $P_{a\succsim b} := \{p \in \Delta(S) \mid a \succsim_p b\}$. Similarly, we define the sets of states $S_{a\sim b} := \{s \in S \mid a \sim_{[s]} b\}$, $S_{a\succ b} := \{s \in S \mid a \succ_{[s]} b\}$ and $S_{a\succsim b} := \{s \in S \mid a \succsim_{[s]} b\}$. We say that (a) *$a$ strictly dominates b* under $\succsim$ if $a \succ_p b$ for all $p \in \Delta(S)$; (b) *a weakly dominates* b under $\succsim$ if $a \succsim_p b$ for all $p \in \Delta(S)$, and $a \succ_p b$ for at least one $p \in \Delta(S)$; (c) a is *equivalent* to b under $\succsim$ if $a \sim_p b$ for all $p \in \Delta(S)$; and (d) $\succsim$ has *preference reversals* on $\{a, b\}$ if there is a belief p with $a \succ_p b$ and another belief q with $b \succ_q a$. Hence, $\succsim$ either exhibits weak dominance, equivalence, or preference reversals on $\{a, b\}$.

In the remainder of this paper we will assume that the conditional preference relation does not have equivalent acts. That is, $\succsim$ either exhibits weak dominance or preference reversals on every pair of acts $\{a, b\}$. In the discussion section we will briefly explain how our analysis can easily be extended to cover equivalent acts.

A conditional preference relation is said to have an *expected utility representation* if there is a utility function, assigning a utility index to every act-state pair (a, s), such that for every possible belief, the DM prefers act a to act b precisely when his expected utility from a is higher than that from b under the belief at hand.

Definition 3.2 (Expected-utility representation) *A conditional preference relation $\succsim$ has an expected utility representation if there is a utility function $u : A \times S \to \mathbf{R}$ such that for every belief $p \in \Delta(S)$ and every two acts $a, b \in A$,*

$$a \succsim_p b \text{ if and only if } \sum_{s \in S} p(s) \cdot u(a, s) \geq \sum_{s \in S} p(s) \cdot u(b, s).$$

In this case, we say that the conditional preference relation $\succsim$ is *represented* by the utility function u. For a given vector $v \in \mathbf{R}^S$ we use the notation $u(a, v) := \sum_{s \in S} v(s) \cdot u(a, s)$. Hence, the condition above can be written as $a \succsim_p b$ if and only if $u(a, p) \geq u(b, p)$.

3.2 Regularity Axioms

We will now impose some very basic axioms on conditional preference relations, to which we refer as *regularity axioms.*

Axiom 3.1 (Completeness) *For every belief p and any two acts $a, b \in A$, either $a \succsim_p b$ or $b \succsim_p a$.*

Axiom 3.2 (Transitivity) *For every belief p and every three acts $a, b, c \in A$ with $a \succsim_p b$ and $b \succsim_p c$, it holds that $a \succsim_p c$.*

Axiom 3.3 (Continuity) *For every two different acts $a, b \in A$ and every two beliefs $p \in P_{a \succ b}$ and $q \in P_{b \succ a}$, there is some $\lambda \in (0, 1)$ such that $(1 - \lambda)p + \lambda q \in P_{a \sim b}$.*

Axiom 3.4 (Preservation of indifference) *For every two different acts $a, b \in A$ and every two beliefs $p \in P_{a \sim b}$ and $q \in P_{a \sim b}$, we have that $(1 - \lambda)p + \lambda q \in P_{a \sim b}$ for all $\lambda \in (0, 1)$.*

Axiom 3.5 (Preservation of strict preference) *For every two different acts $a, b \in A$ and every two beliefs $p \in P_{a \succsim b}$ and $q \in P_{a \succ b}$, we have that $(1 - \lambda)p + \lambda q \in P_{a \succ b}$ for all $\lambda \in (0, 1)$.*

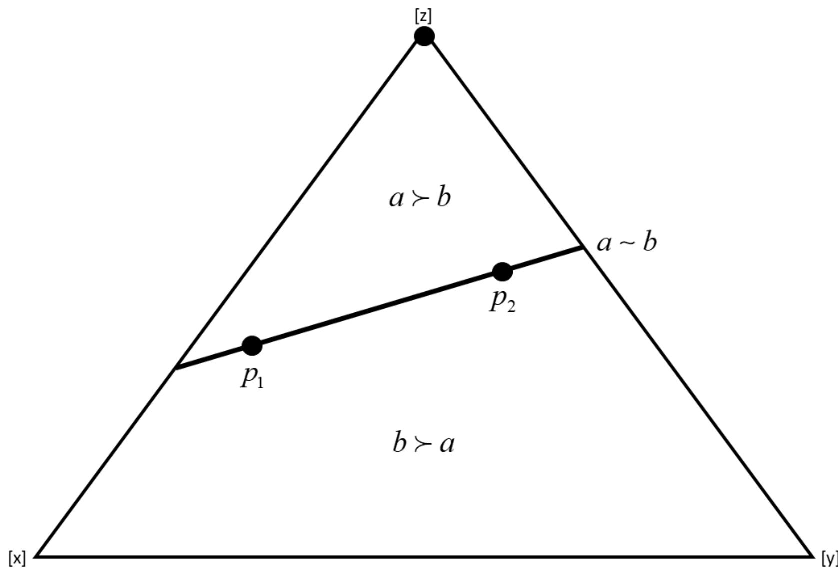

Figure 1: A typical regular conditional preference relation

Completeness and transitivity together resemble the *ranking* axiom in Gilboa and Scmeidler (2003). Our definition of continuity is formally different from Gilboa and Schmeidler's (2003) version, but reveals the same idea. When taken together, our axioms of preservation of indifference and preservation of strict preference correspond precisely to Gilboa and Schmeidler's (2003) axiom of *combination.*

In the remainder of the paper, whenever we say that a conditional preference relation is regular, or satisfies the regularity axioms, we mean that it satisfies completeness, transitivity, continuity, preservation of indifference and preservation of strict preference. See Figure 1 for a typical regular conditional preference relation $\succsim$ with two acts a and b, and three states x, y and z. The area within the triangle represents the set $\Delta(S)$ of all probabilistic beliefs on $S = \{x, y, z\}$, with the probability 1 beliefs $[x], [y]$ and $[z]$ as the extreme points. The two-dimensional plane represents all the vectors in $\mathbf{R}^S$ where the sum of the coordinates is 1, containing the belief simplex $\Delta(S)$ as a subset. Hence, $a \sim_p b$ for all beliefs p on the line segment, $a \succ_p b$ for all beliefs p above the line segment, and $b \succ_p a$ for all beliefs p below the line segment. It may be verified that $\succsim$ satisfies all the regularity axioms.

The following theorem shows that a conditional preference relation on *two acts* has an expected utility representation precisely when it satisfies the regularity axioms.

Theorem 3.1 (Two choices) *Consider a set A consisting of two acts, a finite set of states S, and a conditional preference relation $\succsim$ on (A, S). Then, $\succsim$ has an expected utility representation, if and only if, it satisfies completeness, transitivity, continuity, preservation of indifference and preservation of strict preference.*

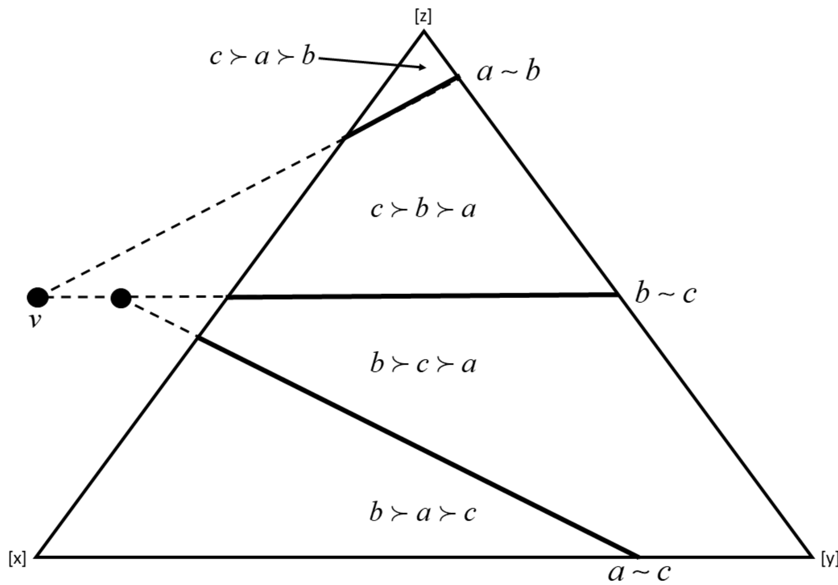

Figure 2: Regularity axioms are not sufficient for expected utility representation

In particular, the conditional preference relation $\succsim$ in Figure 1 has an expected utility representation. One way to generate a utility function u that represents $\succsim$ is as follows: Choose the utilities $u(a,x), u(a,y)$ and $u(a,z)$ arbitrarily. Then, choose the utilities $u(b,x), u(b,y)$ and $u(b,z)$ such that the expected utility for b at the beliefs p_1 and p_2 is equal to the expected utility for a at these beliefs, and such that $u(b,z) < u(a,z)$.

4 Case of No Weak Dominance

In this section we consider the case of more than two acts and show, by means of an example, that the regularity axioms are no longer sufficient to guarantee an expected utility representation. The reason for this failure is that, starting from this conditional preference relation, we cannot uniformly increase the preference intensity for any given act without contradicting the conditional preference relation. This leads to a new axiom, "existence of a uniform preference increase", which we formally present below. It is shown that in the absence of weakly dominated acts, this axiom, together with the regulariy axioms, characterizes precisely those conditional preference relations that have an expected utility representation. We finally prove that in the absence of weakly dominated acts, the utility differences are uniquely given up to a positive multiplicative constant, provided the sets of beliefs where the DM is indifference between two acts do not all coincide.

4.1 Why Regularity Axioms Are Not Sufficient

Consider the conditional preference relation $\succsim$ represented by Figure 2. It may be verified that $\succsim$ satisfies

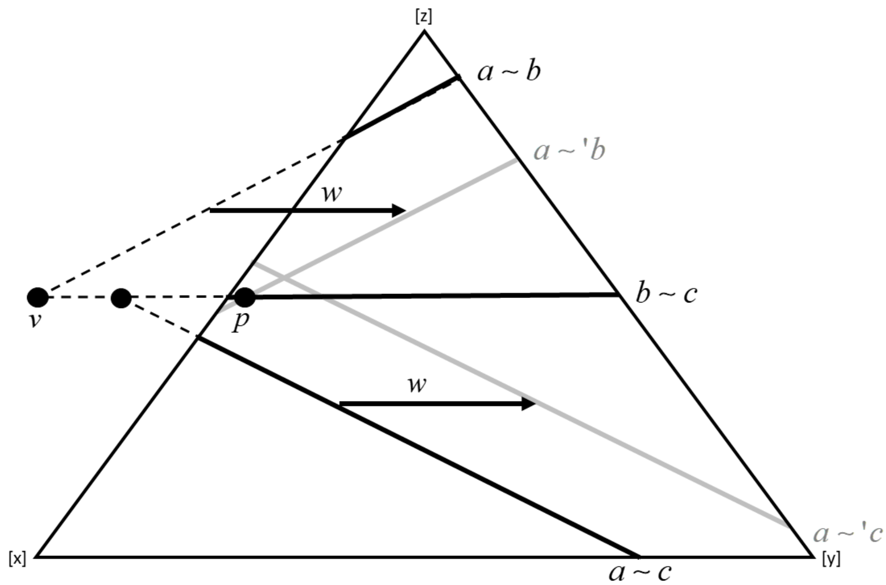

Figure 3: When there is no uniform preference increase

all the regularity axioms. Yet, there is no expected utility representation for $\succsim$. To see why, suppose there would be a utility function u that represents $\succsim$. Then, the induced expected utilities of a and b must be equal on the hyperplane $\langle P_{a\sim b}\rangle$, the expected utilities of b and c must be equal on the hyperplane $\langle P_{b\sim c}\rangle$ and the expected utilities of a and c must be equal on the hyperplane $\langle P_{a\sim c}\rangle$, also at vectors that lie outside the belief simplex. But then, the expected utilities of a and c must be the same at the vector v where $\langle P_{a\sim b}\rangle$ and $\langle P_{b\sim c}\rangle$ intersect, which is impossible since v does not belong to $\langle P_{a\sim c}\rangle$.

This raises the question: What is "wrong" with this conditional preference relation? As it turns out, we cannot uniformly increase the preference intensity for act a by a fixed degree without contradicting the conditional preference relation. To see this, suppose there would be an alternative conditional preference relation $\succsim'$ that uniformly increases the preference intensity for act a by a fixed degree, relative to $\succsim$. Then, the preference intensity between a and b and the preference intensity between a and c should both be raised by the same amount. The indifference set $P_{b\sim c}$ contains precisely those beliefs where the DM is indifferent between b and c. Hence, intuitively, these are precisely the beliefs where his preference intensity between a and b is equal to his preference intensity between a and c. If we move from one belief in $P_{b\sim c}$ to another belief in $P_{b\sim c}$, we thus increase, or decrease, the preference intensity between a and b and the preference intensity between a and c by the same amount. Therefore, the new indifference sets $P_{a\sim' b}$ and $P_{a\sim' c}$ must be obtained from the original indifference sets $P_{a\sim b}$ and $P_{a\sim c}$ by a common parallel shift w that moves from one point in $\langle P_{b\sim c}\rangle$ to another point in $\langle P_{b\sim c}\rangle$. See Figure 3 for an illustration. However, as can be seen from Figure 3, the resulting conditional preference relation $\succsim'$ is not transitive: At the belief p, the DM is indifferent between a and b, and indifferent between b and c, but not indifferent between a and c under $\succsim'$.

In fact, starting from the original conditional preference relation $\succsim$, there is no uniform preference increase for act a. The reason is that any uniform preference increase for a must result in shifting the original indifference sets $P_{a\sim b}$ and $P_{a\sim c}$ along a multiple of the vector w. Hence, if a uniform preference increase for a would exist then, by scaling this preference increase up or down by an appropriate amount, there should also be a uniform preference increase for a where $P_{a\sim' b}$ passes through the belief p in Figure 3. This, as we have seen, is impossible. By a similar reasoning, it can also be verified that there is no uniform preference increase for act b or for act c in this example.

As we will show, the absence of a uniform preference increase is precisely what prevents a regular conditional preference relation from having an expected utility representation, provided there are no weakly dominated acts. In the following subsection we formally define a uniform preference increase, and use it to introduce a new axiom, "existence of a uniform preference increase", which states that a uniform preference increase should exist for at least one of the acts.

4.2 The Axiom "Existence of Uniform Preference Increase"

Imagine the DM holds a conditional preference relation $\succsim$, and decides to uniformly increase his preference intensity for act a by a certain degree which we normalize to 1 That is, for every belief p, and relative to every other act b, the preference intensity for a is increased by 1. How would the new conditional preference relation $\succsim'$ compare to $\succsim$? Our arguments below will be based on two informal principles:

Principle 1: If we move from a belief p to a belief q on a line, then the "preference intensity" between a and b will change linearly.

Principle 2: The DM is indifferent between b and c precisely when the "intensity" by which he prefers a to b is equal to the "intensity" by which he prefers a to c.

The intensity by which the DM prefers a to b may also be negative, which means that he prefers b to a. Of course, "preference intensity" is not formally defined here, but it helps to motivate our new axiom.

Consider a belief $p_{ab} \in P_{a\sim b}$, a belief $p'_{ab} \in P_{a\sim' b}$ and some belief p such that $p = (1-\lambda)p_{ab} + \lambda p'_{ab}$. See Figure 4 for the case where $\lambda > 1$. Here, the numbers 1 and $\lambda - 1$ indicate the relative lengths of the corresponding line segments. Recall that the new conditional preference relation $\succsim'$ increases the preference intensity for a by 1, relative to $\succsim$. As $int_{b\succ' a}(p'_{ab}) = 0$, this implies that $int_{b\succ a}(p'_{ab}) = 1$, where $int_{b\succ' a}(p'_{ab})$ and $int_{b\succ a}(p'_{ab})$ denote the intensity by which the DM prefers b to a at the belief p'_{ab} under $\succsim'$ and $\succsim$, respectively. Since $p = (1-\lambda)p_{ab} + \lambda p'_{ab}$ and $int_{b\succ a}(p_{ab}) = 0$, it follows by principle 1 that $int_{b\succ a}(p) = \lambda$.

Now consider an alternative belief $q_{ab} \in P_{a\sim b}$ and some belief q' such that $p = (1-\lambda)q_{ab} + \lambda q'$. See Figure 4. As $int_{b\succ a}(q_{ab}) = 0$, it follows by principle 1 that

$$\lambda = int_{b\succ a}(p) = \lambda \cdot int_{b\succ a}(q'),$$

which implies that $int_{b\succ a}(q') = 1$, and hence $q' \in P_{a\sim' b}$. We thus see that, whenever

$$(1-\lambda)p_{ab} + \lambda p'_{ab} = (1-\lambda)q_{ab} + \lambda q'$$

for some $p_{ab}, q_{ab} \in P_{a\sim b}$ and $p'_{ab} \in P_{a\sim' b}$, then the belief q' must be in $P_{a\sim' b}$ as well. This means that the indifference sets $P_{a\sim b}$ and $P_{a\sim' b}$ are homothetic with respect to the belief p, and are thus parallel. See Figure

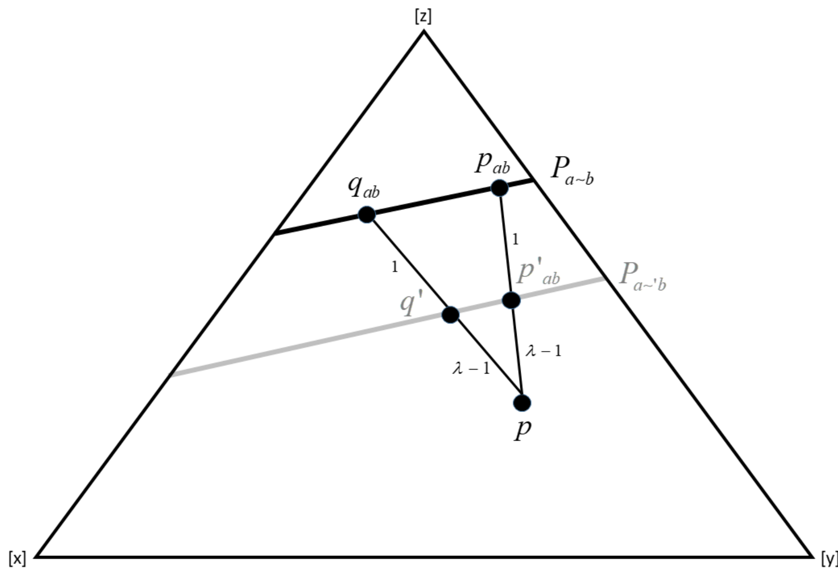

Figure 4: A uniform preference increase generates parallel indifference sets $P_{a\sim b}$ and $P_{a\sim' b}$

4. This bears some similarity with Burghart (2020) who defines, and studies, the notion of homotheticity for preference relations over lotteries. In particular, as $P_{a\sim b}$ and $P_{a\sim' b}$ are parallel, there is for every belief p a unique number λ with $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim' b}$. Since we have seen above that λ reflects the intensity by which the DM prefers b to a at the belief p, we refer to the number $int_{b\succ a}(p) := \lambda$ as the *derived preference intensity* between b and a at p.

Such a derived preference intensity can only exist if both sets $P_{a\sim b}$ and $P_{a\sim' b}$ are non-empty. Therefore, a must not weakly dominate b under $\succsim$, since otherwise a would strictly dominate b under $\succsim'$, and $P_{a\sim' b}$ would be empty. Moreover, b must not strictly dominate a under $\succsim$, since otherwise $P_{a\sim b}$ would be empty. This leads to the definition below.

Definition 4.1 (Domination graph) *The domination graph $DG[\succsim]$ is the directed graph where the vertices are the acts in A, and where the edge (a, b) is present precisely when a does not weakly dominate b, and b does not strictly dominate a.*

Hence, the edge (a, b) is present in $DG[\succsim]$ precisely when the preference intensity between b and a at the various beliefs can be derived from increasing the preference intensity of a in the way illustrated above.

Consider now a third act c, and assume that also (a, c) is in $DG[\succsim]$. Then, at a given belief p the preference intensity between c and a is given by $int_{c\succ a}(p) = \mu$, where μ is the unique number such that $p \in (1-\mu)P_{a\sim c} + \mu P_{a\sim' c}$. By principle 2 it then follows that $p \in P_{b\sim c}$ precisely when $int_{b\succ a}(p) = int_{c\succ a}(p)$. See Figure 5 where $int_{b\succ a}(p) = int_{c\succ a}(p) = \lambda$. This leads to the following definition of a uniform preference increase.

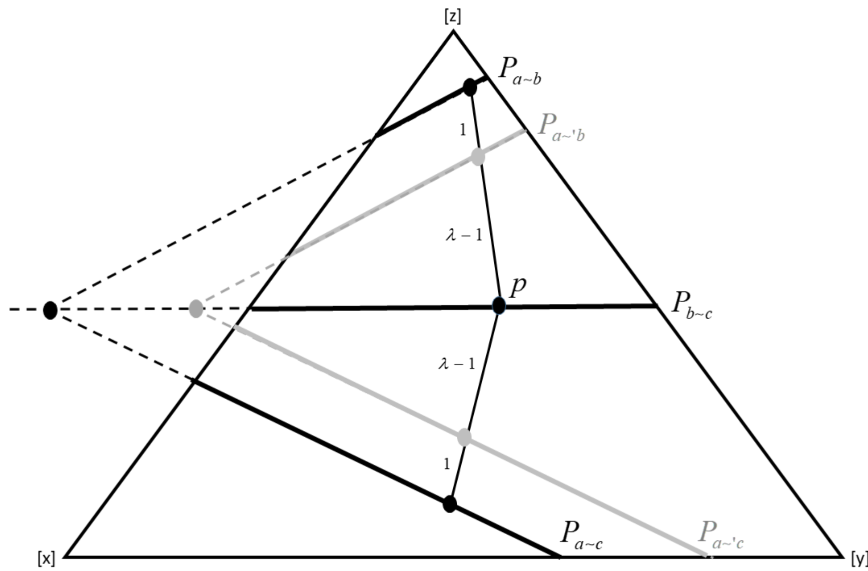

Figure 5: A uniform preference increase for a with three choices

Definition 4.2 (Uniform preference increase) *Consider a conditional preference relation $\succsim$, an act a, and an alternative conditional preference relation $\succsim'$. Then, $\succsim'$ uniformly increases the preference for a relative to $\succsim$ if*

(a) $\succsim'$ is regular and $P_{a\succsim b} \subseteq P_{a\succ' b}$ for all acts $b \neq a$;

(b) for every act $b \neq a$ with $(a,b) \in DG[\succsim]$, the conditional preference relation $\succsim'$ has preference reversals on $\{a,b\}$, and for every belief p there is a unique number $int_{b\succ a}(p)$ such that $p \in (1 - int_{b\succ a}(p)) \cdot P_{a\sim b} + int_{b\succ a}(p) \cdot P_{a\sim' b}$; and

(c) for every two acts $b, c \neq a$ the conditional preference $\succsim'$ coincides with $\succsim$ on $\{b,c\}$, and if $(a,b),(a,c) \in DG[\succsim]$ then

$$p \in P_{b\sim c} \text{ if and only if } int_{b\succ a}(p) = int_{c\succ a}(p).$$

Here, the condition $P_{a\succsim b} \subseteq P_{a\succ' b}$ makes sure that $\succsim'$ *increases* the preference for a. Condition (c), in turn, guarantees that the derived preference intensities between b and a and between c and a are consistent with the DM's preferences between b and c. The following axiom states that one should always be able to find a new conditional preference relation that uniformly increases the preference for one of the acts.

Axiom 4.1 (Existence of a uniform preference increase) *There is an act a and a conditional preference relation $\succsim'$ that uniformly increases the preference for a relative to $\succsim$.*

Intuitively, this axiom requires the DM to think about the hypothetical situation in which one of the acts would become "uniformly more attractive" relative to the other acts, and demands that such a hypothetical scenario be compatible with the current conditional preference relation. As it turns out, this new axiom, together with the regularity axioms, characterizes precisely those conditional preference relations that admit an expected utility representation, provided no act is weakly dominated.

Theorem 4.1 (No weakly dominated choices) *Consider a finite set of acts A, a finite set of states S, and a conditional preference relation $\succsim$ on (A, S) such that no act weakly dominates another act under $\succsim$. Then, $\succsim$ has an expected utility representation, if and only if, it satisfies completeness, transitivity, continuity, preservation of indifference, preservation of strict preference, and existence of a uniform preference increase.*

In Section 5 we will see that these axioms may no longer be sufficient if there are weakly dominated acts.

4.3 Almost Unique Utility Representation

If the conditional preference relation $\succsim$ has an expected utility representation and there are no weakly dominated acts, then "typically" the utility differences will be unique up to a positive multiplicative constant. This is the content of the theorem below.

Theorem 4.2 (Almost unique utility representation) *Consider a finite set of acts A, a finite set of states S, and a conditional preference relation $\succsim$ on (A, S), such that no act weakly dominates another act under $\succsim$, and not all indifference sets $P_{a\sim b}$ are the same. Suppose that $\succsim$ is regular and satisfies the existence of a uniform preference increase. Then, for every two utility functions u, v that represent $\succsim$ there is some $\alpha > 0$ such that $u(a, s) - u(b, s) = \alpha \cdot (v(a, s) - v(b, s))$ for all $a, b \in A$ and all $s \in S$.*

In this case there would thus be $|S| + 1$ degrees of freedom for the expected utility representation. Moreover, under the conditions of the theorem, the utility difference $u(a, p) - u(b, p)$ at a belief p, which is unique up to a positive multiplicative constant, may be viewed as expressing the "preference intensity" between a and b at p. As an example, suppose that $0 < u(a, x) - u(b, x) = 2 \cdot (u(b, y) - u(a, y))$. Then, the DM will be indifferent between a and b at the belief $1/3[x] + 2/3[y]$, which seems to reflect that the intensity by which the DM prefers a to b at x is twice the intensity by which he prefers b to a at y. This indeed corresponds to the fact that the utility difference between a and b at x is twice as large as at y, in absolute terms. However, we will not enter the debates on whether such utility differences, or preference intensities, can be interpreted as reflecting neo-classical cardinal utility (see, for instance, Baccelli and Mongin (2016), Baumol (1958) and Moscati (2018)).

The above interpretation of the utility differences may no longer hold, however, if the conditions in the theorem above are not satisfied. Suppose there are three acts a, b and c, two states x and y, and let $\succsim$ be such that $a \succ_p b \succ_p c$ if $p(x) > 1/2$, $a \sim_p b \sim_p c$ if $p(x) = 1/2$, and $c \succ_p b \succ_p a$ if $p(x) < 1/2$. Hence, the three indifference sets $P_{a\sim b}, P_{a\sim c}$ and $P_{b\sim c}$ are all equal to $\{1/2[x] + 1/2[y]\}$, and thus coincide. Note that the utility functions u, v given by $u(a, x) = 3, u(b, x) = 2, u(c, x) = 0, u(a, y) = -3, u(b, y) = -2, u(c, y) = 0$

and $v(a,x) = 3, v(b,x) = 1, v(c,x) = 0, v(a,y) = -3, v(b,y) = -1, v(c,y) = 0$ both represent $\succsim$. Yet, the utility differences in u and v differ by more than just a multiplicative constant. The reason is that in this case, $\succsim$ does not provide us with sufficiently many data to derive the DM's preference intensity over the three acts at the various beliefs. A similar phenomenon may arise if there are weakly dominated acts.

5 General Case

In this section we start with an example showing that the regularity axioms, in combination with the existence of a uniform preference increase, may no longer guarantee an expected utility representation if weakly dominated acts are allowed. This leads to a new axiom, "existence of coherent uniform preference increases", which is formally presented below. It is shown that this new axiom, in combination with the regularity axioms, characterize in general precisely those conditional preference relations that have an expected utility representation. We also provide a procedure, the "utility design procedure", which can be used to generate a utility function that represents the conditional preference relation, provided it satisfies the axioms below. This procedure explicitly uses the existence of a coherent system of uniform preference increases. We conclude with a numerical example that illustrates the procedure.

5.1 Why Existence of a Uniform Preference Increase is Not Sufficient

Consider the conditional preference relation $\succsim$ in Figure 6, where b strictly dominates c and d strictly dominates a. It may be verified that $\succsim$ satisfies the regularity axioms, and that there is a uniform preference increase for each of the acts a, b, c and d. Despite this, $\succsim$ does not have an expected utility representation. Indeed, if the utility function u were to represent $\succsim$, then the utility of b would be equal to the utility of c at the points v and w outside the belief simplex. This would imply that the utility of b would be equal to the utility of c at the belief p^*. However, the DM strictly prefers b to c at p^*, which is a contradiction.

The problem with $\succsim$ is that a uniform preference increase for a is necessarily *incoherent* with a uniform preference increase for b. To see this, consider uniform preference increases $\succsim^a, \succsim^b$ for a and b, respectively, that increase the preference intensity for the associated act by the same degree, normalized to 1. For a given belief p, let $int_{d\succ b}(p)$ be the unique number such that $p \in (1 - int_{d\succ b}(p)) \cdot P_{b\sim d} + int_{d\succ b}(p) \cdot P_{b\sim^b d}$. As before, this represents the derived preference intensity between d and b at p.

Similarly, let the derived preference intensities $int_{c\succ a}(p)$ and $int_{a\succ b}(p)$ be such that $p \in (1 - int_{c\succ a}(p)) \cdot P_{a\sim c} + int_{c\succ a}(p) \cdot P_{a\sim^a c}$ and $p \in (1 - int_{a\succ b}(p)) \cdot P_{b\sim a} + int_{a\succ b}(p) \cdot P_{b\sim^b a}$. Then, we can "indirectly" derive the preference intensity between c and b by

$$int_{c\succ b}(p) = int_{c\succ a}(p) + int_{a\succ b}(p).$$

Here, we assume that the preference intensity between acts is additive, in the sense that $int_{c\succ b}(p) = int_{c\succ a}(p) + int_{a\succ b}(p)$. If the uniform preference increases $\succsim^a$ and $\succsim^b$ are to be coherent, then $p \in P_{c\sim d}$ if and only if $int_{c\succ b}(p) = int_{d\succ b}(p)$, and hence

$$p \in P_{c\sim d} \text{ if and only if } int_{c\succ a}(p) + int_{a\succ b}(p) = int_{d\succ b}(p). \tag{5.1}$$

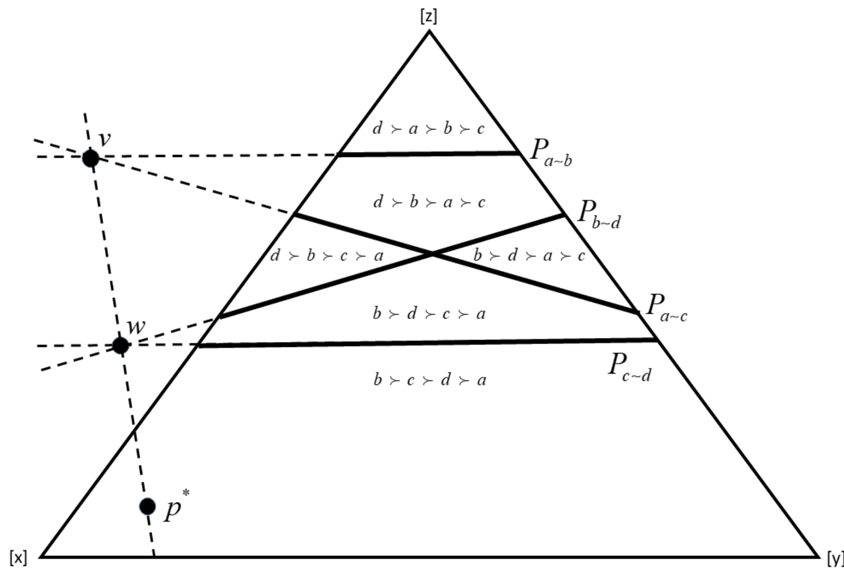

Figure 6: Existence of a uniform preference increase is not sufficient

Consider the point w in Figure 6, outside the belief simplex. Since w is on the line generated by $P_{b\sim d}$ we have that $int_{d\succ b}(w) = 0$. Moreover, since w is on the line generated by $P_{c\sim d}$, it follows from (5.1) that $int_{c\succ a}(w) + int_{a\succ b}(w) = 0$.

Consider next the point v in Figure 6, also outside the belief simplex. As v is on the lines generated by $P_{a\sim b}$ and $P_{a\sim c}$, we have that $int_{c\succ a}(v) = int_{a\succ b}(v) = 0$, and hence, in particular, $int_{c\succ a}(w) + int_{a\succ b}(w) = 0$.

Note that the belief p^* in Figure 6 is on the line through v and w. Since $int_{c\succ a}(v) + int_{a\succ b}(v) = int_{c\succ a}(w) + int_{a\succ b}(w) = 0$, it must be that $int_{c\succ a}(p^*) + int_{a\succ b}(p^*) = 0$ also. This would imply that $int_{c\succ b}(p^*) = 0$, and hence the DM should be indifferent between b and c at p^*. This, however, is a contraction to $\succsim$, and we thus conclude that the uniform preference increases for a and b are necessarily incoherent.

5.2 Existence of Coherent Uniform Preference Increases

Recall the definition of the domination graph $DG[\succsim]$ from Definition 4.1. Consider two acts a and b, and a path $a = c^1 \to c^2 \to ... \to c^K = b$ in $DG[\succsim]$, which we denote by Π_1. Then, the preference intensity between b and a at a belief p can "indirectly" be retrieved on the basis of this path, as follows: For every act d in this path, select a uniform preference increase $\succsim^d$ that uniformly increases the preference intensity for d by 1. By Definition 4.2 (b), there is for every edge (d, e) in this path a unique number $int_{e\succ d}(p)$ such that $p \in (1 - int_{e\succ d}(p) \cdot P_{d\sim e} + int_{e\succ d}(p) \cdot P_{d\sim^d e}$. As we have seen, $int_{e\succ d}(p)$ is the derived preference intensity between e and d. If we assume that the preference intensity between two acts is additive, then the preference intensity between b and a at p that can (directly or indirectly) be retrieved from the path Π_1 is

$$int_{b\succ a}(p) = int_{c^K \succ c^1}(p) = \sum\nolimits_{k=1}^{K-1} int_{c^{k+1}\succ c^k}(p) = \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p).$$

Similarly, if we consider another act $c \neq a, b$, and a path Π_2 from a to c in $DG[\succsim]$, then the preference intensity between c and a at p which can be retrieved from the path Π_2 is

$$int_{c\succ a}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p).$$

If the system of uniform preference increases is to be *coherent,* then $p \in P_{b\sim c}$ if and only if $int_{b\succ a}(p) = int_{c\succ a}(p)$. Hence,

$$p \in P_{b\sim c} \text{ if and only if } \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p),$$

independent of the act a and the chosen paths Π_1, Π_2 from a to b and a to c, respectively. This gives rise to the following definition.

Definition 5.1 (Coherent system of uniform preference increases) *Let $\succsim$ be a conditional preference relation, and $\{\succsim^a \mid a \in A\}$ a system of conditional preference relations. Then, $\{\succsim^a \mid a \in A\}$ is a coherent system of uniform preference increases relative to $\succsim$ if*

(a) $\succsim^a$ is regular for every act a, and $P_{a\succsim b} \subseteq P_{a\succ^a b}$ for every two acts a, b;

(b) for every two acts a, b with $(a, b) \in DG[\succsim]$, the conditional preference relation $\succsim^a$ has preference reversals on $\{a, b\}$, and for every belief p there is a unique number $int_{b\succ a}(p)$ such that $p \in (1 - int_{b\succ a}(p)) \cdot P_{a\sim b} + int_{b\succ a}(p) \cdot P_{a\sim^a b}$; and

(c) for every three acts a, b, c the conditional preference $\succsim^a$ coincides with $\succsim$ on $\{b, c\}$, and for every path Π_1 from a to b and every path Π_2 from a to c in $DG[\succsim]$ we have

$$p \in P_{b\sim c} \textit{ if and only if } \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p).$$

In particular, every conditional preference relation $\succsim^a$ must uniformly increase the preference for a relative to $\succsim$. Moreover, all uniform preference increases $\succsim^a$ should increase the preference intensity of the corresponding act by the same degree. Condition (c) makes sure that the (directly or indirectly) derived preference intensities between b and a and between c and a are always consistent with the DM's preferences between b and c.

In this definition we allow for the empty path Π, from an act to itself, and use the convention that $\sum_{(d,e)\in\Pi} int_{e\succ d}(p) = 0$ in this case. Also, b and c may be equal, and accordingly we set $P_{b\sim b} = \Delta(S)$. In the example of Figure 6, our argument above shows that there is no coherent system of uniform preference increases for the conditional preference relation at hand. The new axiom, which is stronger than the existence of a uniform preference increase, states that there should be a coherent system of uniform preference increases.

Axiom 5.1 (Existence of coherent uniform preference increases) *There is a coherent system* $\{\succsim^a \mid a \in A\}$ *of uniform preference increases relative to* $\succsim$.

This axiom, together with the regularity axioms, characterizes those conditional preference relations that have an expected utility representation.

Theorem 5.1 (Expected utility representation) *Consider a finite set of acts* A*, a finite set of states* S*, and a conditional preference relation* $\succsim$ *on* (A, S)*. Then,* $\succsim$ *has an expected utility representation, if and only if, it satisfies completeness, transitivity, continuity, preservation of indifference, preservation of strict preference, and existence of coherent uniform preference increases.*

In the following subsection we present a procedure to generate a representing utility function, provided the conditional preference relation satisfies the axioms above. This procedure explicitly uses a system of coherent uniform preference increases, the existence of which is guaranteed by the last axiom.

5.3 Utility Design Procedure

To formulate the procedure we rely on the following properties.

Lemma 5.1 (Properties needed for the procedure) *Let* $\succsim$ *be a regular conditional preference relation. Then,*

(a) there is a spanning tree for every undirected connected component of the domination graph $DG[\succsim]$;

(b) for every two undirected connected components G_1 *and* G_2 *in* $DG[\succsim]$*, either* a *strictly dominates* b *for every* a *in* G_1 *and* b *in* G_2*, or* b *strictly dominates* a *for every* a *in* G_1 *and* b *in* G_2;

(c) for every edge (a, b) *in* $DG[\succsim]$*, and every conditional preference relation* $\succsim^a$ *that uniformly increases the preference for* a *relative to* $\succsim$*, there is a belief* $p_1 \in P_{a\sim b}$ *and beliefs* $p_2, ..., p_{|S|} \in P_{a\sim^a b}$ *such that* $p_1, p_2, ..., p_{|S|}$ *are linearly independent.*

The utility design procedure below can be used to construct a utility function that represents the conditional preference relation at hand, provided it satisfies the axioms from Theorem 5.1.

Definition 5.2 (Utility design procedure) *Consider a regular conditional preference relation* $\succsim$*, together with a coherent system* $\{\succsim^a \mid a \in A\}$ *of uniform preference increases. Let* $G^1, ..., G^K$ *be the undirected connected components of* $DG[\succsim]$ *such that* a *strictly dominates* b *for every* $k \in \{1, ..., K-1\}$*, every* a *in* G^k *and every* b *in* G^{k+1}.

For every undirected connected component G^k*, select a spanning tree* T *for* G^k*, with root* r*. Choose some* $\alpha > 0$*. For every* $k \in \{0, ..., depth(T)\}$*, we define the utilities for acts with depth* k *by induction on* k*. If* $depth(a) = 0$*, then* $a = r$ *and we choose the utility* $v(r, s)$ *arbitrarily for every* $s \in S$*. For every* $k > 0$ *and* *act b with depth k, consider the predecessor a with depth $k-1$. Select beliefs $p_1 \in P_{a \sim b}$, $p_2, ..., p_{|S|} \in P_{a \sim^a b}$ such that $p_1, p_2, ..., p_{|S|}$ are linearly independent, and find the unique linear mapping $v_b : \mathbf{R}^S \to \mathbf{R}$ such that*

$$v_b(p_1) = v(a, p_1) \text{ and } v_b(p_k) = v(a, p_k) + \alpha \text{ for all } k \in \{2, ..., |S|\}.$$

Set $v(b, s) := v_b([s])$ for every $s \in S$.

Finally, select numbers $n^1, ..., n^K$ such that $v(a, s) + n^k > v(b, s) + n^{k+1}$ for every $k \in \{1, ..., K-1\}$, every a in G^k and every b in G^{k+1}. Set $u(a, s) := v(a, s) + n^k$ for every $k \in \{1, ..., K\}$, every a in G^k and every $s \in S$.

The ordering of the undirected connected components in terms of strict dominance, the existence of a spanning tree for every undirected connected component, and the selection of linearly independent beliefs $p_1, ..., p_{|S|}$ in the way indicated above, are all possible in the light of Lemma 5.1. For every undirected connected component there are at least $|S| + 1$ degrees of freedom for constructing the associated utilities: First, the "baseline utilities" $v(r, s)$ for the root can be chosen arbitrarily, and moreover we can freely select the "numeraire" $\alpha > 0$, measuring the utility increase associated with each of the uniform preference increases in the system. Additional degrees of freedom within a connected component may arise because of the choice of the particular coherent system of uniform preference increases.

5.4 Illustration of Procedure

We will now illustrate the utility design procedure by means of an example. Consider the conditional preference relation $\succsim$ in Figure 7, with three states and four acts. The vector (1/4, 0, 3/4) denotes the belief that assigns probabilities $1/4, 0$ and $3/4$ to x, y and z, respectively, and similarly for the other vectors. Note that b strictly dominates c and d strictly dominates a. There is a coherent system $\{\succsim^a, \succsim^b, \succsim^c, \succsim^d\}$ of uniform preference increases, represented by the grey lines in Figure 8. Below we will explain why this system is coherent. In the figure, $P_{a \sim^a b}, P_{a \sim^a c}, P_{d \sim^d b}$ and $P_{d \sim^d c}$ are the grey lines just below $P_{a \sim b}, P_{a \sim c}, P_{b \sim d}$ and $P_{c \sim d}$, respectively, whereas $P_{b \sim^b a}, P_{c \sim^c a}, P_{b \sim^b d}$, and $P_{c \sim^c d}$ are the grey lines just above $P_{a \sim b}, P_{a \sim c}, P_{b \sim d}$ and $P_{c \sim d}$, respectively. Also,

$$p_1 = (1/4,\ 0,\ 3/4),\ p_2 = (7/24,\ 0,\ 17/24),\ p_3 = (0,\ 7/24,\ 17/24)$$
$$q_1 = (1/3,\ 0,\ 2/3),\ q_2 = (9/24,\ 0,\ 15/24),\ q_3 = (0,\ 9/16,\ 7/16),$$
$$r_1 = (5/6,\ 0,\ 1/6),\ r_2 = (3/4,\ 0,\ 1/4),\ r_3 = (0,\ 9/14,\ 5/14).$$

The domination graph $DG[\succsim]$ is given by

$$\begin{array}{ccc} a & \longleftrightarrow & b \\ \updownarrow & & \updownarrow \\ c & \longleftrightarrow & d \end{array},$$

where $a \longleftrightarrow b$ means that there is an edge from a to b and an edge from b to a. Hence, there is only one undirected connected component.

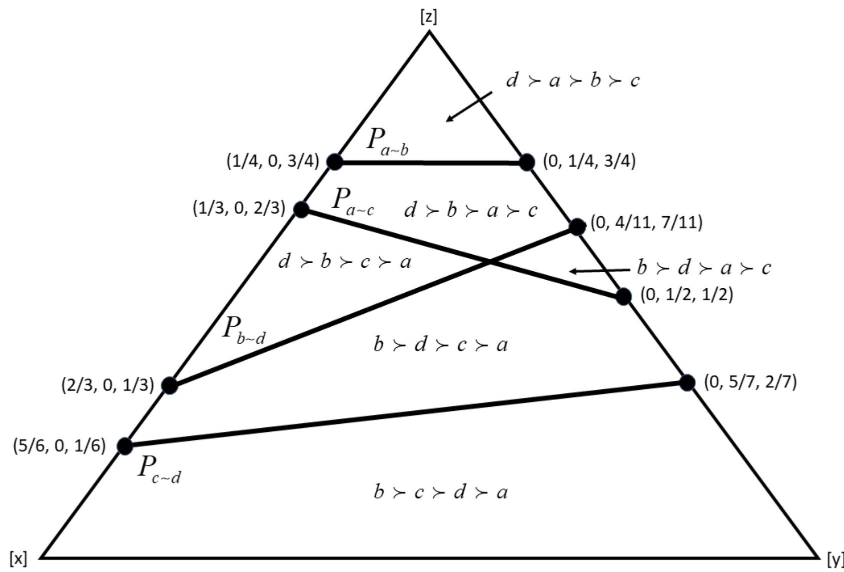

Figure 7: Illustration of utility design procedure

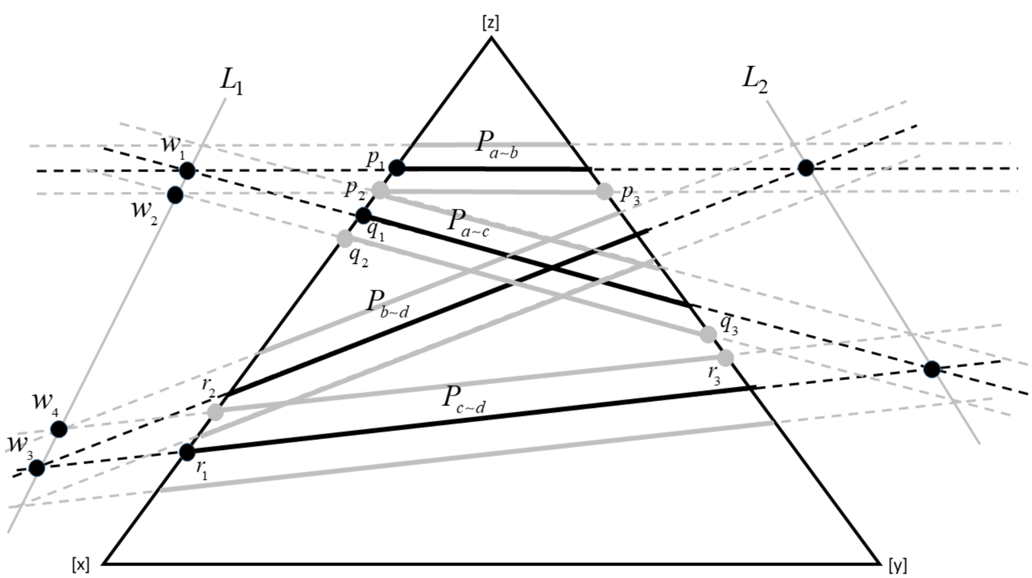

Figure 8: Coherent system of uniform preference increases

We will now show that the system of uniform preference increases is coherent. First, it may be verified that for all $p \in \Delta(S)$

$$int_{e \succ f}(p) + int_{f \succ e}(p) = 0 \text{ for all } (e,f) \in \{(a,b),(b,d),(d,c),(c,a)\} \tag{5.2}$$

$$int_{a \succ c}(p) + int_{c \succ d}(p) + int_{d \succ b}(p) + int_{b \succ a}(p) = 0 \tag{5.3}$$

$$int_{d \succ b}(p) + int_{b \succ a}(p) > 0 \text{ and } int_{c \succ d}(p) + int_{d \succ b}(p) < 0. \tag{5.4}$$

Here, (5.2) holds since $P_{e \sim f} = 1/2 P_{e \sim^e f} + 1/2 P_{f \sim^f e}$ for all $(e,f) \in \{(a,b),(b,d),(d,c),(c,a)\}$.

To see why (5.3) holds, we extend the coordinates $int_{f \succ e}(p)$ to all vectors in $\mathbf{R}^S$ with $\sum_{s \in S} v(s) = 1$, as follows. For every vector $v \in \mathbf{R}^S$ with $\sum_{s \in S} v(s) = 1$, and for all $(e,f) \in \{(a,b),(b,d),(d,c),(c,a)\}$, let $int_{f \succ e}(v)$ be the unique number with $v \in (1 - int_{f \succ e}(v)) \cdot \langle P_{e \sim f} \rangle + int_{f \succ e} \cdot \langle P_{e \sim^e f} \rangle$. Recall that $\langle P_{e \sim f} \rangle$ is the linear span of $P_{e \sim f}$. Consider the vectors w_1 and w_2 on the line L_1 in Figure 8. As $w_1 \in \langle P_{a \sim b} \rangle \cap \langle P_{a \sim c} \rangle$ and $w_2 \in \langle P_{a \sim^a b} \rangle \cap \langle P_{a \sim^a c} \rangle$ it follows that $int_{b \succ a}(w_1) = int_{a \succ c}(w_1) = 0$, $int_{b \succ a}(w_2) = 1$ and $int_{a \succ c}(w_2) = -1$. In particular, $int_{b \succ a}(w_1) + int_{a \succ c}(w_1) = int_{b \succ a}(w_2) + int_{a \succ c}(w_2) = 0$. Since the line L_1 passes through w_1 and w_2, and the coordinate $int_{f \succ e}(v)$ is linear in the vector v for every e, f (see Lemma 8.3 in Appendix A), it follows that $int_{b \succ a}(v) + int_{a \succ c}(v) = 0$ for every v on L_1. Similarly, it can be shown that $int_{c \succ d}(v) + int_{d \succ b}(v) = 0$ for every v on L_1, by considering the vectors w_3 and w_4 on L_1 in Figure 8. In particular, we have that $int_{a \succ c}(v) + int_{c \succ d}(v) + int_{d \succ b}(v) + int_{b \succ a}(v) = 0$ for every v on the line L_1. In a similar fashion, it can also be shown that $int_{a \succ c}(v) + int_{c \succ d}(v) + int_{d \succ b}(v) + int_{b \succ a}(v) = 0$ for every v on the line L_2 in Figure 8. As every belief can be written as a linear combination of a vector in L_1 and a vector in L_2, and the coordinate $int_{f \succ e}(v)$ is linear in the vector v, equation (5.3) follows.

Finally, (5.4) follows from the fact that for every vector v to the left of L_2 we have that $int_{d \succ b}(v) + int_{b \succ a}(v) > 0$, whereas for every vector v to the right of L_1 we have that $int_{c \succ d}(v) + int_{d \succ b}(v) < 0$. It may be verified that (5.2), (5.3) and (5.4) imply that the system of uniform preference increases is coherent.

We now implement the utility design procedure by choosing the coherent system of uniform preference increases from Figure 8, the spanning tree T with root a given by

$$\begin{array}{ccc} a & \longrightarrow & b \\ \downarrow & & \\ c & \longrightarrow & d \end{array}$$

and by setting $\alpha := 1$.

We start by choosing $u(a,x)$, $u(a,y)$ and $u(a,z)$ equal to 0. The acts of depth 1 are b and c. To find the utilities for b, consider the beliefs $p_1 = (1/4,\ 0,\ 3/4) \in P_{a \sim b}$, $p_2 = (7/24,\ 0,\ 17/24) \in P_{a \sim^a b}$ and $p_3 = (0,\ 7/24,\ 17/24) \in P_{a \sim^a b}$. We must find the unique linear mapping $v_b : \mathbf{R}^S \to \mathbf{R}$ with $v_b(p_1) = u(a, p_1)$, $v_b(p_2) = u(a, p_2) + 1$ and $v_b(p_3) = u(a, p_3) + 1$. This gives rise to the system of linear equations

$$1/4 v_b(x) + 3/4 v_b(z) = 0,\ 7/24 v_b(x) + 17/24 v_b(z) = 1 \text{ and } 7/24 v_b(y) + 17/24 v_b(z) = 1,$$

which has the unique solution $v_b(x) = 18$, $v_b(y) = 18$ and $v_b(z) = -6$. Thus, $u(b,x) = u(b,y) = 18$ and $u(b,z) = -6$.

In a similar fashion, we can use the beliefs q_1, q_2 and q_3 in Figure 8 to derive that $u(c, x) = 16$, $u(c, y) = 8$ and $u(c, z) = -8$.

The unique act of depth 2 is d. To find the utilities for d, consider the beliefs $r_1 = (5/6,\ 0,\ 1/6) \in P_{c \sim d}$, $r_2 = (3/4,\ 0,\ 1/4) \in P_{c \sim^c d}$ and $r_3 = (0,\ 9/14,\ 5/14) \in P_{c \sim^c d}$. We must find the unique linear mapping $v_d : \mathbf{R}^S \to \mathbf{R}$ with $v_d(r_1) = u(c, r_1)$, $v_d(r_2) = u(c, r_2) + 1$ and $v_d(r_3) = u(c, r_3) + 1$. This gives rise to the system of linear equations

$$5/6 v_d(x) + 1/6 v_d(z) = 12,\ 3/4 v_d(x) + 1/4 v_d(z) = 11 \text{ and } 9/14 v_d(y) + 5/14 v_d(z) = 46/14,$$

which has the unique solution $v_d(x) = 14$, $v_d(y) = 4$ and $v_d(z) = 2$. Thus, $u(d, x) = 14$, $u(d, y) = 5$ and $u(d, z) = 2$. Altogether, this yields the utility representation u given by

	x	y	z
a	0	0	0
b	18	18	-6
c	16	8	-8
d	14	4	2

.

The coherent system of uniform preference increases on which this representation is built is obtained by uniformly increasing the utility of a particular act by 1.

For the representation there are precisely 4 degrees of freedom: Three because the baseline utilities for a can be chosen freely, and one because the numeraire $\alpha > 0$ is arbitrary. As in this particular example the coherent system of uniform preference increases is unique up to a common proportionality factor, no additional degrees of freedom arise.

6 Strong Transitivity and the Line Property

In Section 4 we have seen that in the absence of weakly dominated acts, a regular conditional preference relation has an expected utility representation precisely when there is a uniform preference increase for at least one of the acts. But is there a way to easily verify whether such a uniform preference increase exists or not? In this section we provide an affirmative answer, by showing that in the absence of weakly dominated acts, the existence of a uniform preference increase is equivalent to two easily verifiable conditions: *strong transitivity* and the *line property,* provided the conditional preference relation is regular. These conditions can be tested directly by only considering the original conditional preference relation, without having to search for a uniform preference increase explicitly.

The first condition, *strong transitivity,* states that for every three acts a, b and c, the linear spans of the indifference sets $P_{a \sim b}$, $P_{a \sim c}$ and $P_{b \sim c}$ are such that the three pairwise intersections coincide.

Axiom 6.1 (Strong transitivity) *For every three acts $a, b, c \in A$ it holds that $\langle P_{a \sim b} \rangle \cap \langle P_{a \sim c} \rangle \subseteq \langle P_{b \sim c} \rangle$.*

See, for instance, Figure 5 where the linear spans of $P_{a\sim b}$, $P_{a\sim c}$ and $P_{b\sim c}$, when restricted to the plane where the sum of the coordinates is 1, all meet at the same point outside the belief simplex $\Delta(S)$, and hence the conditional preference relation at hand satisfies strong transitivity. If $P_{a\sim b}$, $P_{a\sim c}$ and $P_{b\sim c}$ would all meet inside the belief simplex, then strong transitivity would correspond to the usual transitivity of the indifference relation between acts. In that sense, strong transitivity can indeed be viewed as a strong version of transitivity, where this intersection property is also required outside the belief simplex. Note that the conditional preference relation in Figure 2, for which we have informally argued that there is no uniform preference increase, violates strong transitivity.

The diversity axiom in Gilboa and Schmeidler (2003) implies that for every strict ordering between a, b and c there must be a belief at which this ordering is realized. In particular, the indifference sets $P_{a\sim b}$, $P_{a\sim c}$ and $P_{b\sim c}$ must already meet inside the belief simplex, and hence strong transitivity would have no bite beyond plain transitivity under the diversity axiom.

The second condition, which we call the *line property,* is of a more technical nature, and only has bite if there are at least four acts. It states that, whenever we take four acts a, b, c, d, there must be a line that intersects the linear span of each of the six associated indifference sets at a unique point, such that the locations of these six intersection points are related to each other according to a specific formula.

Axiom 6.2 (Line property) *There is a line $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$, with $v, w \in \mathbf{R}^S$, that intersects each of the sets $\langle P_{a\sim b}\rangle$ at a single point $v + \lambda_{ab} w$, and where $\lambda_{ab} \neq \lambda_{ac}$ whenever $P_{a\sim b} \neq P_{a\sim c}$, such that*

$$(\lambda_{ab} - \lambda_{bd})(\lambda_{ac} - \lambda_{bc})(\lambda_{ad} - \lambda_{cd}) = (\lambda_{ab} - \lambda_{bc})(\lambda_{ac} - \lambda_{cd})(\lambda_{ad} - \lambda_{bd})$$

for all $a, b, c, d \in A$.

The following theorem shows that, under regularity and the absence of weakly dominated acts, the existence of a uniform preference increase is equivalent to the two conditions above.

Theorem 6.1 (Strong transitivity and line property) *Let $\succsim$ be a conditional preference relation where no act weakly dominates another act. Then, $\succsim$ has an expected utility representation, if and only if, it is regular, satisfies strong transitivity and satisfies the line property.*

This result has an interesting consequence for the case of two states. If $S = \{x, y\}$, strong transitivity is equivalent to the usual transitivity of the indifference relation between acts, whereas the line property is equivalent to the condition that

$$(p_{ab}(x) - p_{bd}(x))(p_{ac}(x) - p_{bc}(x))(p_{ad}(x) - p_{cd}(x)) = (p_{ab}(x) - p_{bc}(x))(p_{ac}(x) - p_{cd}(x))(p_{ad}(x) - p_{bd}(x)),$$

where p_{ef} is the unique belief in $P_{e\sim f}$ for every $\{a, b, c, d\}$. This leads to the following result.

Corollary 6.1 (Characterization for the case of two states) *Let $S = \{x, y\}$ and $\succsim$ a conditional preference relation where no act weakly dominates another act. Let $P_{a\sim b} = \{p_{ab}\}$ for all acts $a, b \in A$. Then, $\succsim$ has an expected utility representation, if and only if, $\succsim$ is regular and*

$$(p_{ab}(x) - p_{bd}(x))(p_{ac}(x) - p_{bc}(x))(p_{ad}(x) - p_{cd}(x)) = (p_{ab}(x) - p_{bc}(x))(p_{ac}(x) - p_{cd}(x))(p_{ad}(x) - p_{bd}(x))$$

for all $a, b, c, d \in A$.

A direct consequence is that for two acts and three states, *every* regular conditional preference relation for which there are no weakly dominated acts will have an expected utility representation. In fact, this property even holds if we would allow for weakly dominated acts.

7 Discussion

(a) Comparison with Savage. One important difference with the framework of Savage (1954) is that we view the DM's belief as a primitive notion, from which we can derive his preference relation over acts. This is precisely how a conditional preference relation is defined: It takes the belief as an input, and delivers the preferences over acts as an output. One of the beautiful features of the Savage framework is that the DM's belief can be derived from his preferences over acts. That is, Savage views the DM's preferences over acts as the primitive notion, which then induces his belief. There has been a long-standing debate about which of the two, belief or preferences, should be taken as the primitive object, and we do not want to enter this debate here. But the logic that underlies our framework is that the DM first reasons himself towards a belief, then forms his preferences over acts based on this belief, which finally allows him to make a choice based on this preference relation.

Another difference with Savage lies in the role of the utility function. In our model, the utility function generates the DM's preferences over acts for *all* possible beliefs over the states. As the Savage axiom system leads to a unique probabilistic belief over states, the utility function in the Savage framework can only be viewed in combination with this specific belief.

A final difference we would like to stress concerns the uniqueness of the utility representation. Recall from Theorem 4.2 that in the absence of weakly dominated acts, there are $|S| + 1$ degrees of freedom for the utility function in our framework, provided the sets of beliefs where the DM is indifferent beween two acts do not all coincide. Indeed, the utilities for one act can be chosen arbitrarily, leading to $|S|$ degrees of freedom, whereas the utility differences with a are all uniquely given, up to a positive multiplicative constant, leading to one additional degree of freedom. Unless all acts are equivalent, this is also the smallest number of degrees of freedom possible. There may be more degrees of freedom, up to $|A| \cdot |S|$, which would be the case if every act strictly dominates, or is strictly dominated by, another act.

In the Savage framework, on the other hand, the utility representation is always unique up to a positive affine transformation, leaving only two degrees of freedom. The reason is that a DM in the Savage framework holds preferences over *all possible* mappings from states to consequences, providing us with "more data" that restrict the possible utilities compared to a DM in our framework. However, the two degrees of freedom in Savage's framework are only possible because Savage's axiom of small event continuity implies that there are infinitely many states. We assume only finitely many states, but our "richness of data" comes from the fact that a conditional preference relation specifies a preference relation for infinitely many beliefs (if there are at least two states).

(b) Other foundations for expected utility in decision problems and games. The foundation for expected utility that is closest to ours is by Gilboa and Schmeidler (2003). As already stressed in the introduction, they also impose conditions on conditional preference relations. However, their axiom system is sufficient but not necessary for an expected utility representation. More precisely, it singles out those conditional preference relations that can be represented by a *diversified* utility function. By this we mean a utility matrix where no row is weakly dominated by, or equivalent to, an affine combination of at most three other rows. The crucial axiom in their analysis is *diversity,* which states that for every strict ordering of at most four acts there must be at least one belief for which that ordering obtains in the conditional preference relation at hand.

In contrast, we impose no restrictions on the utility matrix that can be used to represent the conditional preference relation. In particular, we allow for non-diversified utility matrices and, correspondingly, allow for non-diversified conditional preference relations. Note that all examples in this paper with three or four acts were examples of non-diversified conditional preference relations, having a non-diversified utility representation. By definition, diversity does not allow for weak dominance between acts. It may also be verified that the diversity condition exludes all cases with two states and more than two acts, and all cases with three states and more than three acts. Indeed, if we have two states and at least three acts, then there are 6 possible strict orderings on three acts, but at most 4 of these orderings will be possible in a regular conditional preference relation. Similarly, if we have three states and at least four acts, then there are 24 possible strict orderings on four acts, but at most 16 of these will be possible in a regular conditional preference relation. On the other hand, Gilboa and Schmeidler (2003) allow for infinitely many, even uncountably many, acts and states, whereas we do not.

Fishburn (1976) and Fishburn and Roberts (1978) concentrate on games, and assume that every player holds a preference relation over the combinations of randomized choices – or mixed strategies – of all the players. Combinations of mixed strategies may be viewed as lotteries with objective probabilities on the set of possible (pure) choice combinations in the game. By imposing certain axioms on these preference relations over mixed strategy combinations, they are able to identify those that admit an expected utility representation. It may thus be viewed as a generalization of von Neumann and Morgenstern's (1947) axiomatic characterization of expected utility for lotteries. The crucial difference with our approach is that we do not consider randomizations over choices, and that we use conditional preference relations as the primitive, rather than preferences over lotteries with objective probabilities.

In Aumann and Drèze (2002), a game is modelled as a mapping that assigns to every choice combination by the players a lottery over consequences for each of the players. The DM (a player in the game) is then assumed to hold a preference relation on the probability distributions over such mappings. Aumann and Drèze (2005) take a different approach, by supposing that the DM in a game holds a preference relation on lotteries which are defined over his own choices and over the possible consequences in the game. In both papers, it is shown that certain axioms on the preference relation lead to an expected utility representation that involves a unique, or essentially unique, probabilistic belief for the DM about the opponents' choice combinations. In that sense, these results are similar to Savage (1954). Another similarity is that also the models by Aumann and Drèze require the DM to hold preferences over objects that do not correspond to acts in the decision problem at hand.

Much in the spirit of the present paper, Mariotti (1995) also points out that a DM in Savage (1954) is required to hold preferences over acts that do not belong to his actual decision problem, and finds this problematic. Mariotti (1995) goes even further, and shows that certain game-theoretic principles are inconsistent with the axioms of completeness and monotonicity in Savage's framework, thus establishing a degree of "incompatibility" between games on the one hand and the framework of Savage on the other hand.

(c) Comparison with von Neumann and Morgenstern. In their foundation for expected utility, von Neumann and Morgenstern (1947) (vNM from now on) concentrate on lotteries, which are objective probability distributions on a fixed set of possible consequences. They impose axioms on preference relations over such lotteries, and show that they single out those preference relations admitting an expected utility representation. Despite the conceptual difference between objective probabilities and subjective beliefs, as we use them in our framework, there is a strong mathematical relationship between our approach and that of von vNM. In our framework, an act a, in combination with a belief p, can mathematically be identified with a vNM lottery over consequences $(b, s) \in A \times S$, where every consequence (b, s) occurs with probability $p(s)$ if $b = a$, and occurs with probability zero if $b \neq a$. Call this lottery $l(a, p)$. Now consider a conditional preference relation, which is a collection of preference relations $\succsim_p$ over acts, one for each possible probabilistic belief p. It can thus be identified with a partial preference relation in vNM where only lotteries $l(a, p)$ and $l(b, p)$ corresponding to the same belief p are ranked. In that sense, vNM require the DM to rank more lotteries than in our framework. As most of the vNM axioms require comparisons between lotteries that are not ranked within our framework, the vNM axiom system is not directly applicable to our setup.

(d) Comparison with state-dependent preference theory. Our notion of a conditional preference relation requires the DM to envisage a preference relation over acts for *every* probabilistic belief over states that is possible. Assuming that the DM holds an "actual" belief, this means that the DM must also reason about preference relations he would have contingent on *hypothetical* beliefs. In a decision problem this makes sense as the DM could change his belief in the light of new information, or because of new insights during his reasoning process. If the DM anticipates on revising his belief in the future then it seems natural for him to think about the preferences he would have as a result of changing beliefs. In games the need for preferences contingent on hypothetical beliefs is even more prominent, as the DM will be uncertain about the actual belief that his opponent holds in the game. The DM must therefore reason about the preference relation over acts that his opponent will hold for *every* possible belief the opponent may have. By fixing a utility matrix in a game with complete information, we thus assume that the players know the opponents' conditional preference relations, without knowing their precise beliefs.

Our axiom system singles out precisely those conditional preference relations where the DM's preferences at the various beliefs can all be represented by the *same* utility function. It thus imposes some "consistency" between the preference relations at the different beliefs. In this sense, our analysis is similar to the state-dependent preference framework that has been presented in Karni, Schmeidler and Vind (1983) and Wakker (1987), for instance. Different from Savage, these papers assume that the DM's preferences over consequences may depend on the specific state. In their consistency axiom, Karni, Schmeidler and Vind (1983) assume the DM to envisage a preference relation over acts contingent on a *hypothetical* belief, and require the DM's actual preference relation to be consistent with the latter preference relation. In our setting, we impose even more: We require the DM's actual preference relation over acts, based on his actual belief, to be consistent with the entire system of preference relations contingent on *each* of the other possible beliefs.

(e) Comparison with case-based decision theory. Case-based decision theory, as originally formulated in Gilboa and Schmeidler (1995), assumes that the DM evaluates an act based on how this act performed in previous decision problems. More precisely, assume that C represents the collection of decision problems, or cases, the DM faced in the past, and that $s(c)$ measures the similarity of decision problem c to the present decision problem. Then, the desirability of an act a in the present decision problem is measured by $\sum_{c \in C} s(c) \cdot u(a, c)$, where $u(a, c)$ is the utility that selecting act a generated in decision problem c.

Our framework can be embedded into case-based decision theory as follows: If a conditional preference relation is represented by a utility function u, then the desirability of an act a in the present decision problem, for a given $p \in \Delta(S)$, is given by $\sum_{s \in S} p(s) \cdot u(a, s)$. Now suppose that the states s represent decision problems that the DM faced in the past, and that $p(s)$ measures the similarity of problem s to the decision problem he is facing now. Then, the measure for the desirability of act a resembles exactly that in Gilboa and Schmeidler (1995).

Alternatively, one could still interpret p as a probabilistic belief over states, and identify every state s with the degenerate belief $[s]$ that assigns probability 1 to s. Suppose that, for some reason, the DM has had each of these degenerate beliefs $[s]$ in the past, and remembers the utility $u(a, s)$ that each act a generated under that belief. Then, every belief $[s]$ can be viewed as a case in the Gilboa-Schmeidler framework. If the DM's actual belief is p, then the belief probability $p(s)$ can be viewed as the similarity of the actual belief p to the past belief $[s]$. Also in this scenario, the measure for the desirability of act a in the actual problem, with the actual belief p, coincides with that of the Gilboa-Schmeidler framework.

(f) Utility differences as preference intensities. In Theorem 4.2 we have shown that in the absence of weakly dominated acts, the utility differences are unique up to a positive multiplicative constant, provided the sets of beliefs where the DM is indifferent between two acts do not all coincide. In that case, the expected utility difference between two acts a and b at a state s may be interpreted as the "preference intensity" between a and b at the state s. This is similar to how utility differences are interpreted in Anscombe and Aumann (1963) and Wakker (1989). The *state independence* axiom in Anscombe and Aumann (1963) states that the preference relation over objective lotteries on consequences must be independent of the state. This implies, in turn, that the utility differences between two consequences must be the same at every state, and these may be viewed as expressing the "preference intensity" between the two consequences.

The key condition in Wakker's (1989) axiom system is *state independent preference intensity.* The main idea is that the "preference intensity" between two consequences c_1 and c_2 at a state s can be measured by taking two acts, where one is strictly preferred to the other, and replacing the two acts at state s by c_1 and c_2, respectively, such that the DM becomes indifferent between the two new acts. State independent preference intensity requires that if the preference intensities between c_1 and c_2 and between c_3 and c_4 coincide at one state, then they must coincide at all states. In that case, the utility difference between two consequences will always be the same at all states, and may thus be viewed as expressing the "preference intensity" between the two consequences.

Analogously, the key condition in our axiom system, *existence of coherent uniform preference increases,* guarantees that the preference intensity between any two acts and at any belief is always uniquely defined (up to a scaling constant), provided there are no weakly dominated acts, and the sets of beliefs where the

DM is indifferent between two acts do not all coincide. To see this, consider a coherent system of uniform preference increases, which exists by the axiom. Then, at any belief, and for every two acts a and b, there are in general various chains of uniform preference increases in the system that can be used to derive the preference intensity between a and b. See Section 5.2 for the details. However, the axiom ensures that all these different chains result in the same preference intensity, thus establishing a unique preference intensity between every two acts and at every belief.

(g) Belief revision. A conditional preference relation does not only specify the DM's preferences over acts for a given belief, but also describes how these preferences would change if he were to *revise* his belief in the light of new information. In a dynamic decision problem or game it may happen, for instance, that some state is ruled out by some new information, forcing the DM to change his belief in response. And such information events may even take place sequentially, such that more and more states can be ruled out. The notion of a conditional preference relation is thus able to describe how the DM's preferences would change as a result of belief revision during the course of a dynamic decision problem or game.

(h) Game theory with conditional preference relations. In principle we could build an entire theory of games based on conditional preference relations, which may or may not satisfy our system of axioms. In a game, the DM would be a player i, his set of acts A_i would be the set of actions in the game, and the states would be the set $S_i = \times_{j \neq i} A_j$ of opponents' action profiles. Fix a conditional preference relation $\succsim^i$ for every player i. A *Nash equilibrium* (Nash (1950, 1951)) could be defined as a tuple of probability distributions $(\sigma_i)_{i \in I}$, with $\sigma_i \in \Delta(A_i)$ for every player i, such that $\sigma_i(a_i) > 0$ only if a_i is optimal for the induced preference relation $\succsim^i_{\sigma_{-i}}$. Here, σ_{-i} denotes the product of the probability distributions σ_j for $j \neq i$, which is a probability distribution over A_{-i} and hence a belief for player i. With this definition, a Nash equilibrium is thus interpreted as a tuple of beliefs about the opponents' actions, as in Aumann and Brandenburger (1995).

Similarly, *correlated rationalizability* (Brandenburger and Dekel (1987), Bernheim (1984), Pearce (1984)) could be defined by the recursive procedure where $A_i^0 := A_i$ for all players i, and

$$A_i^k := \{a_i \in A_i^{k-1} \mid a_i \text{ optimal for } \succsim^i_{p_i} \text{ for some } p_i \in \Delta(A_{-i}^{k-1})\}$$

for every $k \geq 1$. In fact, most – if not all – concepts in game theory could be generalized in terms of conditional preference relations.

(i) Equivalent acts. In this paper we have restricted attention to scenarios where no two acts are equivalent. In fact, our entire analysis can easily be extended to the case where equivalent acts are allowed. Suppose we start with a set of acts A where some acts are equivalent. Then, we can partition A into equivalence classes $\{A_1, A_2, ..., A_K\}$ with representative acts $a_1, a_2, ..., a_K$, and subsequently restrict the conditional preference relation $\succsim$ to the set $A^* = \{a_1, a_2, ..., a_K\}$, resulting in a new conditional preference relation $\succsim^*$. Then, Theorem 5.1 can be generalized as follows: The conditional preference relation $\succsim$ has an expected utility representation, if and only if, $\succsim^*$ is regular and satisfies the existence of coherent uniform preference increases. The proof is easy: If $\succsim^*$ is regular and satisfies the existence of coherent uniform preference increases, then by Theorem 5.1 it is represented by a utility function u. Extend u to a utility function v on $A \times S$ by setting $v(a,s) := u(a_k, s)$ for all acts $a \in A$ and all $s \in S$, where $a \in A_k$. Clearly, v will then represent $\succsim$. In the same way, the other results in this paper can also be extended to cases that allow for equivalent acts.

8 Appendix

In the appendix we start by providing the proofs of Section 5, before giving the proofs for Sections 3 and 4. The reason is that Theorem 5.1 is the main, and most general, result in this paper. Various parts in the proof of Theorem 5.1 also occur in the proofs of Theorems 3.1 and 4.1. Rather than fully repeating these steps in the proofs of Theorems 3.1 and 4.1, we simply refer back to the corresponding steps in the proof of Theorem 5.1.

8.1 Appendix A: Proofs of Section 5

We first prove Lemma 5.1, and subsequently Theorem 5.1.

8.1.1 Proof of Lemma 5.1

To prove Lemma 5.1 we need the following properties.

Lemma 8.1 (Linear structure of indifference sets) *Suppose that the conditional preference relation $\succsim$ is regular. Then, for every pair of acts a, b the following properties hold:*

(a) $P_{a\sim b} = \langle P_{a\sim b}\rangle \cap \Delta(S)$;

(b) if $\succsim$ has preference reversals on $\{a,b\}$, then $\langle P_{a\sim b}\rangle$ is a hyperplane, there are $|S|-1$ linearly independent beliefs in $P_{a\sim b}$, and there is a full support belief $p \in P_{a\sim b}$ with $p(s) > 0$ for all $s \in S$;

(c) if a weakly dominates b under $\succsim$ then $P_{a\sim b} = \{p \in \Delta(S) \mid \sum_{s\in S_{a\sim b}} p(s) = 1\}$; and

(d) there is a hyperplane H such that $P_{a\sim b} = H \cap \Delta(S)$.

Proof. (a) Clearly, $P_{a\sim b} \subseteq \langle P_{a\sim b}\rangle \cap \Delta(S)$. It remains to show that $\langle P_{a\sim b}\rangle \cap \Delta(S) \subseteq P_{a\sim b}$. We prove, by induction on k, that every $p \in \langle P_{a\sim b}\rangle \cap \Delta(S)$ which can be written as the linear combination of k elements in $P_{a\sim b}$, is in $P_{a\sim b}$. For $k = 1$ this is clear.

Take some $k \geq 2$, and assume that the statement above is true for $k-1$. Consider a $p \in \langle P_{a\sim b}\rangle \cap \Delta(S)$ that can be written as the linear combination of k elements in $P_{a\sim b}$. That is, $p = \lambda_1 p_1 + ... + \lambda_k p_k$, with $p_1, ..., p_k \in P_{a\sim b}$ and $\lambda_1, ..., \lambda_k \neq 0$. Assume, without loss of generality, that $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_k$. As $p \in \Delta(S)$ and $p_1, ..., p_k \in \Delta(S)$, we have that $\sum_{s\in S} p(s) = \sum_{s\in S} p_m(s) = 1$ for all m, and hence $\lambda_1 + ... + \lambda_k = 1$. Thus, $\lambda_1 \leq \frac{1}{2}$ and p can be written as

$$p = \lambda_1 p_1 + (1-\lambda_1)w, \text{ with } w = \frac{1}{1-\lambda_1}(\lambda_2 p_2 + ... + \lambda_k p_k).$$

We show that $w \in \Delta(S)$. By construction, $\sum_{s\in S} w(s) = 1$, and it thus remains to show that $w(s) \geq 0$ for all s. We distinguish two cases: If $\lambda_1 > 0$, then $\lambda_m > 0$ for all m. As $p_2, ..., p_k \in \Delta(S)$, it follows that $w(s) \geq 0$ for all s. Suppose now that $\lambda_1 < 0$. Then, $w = \frac{1}{1-\lambda_1}(p - \lambda_1 p_1)$. As $p, p_1 \in \Delta(S)$ and $\lambda_1 < 0$, it follows that $w(s) \geq 0$ for all s. We thus conclude that $w \in \Delta(S)$. Hence, $w \in \langle P_{a\sim b}\rangle \cap \Delta(S)$ is the linear combination of $k-1$ elements in $P_{a\sim b}$. By our induction assumption, $w \in P_{a\sim b}$. Therefore, $p = \lambda_1 p_1 + (1-\lambda_1)w$ is in $\Delta(S)$ with $p_1, w \in P_{a\sim b}$.

We will now show that $p \in P_{a\sim b}$. If $\lambda_1 \in [0,1]$ it follows by preservation of indifference. Suppose now that either $\lambda_1 < 0$ or $\lambda_1 > 1$. Assume first that $\lambda_1 < 0$. Then $w = \frac{1}{1-\lambda_1}(p - \lambda_1 p_1)$, where $\frac{1}{1-\lambda_1}, -\frac{\lambda_1}{1-\lambda_1} \in (0,1)$. Suppose, contrary to what we want to show, that $p \notin P_{a\sim b}$. Since $p_1 \in P_{a\sim b}$, it would follow by preservation of strict preference that $w \notin P_{a\sim b}$, which is a contradiction. Hence, we conclude that $p = \lambda_1 p_1 + (1-\lambda_1)w \in P_{a\sim b}$. If $\lambda_1 > 1$, then $p_1 = \frac{1}{\lambda_1}(p + (\lambda_1 - 1)w)$, where $\frac{1}{\lambda_1}, \frac{\lambda_1 - 1}{\lambda_1} \in (0,1)$. As $p_1, w \in P_{a\sim b}$, it follows by the same argument as above that $p \in P_{a\sim b}$.

Hence, every belief p that can be written as the linear combination of k elements in $P_{a\sim b}$ is again in $P_{a\sim b}$. By induction on k we conclude that $\langle P_{a\sim b}\rangle \cap \Delta(S) \subseteq P_{a\sim b}$.

(b) As $\succsim$ has preference reversals on $\{a,b\}$, the sets $S_{a\succ b}$ and $S_{b\succ a}$ must both be non-empty. Indeed, suppose that $S_{a\succ b}$ would be empty. Then, $[s] \in P_{b\succsim a}$ for all $s \in S$. But then it would follow by preservation of indifference and preservation of strict preference that $p \in P_{b\succsim a}$ for all beliefs $p \in \Delta(S)$, which would be a contradiction to our assumption that $\succsim$ has preference reversals on $\{a,b\}$.

Fix some states $y \in S_{a\succ b}$ and $z \in S_{b\succ a}$. By continuity, there must be some $\lambda_{yz} \in (0,1)$ such that

$$p_{yz} := (1-\lambda_{yz})[y] + \lambda_{yz}[z] \in P_{a\sim b}. \tag{8.1}$$

Similarly, for every $s \in S_{b\succ a}\backslash\{z\}$ there is some $\lambda_{ys} \in (0,1)$ such that

$$p_{ys} := (1-\lambda_{ys})[y] + \lambda_{ys}[s] \in P_{a\sim b}, \tag{8.2}$$

and for every $s \in S_{a\succ b}\backslash\{y\}$, there is some $\lambda_{zs} \in (0,1)$ such that

$$p_{zs} := (1-\lambda_{zs})[z] + \lambda_{zs}[s] \in P_{a\sim b}. \tag{8.3}$$

Consider the set

$$B := \{[s] \mid s \in S_{a\sim b}\} \cup \{p_{yz}\} \cup \{p_{ys} \mid s \in S_{b\succ a}\backslash\{z\}\} \cup \{p_{zs} \mid s \in S_{a\succ b}\backslash\{y\}\}, \tag{8.4}$$

which contains $|S|-1$ vectors in $P_{a\sim b}$. We show that all vectors in B are linearly independent.

Take some numbers α_s for $s \in S_{a\sim b}$, some number α_{yz}, some numbers α_{ys} for $s \in S_{b\succ a}\backslash\{z\}$ and some numbers α_{zs} for $s \in S_{a\succ b}\backslash\{y\}$ such that

$$\sum_{s\in S_{a\sim b}} \alpha_s[s] + \alpha_{yz}p_{yz} + \sum_{s\in S_{b\succ a}\backslash\{z\}} \alpha_{ys}p_{ys} + \sum_{s\in S_{a\succ b}\backslash\{y\}} \alpha_{zs}p_{zs} = \underline{0}.$$

By (8.1), (8.2) and (8.3), this sum is equal to

$$\sum_{s\in S_{a\sim b}} \alpha_s[s] + \alpha_{yz}((1-\lambda_{yz})[y] + \lambda_{yz}[z]) + \sum_{s\in S_{b\succ a}\backslash\{z\}} \alpha_{ys}((1-\lambda_{ys})[y] + \lambda_{ys}[s])$$
$$+ \sum_{s\in S_{a\succ b}\backslash\{y\}} \alpha_{zs}((1-\lambda_{zs})[z] + \lambda_{zs}[s])$$
$$= \sum_{s\in S_{a\sim b}} \alpha_s[s] + \left[\alpha_{yz}(1-\lambda_{yz}) + \sum_{s\in S_{b\succ a}\backslash\{z\}} \alpha_{ys}(1-\lambda_{ys})\right][y]$$
$$+ \left[\alpha_{yz}\lambda_{yz} + \sum_{s\in S_{a\succ b}\backslash\{y\}} \alpha_{zs}(1-\lambda_{zs})\right][z] + \sum_{s\in S_{b\succ a}\backslash\{z\}} \alpha_{ys}\lambda_{ys}[s] + \sum_{s\in S_{a\succ b}\backslash\{y\}} \alpha_{zs}\lambda_{zs}[s] = \underline{0}.$$

As all vectors in $\{[s] \mid s \in S\}$ are linearly independent, it follows that $\alpha_s = 0$ for all $s \in S_{a\sim b}$, that $\alpha_{ys}\lambda_{ys} = 0$ for all $s \in S_{b\succ a}\backslash\{z\}$, and that $\alpha_{zs}\lambda_{zs} = 0$ for all $s \in S_{a\succ b}\backslash\{y\}$. Since $\lambda_{ys} \in (0,1)$ for all $s \in S_{b\succ a}\backslash\{z\}$ and $\lambda_{zs} \in (0,1)$ for all $s \in S_{a\succ b}\backslash\{y\}$, it follows that $\alpha_{ys} = 0$ for all $s \in S_{b\succ a}\backslash\{z\}$ and $\alpha_{zs} = 0$ for all $s \in S_{a\succ b}\backslash\{y\}$. The sum above thus reduces to

$$\alpha_{yz}(1-\lambda_{yz})[y] + \alpha_{yz}\lambda_{yz}[z] = \underline{0}.$$

As $\lambda_{yz} \in (0,1)$, this implies that $\alpha_{yz} = 0$. We thus see that all coefficients in the linear combination above must be 0, and hence the vectors in B are linearly independent. As such, $P_{a\sim b}$ contains $|S|-1$ linearly independent beliefs.

As $P_{a\sim b} \neq \Delta(S)$, it follows that $\langle P_{a\sim b}\rangle$ is a hyperplane in $\mathbf{R}^S$. Consider the set B in (8.4). Then, the belief $p := 1/(|S|-1)\sum_{q\in B} q$ is a full support belief in $P_{a\sim b}$, with $p(s) > 0$ for all $s \in S$.

(c) Let $A = \{p \in \Delta(S) \mid \sum_{s\in S_{a\sim b}} p(s) = 1\}$. To show that $P_{a\sim b} \subseteq A$, take some $p \in P_{a\sim b}$. Assume, contrary to what we want to show, that $p \notin A$. Then, $p(s) > 0$ for some $s \in S_{a\succ b}$. As $p = \sum_{s\in S_{a\sim b}} p(s)\cdot[s] + \sum_{s\in S_{a\succ b}} p(s)\cdot[s]$ it follows by preservation of indifference and preservation of strict preference that $p \in P_{a\succ b}$. This is a contradiction to the assumption that $p \in P_{a\sim b}$. We thus conclude that $p \in A$. The inclusion $A \subseteq P_{a\sim b}$ follows directly by preservation of indifference. We thus see that $P_{a\sim b} = A$.

(d) If $\succsim$ has preference reversals on $\{a,b\}$ then we know from (b) that $\langle P_{a\sim b}\rangle$ is a hyperplane. By choosing $H = \langle P_{a\sim b}\rangle$ we know by (a) that $P_{a\sim b} = H \cap \Delta(S)$. Suppose next that a weakly dominates b under $\succsim$. Then, we know by (c) that

$$P_{a\sim b} = \{p \in \Delta(S) \mid \textstyle\sum_{s\in S_{a\sim b}} p(s) = 1\}. \tag{8.5}$$

Let $n \in \mathbf{R}^S$ be the vector with $n(s) = 0$ for all $s \in S_{a\sim b}$ and $n(s) = 1$ for all $s \in S\backslash S_{a\sim b}$. As a is not equivalent to b under $\succsim$ we know that $S_{a\sim b} \neq S$, and hence $n \neq \underline{0}$. Define $H := \{v \in \mathbf{R}^S \mid v \cdot n = 0\}$, which is a hyperplane. Then, it follows from (8.5) that $P_{a\sim b} = H \cap \Delta(S)$. ■

Proof of Lemma 5.1. (a) Consider an undirected connected component G of $DG[\succsim]$, with set of acts D. Let $D^1, D^2, ...$ be the sets of acts recursively defined by

$$D^1 := \{a \in D \mid a \text{ does not weakly dominate any } b \in D\},$$
$$D^k := \{a \in D \backslash (D^1 \cup ... \cup D^{k-1}) \mid a \text{ does not weakly dominate any } b \in D \backslash (D^1 \cup ... \cup D^{k-1})\}$$

for every $k \geq 2$. Let K be the smallest number such that $D^1 \cup ... \cup D^K = D$. Then, $(D^1, ..., D^K)$ is a partition of D. By construction, no two acts in D^k weakly dominate each other, and hence, for every $a, b \in D^k$ there is an edge from a to b and an edge from b to a. We show the following three properties.

Property 1. For every $k \in \{1, ..., K-1\}$, $a \in D^1 \cup ... \cup D^k$ and $b \in D^{k+1}$, act a does not weakly dominate b.
Proof of property 1. As $a \in D^1 \cup ... \cup D^k$, act a does not weakly dominate any act in $D \backslash (D^1 \cup ... \cup D^{k-1})$. In particular, a does not weakly dominate any act $b \in D^{k+1}$.

Property 2. For every $k \in \{1, ..., K-1\}$ and $b \in D^{k+1} \cup ... \cup D^K$ there is some $a \in D^k$ such that b weakly dominates a.
Proof of property 2. Take some $b \in D^{k+1} \cup ... \cup D^K$. As $b \notin D^k$, act b weakly dominates some $a^1 \in D \backslash (D^1 \cup ... \cup D^{k-1}) = D^k \cup ... \cup D^K$. If $a^1 \in D^k$ then the proof is complete. If $a^1 \notin D^k$, then a^1 weakly dominates some $a^2 \in D^k \cup ... \cup D^K$. And so on. In this way we obtain a chain of acts $(b, a^1, a^2, ...)$ in $D^k \cup ... \cup D^K$ where b weakly dominates a^1, a^1 weakly dominates a^2, and so on, and where every $a^m \notin D^k$ is followed by another act. But then, there must be an act $a^m \in D^k$ in this chain. Indeed, if all acts were in $D^{k+1} \cup ... \cup D^K$ then the chain would be infinite, and hence there would be a cycle $(a^m, a^{m+1}, ..., a^{m+n} = a^m)$. By transitivity, this would imply that a^m would weakly dominate itself, which is impossible. Hence, we must have a finite chain $(b, a^1, ..., a^m)$ where $a^m \in D^k$. Thus, by transitivity, b weakly dominates $a^m \in D^k$.

Property 3. For every $k \in \{1, ..., K-1\}$ there is some $a \in D^1 \cup ... \cup D^k$ and *$b \in D^{k+1}$ such that b does not strictly dominate a.*
Proof of property 3. Assume the statement is not true. Then, every act in D^{k+1} would strictly dominate every act in $D^1 \cup ... \cup D^k$. But then, by property 2, every act in $D^{k+1} \cup ... \cup D^K$ would strictly dominate every act in $D^1 \cup ... \cup D^k$. As such, there would be no edges between $D^1 \cup ... \cup D^k$ and $D^{k+1} \cup ... \cup D^K$. This, however, would contradict the assumption that G is an undirected connected component.

In the light of properties 1 and 3, there is for every $k \in \{1, ..., K-1\}$ some $a \in D^1 \cup ... \cup D^k$ and $b \in D^{k+1}$ such that there is an edge from a to b. Hence, there are acts $r^1, a^1, r^2, a^2, ..., r^K$ such that $r^k \in D^k$ for all $k \in \{1, ..., K\}$, $a^k \in D^1 \cup ... \cup D^k$ for all $k \in \{1, ..., K-1\}$, and there is an edge from a^k to r^{k+1} for every $k \in \{1, ..., K-1\}$. Here, it is possible that $r^k = a^m$ for some k, m. We now construct the spanning tree T as follows. We start by selecting the edges from r^1 to every other act in D^1. Then, we select the edge from a^1 to r^2. Subsequently, we select the edges from r^2 to every other edge in D^2, and the edge from a^2 to r^3, and so on. Finally, we select the edges from r^K to all other acts in D^K. This way, we obtain a spanning tree for the undirected connected component G.

(b) Consider two different undirected connected components G_1 and G_2 of $DG[\succsim]$. Take some acts a in G_1 and b in G_2. As there is no edge between a and b, either a strictly dominates b, or b strictly dominates a.

Assume, without loss of generality, that a strictly dominates b. Now take some arbitrary acts c in G_1 and d in G_2. We show that c strictly dominates d.

To do so, we first show that c strictly dominates b. Suppose not. Since there is no edge between c and b, it would follow that b strictly dominates c. Since a strictly dominates b and b strictly dominates c, it would follow that a strictly dominates c, which is a contradiction since a and c are in the same undirected connected component G_1.

Now, suppose that c does not strictly dominate d. As there is no edge between c and d, it must be that d strictly dominates c. As c strictly dominates b, it would follow that d strictly dominates b, which is a contradiction since b and d are in the same undirected connected component G_2. Hence, c strictly dominates d, which completes the proof for (b).

(c) Take two acts a and b such that (a,b) is in $DG[\succsim]$, and consider a uniform preference increase $\succsim^a$ for a. Since (a,b) is in $DG[\succsim]$ we have that b does not strictly dominate a, and hence $P_{a\sim b}$ is non-empty. Moreover, by part (b) in Definition 4.2, $\succsim^a$ has preference reversals on $\{a,b\}$. By Lemma 8.1 (a) and (b), we then know that $P_{a\sim^a b} = \langle P_{a\sim^a b}\rangle \cap \Delta(S)$, and that $P_{a\sim^a b}$ has $|S|-1$ linearly independent beliefs $p_2,...,p_{|S|}$. Select an arbitrary belief $p_1 \in P_{a\sim b}$. As $P_{a\sim b} \subseteq P_{a\succ^a b}$, it follows that $p_1 \notin \langle P_{a\sim^a b}\rangle$. Since $p_2,...,p_{|S|}$ are linearly independent beliefs in $\langle P_{a\sim^a b}\rangle$, we conclude that $p_1,p_2,...,p_{|S|}$ are linearly independent. This completes the proof. ■

8.1.2 Proof of Theorem 5.1

To prove Theorem 5.1 we need four preparatory results. The first states that, starting from a conditional preference relation that is represented by a utility function, we can always generate a uniform preference increase by lifting the utility of a given act by a small, constant amount.

Lemma 8.2 (Uniform preference increase from utility increase) *Let $u : A \times S \to \mathbf{R}$ be a utility function that represents the conditional preference relation $\succsim$. For every $\alpha > 0$, let the utility function u^α be given by $u^\alpha(a,s) := u(a,s) + \alpha$ for all $s \in S$, and $u^\alpha(b,s) := u(b,s)$ for every $b \neq a$ and every $s \in S$, and let $\succsim^\alpha$ be the conditional preference relation induced by u^α. Then, there is an $\varepsilon > 0$ such that for every $\alpha \in (0,\varepsilon)$, the conditional preference relation $\succsim^\alpha$ uniformly increases the preference for a relative to $\succsim$.*

Proof. Take some $b \neq a$. It may be easily be verified that $\succsim^\alpha$ is regular, and that $P_{a\succsim b} \subseteq P_{a\succ^\alpha b}$. To prove condition (b) in Definition 4.2, suppose that $(a,b) \in DG[\succsim]$, that is, a does not weakly dominate b and b does not strictly dominate a. Then, α can be chosen small enough such that $\succsim^\alpha$ has preference reversals on $\{a,b\}$. We will now show that, by choosing α small enough, we can guarantee that for every $p \in \Delta(S)$ there is a unique λ with $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim^\alpha b}$.

Observe that $P_{a\sim b} = H \cap \Delta(S)$ and $P_{a\sim^\alpha b} = H^\alpha \cap \Delta(S)$, where H and H^α are the sets given by

$$H := \{v \in \mathbf{R}^S \mid u(a,v) = u(b,v)\} \text{ and } H^\alpha := \{v \in \mathbf{R}^S \mid u(a,v) + \alpha = u(b,v)\}.$$

Consider the vector $n \in \mathbf{R}^S$ given by $n(s) := u(a,s) - u(b,s)$ for all $s \in S$. Since $\succsim$ has preference reversals on $\{a,b\}$ we know that $n \neq \underline{0}$. Moreover, by construction, $H = \{v \in \mathbf{R}^S \mid v \cdot n = 0\}$ and $H^\alpha = \{v \in \mathbf{R}^S \mid v \cdot n = -\alpha\}$, which implies that H and H^α are parallel hyperplanes.

As $\succsim^\alpha$ has preference reversals on $\{a,b\}$, the set $S_{b \succ^\alpha a}$ is non-empty. Hence, there is an $\varepsilon > 0$ such that

$$S_{b \succ a} \subseteq S_{b \succ^\alpha a} \text{ for every } \alpha \in (0, \varepsilon). \tag{8.6}$$

Take some $\alpha \in (0,\varepsilon)$. We show that for every $p \in \Delta(S)$ there is a unique λ with $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a \sim^\alpha b}$.

As $\succsim$ is regular, there is for every $x \in S_{a \succ b}$ and $y \in S_{b \succ a}$ a unique number $\lambda_{xy} \in (0,1)$ such that the belief $p_{xy} = (1-\lambda_{xy})[x] + \lambda_{xy}[y]$ is in $P_{a\sim b}$. Then,

$$P_{a\sim b} = conv(\{[s] \mid s \in S_{a\sim b}\} \cup \{p_{xy} \mid x \in S_{a\succ b},\ y \in S_{b \succ a}\}), \tag{8.7}$$

where *conv* denotes the convex hull. By (8.6) and the fact that $S_{a \succ b} \subseteq S_{a \succ^\alpha b}$, there is for every $x \in S_{a\succ b}$ and $y \in S_{b\succ a}$ a unique number $\mu_{xy} \in (0,1)$ such that the belief $q_{xy} = (1-\mu_{xy})[x] + \mu_{xy}[y]$ is in $P_{a\sim^\alpha b}$. Moreover, since $S_{a\sim b} \subseteq S_{a \succ^\alpha b}$ there is for every $s \in S_{a\sim b}$ and $y \in S_{b \succ a}$ a unique number $\mu_{sy} \in (0,1)$ such that the belief $q_{sy} = (1-\mu_{sy})[s] + \mu_{sy}[y]$ is in $P_{a\sim^\alpha b}$. Thus,

$$P_{a\sim^\alpha b} = conv(\{q_{sy} \mid s \in S_{a\succ b} \cup S_{a\sim b},\ y \in S_{b\succ a}\}). \tag{8.8}$$

Take some $p \in \Delta(S)$. We show that there is some $\lambda \in \mathbf{R}$ with $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim^\alpha b}$. We distinguish three cases.

Case 1. Suppose that $p \cdot n \geq 0$, where n is the vector defined above such that $H = \{v \in \mathbf{R}^S \mid v \cdot n = 0\}$ and $H^\alpha = \{v \in \mathbf{R}^S \mid v \cdot n = -\alpha\}$. Take some $q \in P_{a\sim^\alpha b}$. As $P_{a\sim^\alpha b} = H^\alpha \cap \Delta(S)$, we have that $n \cdot q = -\alpha$. Hence, there is some $\mu \in [0,1)$ such that $n \cdot ((1-\mu)p + \mu q) = 0$. As such, $r := (1-\mu)p + \mu q \in H \cap \Delta(S) = P_{a\sim b}$, and hence $p = \frac{1}{1-\mu} r - \frac{\mu}{1-\mu} q$, with $r \in P_{a\sim b}$ and $q \in P_{a\sim^\alpha b}$. Thus, $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim^\alpha b}$ for $\lambda = -\frac{\mu}{1-\mu}$.

Case 2. Suppose that $p \cdot n \leq -\alpha$. Take some $q \in P_{a\sim b}$. As $P_{a\sim b} = H \cap \Delta(S)$, we have that $n \cdot q = 0$. Hence, there is some $\mu \in [0,1)$ such that $n \cdot ((1-\mu)p + \mu q) = -\alpha$. As such, $r := (1-\mu)p + \mu q \in H^\alpha \cap \Delta(S) = P_{a\sim^\alpha b}$, and hence $p = \frac{1}{1-\mu} r - \frac{\mu}{1-\mu} q$, with $r \in P_{a\sim^\alpha b}$ and $q \in P_{a\sim b}$. Thus, $p \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim^\alpha b}$ for $\lambda = \frac{1}{1-\mu}$.

Case 3. Suppose that $p \cdot n \in (-\alpha, 0)$. Then, there is some $\lambda \in (0,1)$ such that $p \cdot n = (1-\lambda)0 + \lambda(-\alpha)$. Hence, $p \in H^{\lambda\alpha} = \{v \in \mathbf{R}^S \mid v \cdot n = -\lambda\alpha\}$. By (8.7) and (8.8) it follows that

$$H^{\lambda\alpha} \cap \Delta(S) = conv(\{(1-\lambda)[s] + \lambda q_{sy} \mid s \in S_{a\sim b},\ y \in S_{b\succ a}\} \cup \{(1-\lambda)p_{xy} + \lambda q_{xy} \mid x \in S_{a\succ b},\ y \in S_{b\succ a}\}).$$

As $p \in H^{\lambda\alpha} \cap \Delta(S)$, there is for every $s \in S_{a\succ b} \cup S_{a\sim b}$, $y \in S_{b\succ a}$ a number $\mu_{sy} \geq 0$, with $\sum_{s \in S_{a\succ b} \cup S_{a\sim b}, y \in S_{b\succ a}} \mu_{sy} = 1$, such that

$$\begin{aligned} p = & \sum_{s \in S_{a\sim b},\ y \in S_{b\succ a}} \mu_{sy} \cdot ((1-\lambda)[s] + \lambda q_{sy}) + \sum_{x \in S_{a\succ b},\ y \in S_{b\succ a}} \mu_{xy} \cdot ((1-\lambda)p_{xy} + \lambda q_{xy}) \\ & = (1-\lambda) \cdot (\sum_{s \in S_{a\sim b},\ y \in S_{b\succ a}} \mu_{sy} \cdot [s] + \sum_{x \in S_{a\succ b},\ y \in S_{b\succ a}} \mu_{xy} \cdot p_{xy}) \\ & \quad + \lambda \cdot (\sum_{s \in S_{a\sim b},\ y \in S_{b\succ a}} \mu_{sy} \cdot q_{sy} + \sum_{x \in S_{a\succ b},\ y \in S_{b\succ a}} \mu_{xy} \cdot q_{xy}), \end{aligned}$$

which is in $(1-\lambda)P_{a\sim b}+\lambda P_{a\sim^\alpha b}$ by (8.7) and (8.8).

By Cases 1, 2 and 3, we thus conclude that for every $p\in\Delta(S)$ there is some $\lambda\in\mathbf{R}$ with $p\in(1-\lambda)P_{a\sim b}+\lambda P_{a\sim^\alpha b}$. It remains to show that λ is unique. Suppose that λ,μ are such that $p\in(1-\lambda)P_{a\sim b}+\lambda P_{a\sim^\alpha b}$ and $p\in(1-\mu)P_{a\sim b}+\mu P_{a\sim^\alpha b}$. Since $P_{a\sim b}=H\cap\Delta(S)$ and $P_{a\sim^\alpha b}=H^\alpha\cap\Delta(S)$ it follows that $u(b,p)-u(a,p)=\lambda\alpha=\mu\alpha$, and hence $\lambda=\mu$. Thus, λ is unique. This shows that condition (b) in Definition 4.2 is satisfied if α is chosen small enough.

It remains to show condition (c) in Definition 4.2. Take some $b,c\neq a$. Then, $\succsim^\alpha$ coincides with $\succsim$ on $\{b,c\}$. Suppose now that $(a,b),(a,c)\in DG[\succsim]$. Then, we know from above that for every $p\in\Delta(S)$ there are unique numbers $int_{b\succ a}(p), int_{c\succ a}(p)$ such that $p\in(1-int_{b\succ a}(p))P_{a\sim b}+int_{b\succ a}(p)P_{a\sim^\alpha b}$ and $p\in(1-int_{c\succ a}(p))P_{a\sim c}+int_{c\succ a}(p)P_{a\sim^\alpha c}$. As $P_{a\sim b}=H\cap\Delta(S)$ and $P_{a\sim^\alpha b}=H^\alpha\cap\Delta(S)$, it follows that $u(b,p)-u(a,p)=int_{b\succ a}(p)\alpha$. Similarly, $u(c,p)-u(a,p)=int_{c\succ a}(p)\alpha$. As such, $u(b,p)-u(c,p)=(int_{b\succ a}(p)-int_{c\succ a}(p))\alpha$, and hence $P_{b\sim c}=\{p\in\Delta(S)\mid int_{b\succ a}(p)=int_{c\succ a}(p)\}$. Thus, condition (c) in Definition 4.2 holds if α is chosen small enough. As a consequence, $\succsim^\alpha$ uniformly increases the preference for a relative to $\succsim$ if α is chosen small enough, which completes the proof. ■

For the second result, consider a conditional preference relation $\succsim'$ that uniformly increases the preference for a relative to $\succsim$. Then, the preference intensity mapping $int_{b\succ a}$, which assigns to every belief p the coordinates $1-int_{b\succ a}(p)$ and $int_{b\succ a}(p)$ with respect to the sets $P_{a\sim b}$ and $P_{a\sim' b}$, is linear in the belief.

Lemma 8.3 (Preference intensity mapping is linear) *Let $\succsim$ be a regular conditional preference relation, and $\succsim'$ a conditional preference relation that uniformly increases the preference for a relative to $\succsim$. Consider an act b such that $(a,b)\in DG[\succsim]$. For every $p\in\Delta(S)$, let $int_{b\succ a}(p)$ be the unique number such that $p\in(1-int_{b\succ a}(p))\cdot P_{a\sim b}+int_{b\succ a}(p)\cdot P_{a\sim' b}$. Then, for every $p,q\in\Delta(S)$, and every $\mu\in\mathbf{R}$ such that $(1-\mu)p+\mu q\in\Delta(S)$, it holds that $int_{b\succ a}((1-\mu)p+\mu q)=(1-\mu)\cdot int_{b\succ a}(p)+\mu\cdot int_{b\succ a}(q)$.*

Proof. Let $p,q\in\Delta(S)$, and $\mu\in\mathbf{R}$ such that $r:=(1-\mu)p+\mu q\in\Delta(S)$. As $p\in(1-int_{b\succ a}(p))\cdot P_{a\sim b}+int_{b\succ a}(p)\cdot P_{a\sim' b}$ and $q\in(1-int_{b\succ a}(q))\cdot P_{a\sim b}+int_{b\succ a}(q)\cdot P_{a\sim' b}$, it follows that

$$r\in((1-\mu)(1-int_{b\succ a}(p))+\mu(1-int_{b\succ a}(q)))\cdot P_{a\sim b}+((1-\mu)\cdot int_{b\succ a}(p)+\mu\cdot int_{b\succ a}(q))\cdot P_{a\sim' b},$$

which implies that $int_{b\succ a}(r)=(1-\mu)\cdot int_{b\succ a}(p)+\mu\cdot int_{b\succ a}(q)$. This completes the proof. ■

The next result shows that, on the basis of a coherent system of uniform preference increases, we are not only able to derive the indifference sets $P_{b\sim c}$, but also the strict preference sets $P_{b\succ c}$.

Lemma 8.4 (Property of coherent uniform preference increases) *Let $\succsim$ be a regular conditional preference relation, and $\{\succsim^a\mid a\in A\}$ a coherent system of uniform preference increases. For every belief $p\in\Delta(S)$ and edge (d,e) in $DG[\succsim]$, let $int_{e\succ d}(p)$ be such that $p\in(1-int_{e\succ d}(p))\cdot P_{d\sim e}+int_{e\succ d}(p)\cdot P_{d\sim^d e}$. Then, for every three acts a,b,c, for every path Π_1 in $DG[\succsim]$ from a to b, and every path Π_2 in $DG[\succsim]$ from a to c,*

$$P_{b\succ c}=\{p\in\Delta(S)\mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p)>\sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}.$$

Proof. We prove the statement by induction on the total number of edges in Π_1 and Π_2. If this total number of edges is 0, then both Π_1 and Π_2 are the empty path, and hence $b = c = a$. Accordingly, $P_{b \succ c}$ and $\{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}$ are both the empty set, and the statement holds trivially.

Suppose now that the total number of edges in Π_1 and Π_2 is $k > 0$, and assume that the statement holds whenever the total number of edges is smaller than k. Since either Π_1 or Π_2 is not empty, we may assume without loss of generality that Π_1 is not the empty path. Let (f, b) be the last edge in Π_1, and let Π_1' be the path obtained from Π_1 by deleting the last edge (f, b). As $\{\succsim^a \mid a \in A\}$ is a coherent system of uniform preference increases we know, by definition, that

$$P_{b\sim c} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}. \tag{8.9}$$

We distinguish two cases: (1) $P_{f\sim b} \not\subseteq P_{b\sim c}$, and (2) $P_{f\sim b} \subseteq P_{b\sim c}$.

Case 1. Suppose that $P_{f\sim b} \not\subseteq P_{b\sim c}$. As (f, b) is an edge in $DG[\succsim]$ we know that $P_{f\sim b}$ is not empty, and hence we can take some $p^* \in P_{f\sim b} \backslash P_{b\sim c}$. Assume, without loss of generality, that $p^* \in P_{b\succ c}$. Then, by transitivity, $p^* \in P_{f\succ c}$. Note that Π_1' is a path from a to f, Π_2 is a path from a to c, and that the total number of edges in Π_1' and Π_2 is $k-1$. Since $p^* \in P_{f\succ c}$ we conclude by the induction assumption that

$$\sum_{(d,e)\in\Pi_1'} int_{e\succ d}(p^*) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*). \tag{8.10}$$

On the other hand, we know that $p^* \in P_{f\sim b}$ and hence $int_{b\succ f}(p^*) = 0$. Together with (8.10) we thus conclude that

$$\sum_{(d,e)\in\Pi_1} int_{e\succ d}(p^*) = \sum_{(d,e)\in\Pi_1'} int_{e\succ d}(p^*) + int_{b\succ f}(p^*) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*).$$

Hence,

$$p^* \in P_{b\succ c} \text{ and } \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p^*) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*).$$

Since $\succsim$ is regular, we know that

$$P_{b\succ c} = \{p \in \Delta(S) \mid (1-\lambda)p + \lambda p^* \notin P_{b\sim c} \text{ for every } \lambda \in [0,1]\}. \tag{8.11}$$

Moreover, in view of (8.9) and Lemma 8.3 we can similarly conclude that

$$\{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\} = \{p \in \Delta(S) \mid (1-\lambda)p + \lambda p^* \notin P_{b\sim c} \text{ for every } \lambda \in [0,1]\}.$$

Together with (8.11) we conclude that

$$P_{b\succ c} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}, \tag{8.12}$$

which completes Case 1.

Case 2. Suppose that $P_{f\sim b} \subseteq P_{b\sim c}$. Since $P_{f\sim b}$ is non-empty, we conclude that $P_{b\sim c}$ is non-empy, and hence either (b,c) or (c,b) is in $DG[\succsim]$. Assume, without loss of generality, that (b,c) is in $DG[\succsim]$. Let Π_1'' be the path obtained by putting the edge (b,c) after Π_1. Then, Π_1'' and Π_2 are both paths from a to c. As $\{\succsim^a \mid a \in A\}$ is a coherent system of uniform preference increases, we know that

$$\Delta(S) = P_{c\sim c} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1''} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}. \tag{8.13}$$

Since b and c are not equivalent, there is some $p^* \in \Delta(S)\backslash P_{b\sim c}$. Suppose that $p^* \in P_{b\succ c}$. Since, by definition of a uniform preference increase, $P_{b\succsim c} \subseteq P_{b\succ^b c}$, it follows that $P_{b\sim^b c} \subseteq P_{c\succ b}$. As $p^* \in P_{b\succ c}$ and $\succsim$ is regular, we conclude that $int_{c\succ b}(p^*) < 0$. Since, by (8.13),

$$\sum_{(d,e)\in\Pi_1} int_{e\succ d}(p^*) + int_{c\succ b}(p^*) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*)$$

it follows that

$$\sum_{(d,e)\in\Pi_1} int_{e\succ d}(p^*) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*).$$

We have thus found some $p^* \in P_{b\succ c}$ with $\sum_{(d,e)\in\Pi_1} int_{e\succ d}(p^*) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p^*)$. In the same way as in Case 1, we can then prove (8.12).

If $p^* \in P_{c\succ b}$ we conclude, in a similar way as above, that $int_{c\succ b}(p^*) > 0$, and

$$P_{c\succ b} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) < \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}.$$

Together with (8.9) this implies (8.12). This completes Case 2. By induction, (8.12) thus holds in general.■

The last result states that the indifference sets of a uniform preference increase are always parallel to the difference sets of the original conditional preference relation.

Lemma 8.5 (Geometry of uniform preference increases) *Let $\succsim$ be a regular conditional preference relation, and $\succsim'$ a conditional preference relation that uniformly increases the preference for a relative to $\succsim$. Let $b \neq a$ be such that (a,b) is in $DG[\succsim]$. Then, for every $p_{ab} \in P_{a\sim b}$ and $p'_{ab} \in P_{a\sim' b}$,*

$$P_{a\sim b} = (\langle P_{a\sim' b}\rangle + \{p_{ab} - p'_{ab}\}) \cap \Delta(S).$$

Proof. Set $A := (\langle P_{a\sim' b}\rangle + \{p_{ab} - p'_{ab}\}) \cap \Delta(S)$. We first show that $P_{a\sim b} \subseteq A$. Take some $q_{ab} \in P_{a\sim b}$. As $(a,b) \in DG[\succsim]$, we know by the definition of a uniform preference increase that $\succsim'$ has preference reversals on $\{a,b\}$. Hence, by Lemma 8.1 (b), there is a full support belief $p''_{ab} \in P_{a\sim' b}$ with $p''_{ab}(s) > 0$ for all $s \in S$. Let $q := \frac{1}{2}p_{ab} + \frac{1}{2}p''_{ab}$. Recall, by Lemma 8.1 (a) and (b), that $P_{a\sim' b} = \langle P_{a\sim' b}\rangle \cap \Delta(S)$ where $\langle P_{a\sim' b}\rangle$ is a hyperplane. As p''_{ab} is a full support belief, there is some $\varepsilon > 0$ small enough and some $\lambda \in \mathbf{R}$ such that $q \in (1-\lambda)\{(1-\varepsilon)p_{ab}+\varepsilon q_{ab}\}+\lambda P_{a\sim' b}$. In particular, $q \in (1-\lambda)P_{a\sim b}+\lambda P_{a\sim' b}$. Recall that $q \in \frac{1}{2}P_{a\sim b}+\frac{1}{2}P_{a\sim' b}$. Since, by definition of a uniform preference increase, there is a unique λ with $q \in (1-\lambda)P_{a\sim b}+\lambda P_{a\sim' b}$, it must be that $\lambda = \frac{1}{2}$, and hence $q \in \frac{1}{2}\{(1-\varepsilon)p_{ab}+\varepsilon q_{ab}\} + \frac{1}{2}P_{a\sim' b}$. That is, there is some $q'_{ab} \in P_{a\sim' b}$ with

$$q = \tfrac{1}{2}((1-\varepsilon)p_{ab}+\varepsilon q_{ab}) + \tfrac{1}{2}q'_{ab}.$$

As $q = \frac{1}{2}p_{ab} + \frac{1}{2}p''_{ab}$, it follows that $(1-\varepsilon)p_{ab} + \varepsilon q_{ab} + q'_{ab} = p_{ab} + p''_{ab}$, and hence

$$q_{ab} = \tfrac{1}{\varepsilon}(p''_{ab} - q'_{ab}) + p_{ab} = \tfrac{1}{\varepsilon}(p''_{ab} + \varepsilon p'_{ab} - q'_{ab}) + p_{ab} - p'_{ab} \in A.$$

Since this holds for every $q_{ab} \in P_{a\sim b}$, it follows that $P_{a\sim b} \subseteq A$.

We next show that $A \subseteq P_{a\sim b}$. Take some $p \in A$. We know from above that there is some full support belief $p''_{ab} \in P_{a\sim' b}$. Let $q := \frac{1}{2}p_{ab} + \frac{1}{2}p''_{ab}$. Similarly as above, we can conclude that there is some $\varepsilon > 0$ small enough and some $\lambda \in \mathbf{R}$ such that $q \in (1-\lambda)\{(1-\varepsilon)p_{ab} + \varepsilon p\} + \lambda P_{a\sim' b}$. Hence, there is some $q'_{ab} \in P_{a\sim' b}$ such that

$$q = (1-\lambda)((1-\varepsilon)p_{ab} + \varepsilon p) + \lambda q'_{ab}. \tag{8.14}$$

Moreover, by choosing $\varepsilon > 0$ small enough, we can guarantee that q'_{ab} is close to p''_{ab}, and hence will be a full support belief as well. As $p \in A$, there is some $v'_{ab} \in \langle P_{a\sim' b}\rangle$ such that $p = v'_{ab} + p_{ab} - p'_{ab}$, and hence, by (8.14),

$$q = (1-\lambda)(p_{ab} + \varepsilon(v'_{ab} - p'_{ab})) + \lambda q'_{ab}.$$

Together with the fact that $q = \frac{1}{2}p_{ab} + \frac{1}{2}p''_{ab}$, this yields

$$\tfrac{1}{2}p_{ab} + \tfrac{1}{2}p''_{ab} = (1-\lambda)(p_{ab} + \varepsilon(v'_{ab} - p'_{ab})) + \lambda q'_{ab},$$

and hence

$$(2\lambda - 1)p_{ab} = (2-2\lambda)\varepsilon(v'_{ab} - p'_{ab}) + 2\lambda q'_{ab} - p''_{ab}. \tag{8.15}$$

Note that the right-hand side of (8.15) is in $\langle P_{a\sim' b}\rangle$. If $\lambda \neq \frac{1}{2}$, then it would follow that $p_{ab} \in \langle P_{a\sim' b}\rangle \cap \Delta(S)$. Since, by Lemma 8.1 (a), $\langle P_{a\sim' b}\rangle \cap \Delta(S) = P_{a\sim' b}$, this would mean that $p_{ab} \in P_{a\sim' b}$. This, however, would be a contradiction since $p_{ab} \in P_{a\sim b} \subseteq P_{a\succ' b}$. We thus conclude that $\lambda = \frac{1}{2}$. By (8.14) we obtain that

$$q = \tfrac{1}{2}((1-\varepsilon)p_{ab} + \varepsilon p) + \tfrac{1}{2}q'_{ab}. \tag{8.16}$$

As $\succsim'$ is a uniform preference increase for a and (a,b) is in $DG[\succsim]$, there is some μ such that $(1-\varepsilon)p_{ab}+\varepsilon p \in (1-\mu)P_{a\sim b} + \mu P_{a\sim' b}$. That is, there are $r_{ab} \in P_{a\sim b}$ and $r'_{ab} \in P_{a\sim' b}$ with $(1-\varepsilon)p_{ab} + \varepsilon p = (1-\mu)r_{ab} + \mu r'_{ab}$. Together with (8.16) this yields

$$q = \tfrac{1}{2}((1-\mu)r_{ab} + \mu r'_{ab}) + \tfrac{1}{2}q'_{ab} = \tfrac{1}{2}(1-\mu)r_{ab} + \tfrac{1}{2}(1+\mu)\Big(\frac{\frac{1}{2}\mu}{\frac{1}{2}(1+\mu)}r'_{ab} + \frac{\frac{1}{2}}{\frac{1}{2}(1+\mu)}q'_{ab}\Big). \tag{8.17}$$

Recall that q'_{ab} is a full support belief. If we choose $\varepsilon > 0$ small enough, then $|\mu|$ will be small enough such that $w'_{ab} := \frac{\frac{1}{2}\mu}{\frac{1}{2}(1+\mu)} r'_{ab} + \frac{\frac{1}{2}}{\frac{1}{2}(1+\mu)} q'_{ab}$ is in $\Delta(S)$. As $w'_{ab} \in \langle P_{a\sim' b}\rangle$, and $P_{a\sim' b} = \langle P_{a\sim' b}\rangle \cap \Delta(S)$ by Lemma 8.1 (a), it follows that $w'_{ab} \in P_{a\sim' b}$. Since $r_{ab} \in P_{a\sim b}$, it follows from (8.17) that $q \in \frac{1}{2}(1-\mu)P_{a\sim b} + \frac{1}{2}(1+\mu)P_{a\sim' b}$. Recall from above that $q = \frac{1}{2}p_{ab} + \frac{1}{2}p''_{ab}$, and hence $q \in \frac{1}{2}P_{a\sim b} + \frac{1}{2}P_{a\sim' b}$. Since $\succsim'$ is a uniform preference increase for a and $(a, b) \in DG[\succsim]$, there must be a unique number λ such that $q \in (1-\lambda)P_{a\sim b} + \lambda P_{a\sim' b}$, and hence $\frac{1}{2}(1-\mu) = \frac{1}{2}$, which yields $\mu = 0$. As such, $(1-\varepsilon)p_{ab} + \varepsilon p = r_{ab}$, which implies that $p \in \langle P_{a\sim b}\rangle \cap \Delta(S)$. Since, by Lemma 8.1 (a), $P_{a\sim b} = \langle P_{a\sim b}\rangle \cap \Delta(S)$, we conclude that $p \in P_{a\sim b}$. As this holds for every $p \in A$, we conclude that $A \subseteq P_{a\sim b}$.

Together with the insight above that $P_{a\sim b} \subseteq A$ we thus see that $P_{a\sim b} = A$, which completes the proof.■

Proof of Theorem 5.1. **(a)** Suppose that $\succsim$ is represented by a utility function u. Then, it is easily verified that $\succsim$ is regular. It remains to show that $\succsim$ satisfies the existence of coherent uniform preference increases. Fix a number $\alpha > 0$, and for every act a let u^a be the utility function where $u^a(a, s) := u(a, s) + \alpha$ for all $s \in S$, and $u^a(b, s) := u(b, s)$ for all $b \neq a$ and all $s \in S$. Let $\succsim^a$ be the conditional preference relation induced by u^a.

We show that $\{\succsim^a \mid a \in A\}$ is a coherent system of uniform preference increases relative to $\succsim$, provided α is chosen small enough. It is easily verified that every $\succsim^a$ satisfies condition (a) in Definition 5.1. By Lemma 8.2, we can choose α small enough such that $\succsim^a$ satisfies condition (b) in Definition 5.1.

To prove condition (c) in Definition 5.1, observe first that for every three acts a, b, c, the conditional preference relation $\succsim^a$ coincides with $\succsim$ on $\{b, c\}$. Now, consider a path Π_1 from a to b and a path Π_2 from a to c in the domination graph $DG[\succsim]$. For every $p \in \Delta(S)$ and every edge (d, e) in Π_1 and Π_2, let $int_{e\succ d}(p)$ be the unique number such that $p \in (1 - int_{e\succ d}(p)) \cdot P_{d\sim e} + int_{e\succ d}(p) \cdot P_{d\sim^d e}$. As $u(d, q) = u(e, q)$ for all $q \in P_{d\sim e}$ and $u(d, q) + \alpha = u(e, q)$ for all $q \in P_{d\sim^d e}$, it follows that $u(e, p) - u(d, p) = int_{e\succ d}(p) \cdot \alpha$. Hence,

$$u(b, p) - u(a, p) = \sum_{(d,e)\in\Pi_1} u(e, p) - u(d, p) = \alpha \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p).$$

Similarly, $u(c, p) - u(a, p) = \alpha \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)$, and hence

$$P_{b\sim c} = \{p \in \Delta(S) \mid u(b, p)-u(a, p) = u(c, p)-u(a, p)\} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}.$$

This implies that $\{\succsim^a \mid a \in A\}$ is a coherent system of uniform preference increases relative to $\succsim$.

(b) Let $\succsim$ be a regular conditional preference relation that satisfies the existence of coherent uniform preference increases. Select a coherent system $\{\succsim^a \mid a \in A\}$ of uniform preference increases relative to $\succsim$, and choose an undirected connected component G, with set of acts D, in the domination graph $DG[\succsim]$. Choose, moreover, a spanning tree T for G with root r, some utilities $v(r, s)$ for $s \in S$, and some $\alpha > 0$. Let v be the corresponding utility function for the acts in D generated by the utility design procedure. We show that v represents $\succsim$ on D.

Take a pair of acts a, b in D, let Π_1 be the path from r to a, and Π_2 the path from r to b. Note that Π_1 or Π_2 may be the empty path if $r = a$ or $r = b$. As the system $\{\succsim^c \mid c \in A\}$ of uniform preference increases is coherent, we know that

$$P_{a\sim b} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}. \tag{8.18}$$

Moreover, by Lemma 8.4,

$$P_{a\succ b} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}. \tag{8.19}$$

Consider an edge (d, e) in Π_1, and let $p_1 \in P_{d\sim e}$ and $p_2, ..., p_{|S|} \in P_{d\sim^d e}$ be the beliefs selected by the utility design procedure, such that $p_1, p_2, ..., p_{|S|}$ are linearly independent. Then, by definition, $v(e, p_1) = v(d, p_1)$ and $v(e, p_k) = v(d, p_k) + \alpha$ for all $k \in \{2, ..., |S|\}$. By Lemma 8.5 we know that

$$P_{d\sim e} = (\langle P_{d\sim^d e}\rangle + \{p_1 - p_2\}) \cap \Delta(S). \tag{8.20}$$

Recall from Lemma 8.1 (a) that $P_{d\sim^d e} = \langle P_{d\sim^d e}\rangle \cap \Delta(S)$. Moreover, by Lemma 8.1 (b), $\{p_2, ..., p_{|S|}\}$ is a basis for $\langle P_{d\sim^d e}\rangle$. Hence, $P_{d\sim^d e}$ contains precisely those $p \in \Delta(S)$ that can be written as $p = \sum_{k=2}^{|S|} \lambda_k p_k$, with $\sum_{k=2}^{|S|} \lambda_k = 1$. Since, by the design of the utilities, $v(e, p_k) = v(d, p_k) + \alpha$ for all $k \in \{2, ..., |S|\}$, and $v(e, p_1) = v(d, p_1)$, it follows that

$$P_{d\sim^d e} = \{p \in \Delta(S) \mid v(e, p) = v(d, p) + \alpha\}. \tag{8.21}$$

Moreover, by (8.20), $P_{d\sim e}$ contains precisely those $p \in \Delta(S)$ that can be written as $p = \sum_{k=2}^{|S|} \lambda_k p_k + p_1 - p_2$, with $\sum_{k=2}^{|S|} \lambda_k = 1$. Since, by design, $v(e, p_k) = v(d, p_k) + \alpha$ for all $k \in \{2, ..., |S|\}$, and $v(e, p_1) = v(d, p_1)$, it follows that

$$P_{d\sim e} = \{p \in \Delta(S) \mid v(e, p) = v(d, p)\}. \tag{8.22}$$

From (8.21) and (8.22) we conclude that $v(e, p) - v(d, p) = int_{e\succ d}(p) \cdot \alpha$ for all $p \in \Delta(S)$, and hence

$$v(a, p) - v(r, p) = \sum_{(d,e)\in\Pi_1} (v(e, p) - v(d, p)) = \alpha \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) \tag{8.23}$$

for all $p \in \Delta(S)$. Similarly,

$$v(b, p) - v(r, p) = \sum_{(d,e)\in\Pi_2} (v(e, p) - v(d, p)) = \alpha \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p). \tag{8.24}$$

Let $P_{v(a)=v(b)}$ and $P_{v(a)>v(b)}$ be the sets of beliefs p where $v(a, p) = v(b, p)$ and $v(a, p) > v(b, p)$, respectively. Then, by (8.23) and (8.24),

$$P_{v(a)=v(b)} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) = \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\} \text{ and} \tag{8.25}$$

$$P_{v(a)>v(b)} = \{p \in \Delta(S) \mid \sum_{(d,e)\in\Pi_1} int_{e\succ d}(p) > \sum_{(d,e)\in\Pi_2} int_{e\succ d}(p)\}. \tag{8.26}$$

If we compare (8.18), (8.19), (8.25) and (8.26) we see that $P_{v(a)=v(b)} = P_{a\sim b}$ and $P_{v(a)>v(b)} = P_{a\succ b}$, which implies that v represents $\succsim$ on $\{a, b\}$. Hence, v represents $\succsim$ on every pair in D, and thus represents $\succsim$ on D. As this holds for every undirected connected component in $DG[\succsim]$, it follows that the utility function u produced by the utility design procedure represents $\succsim$. This completes the proof. ∎

8.2 Appendix B: Proofs of Sections 3 and 4

In this section we prove Theorem 3.1, Theorem 4.1 and Theorem 4.2.

Proof of Theorem 3.1. **(a)** Suppose that $\succsim$ is represented by a utility function u. Then, by Theorem 5.1, $\succsim$ is regular.

(b) Suppose that $\succsim$ is regular. We distinguish two cases: (1) a weakly dominates b, or vice versa, (2) there are preference reversals between a and b.

Case 1. Suppose, without loss of generality, that a weakly dominates b. Then, by Lemma 8.1 (c), $P_{a\sim b} = \{p \in \Delta(S) \mid \sum_{s\in S_{a\sim b}} p(s) = 1\}$. Construct a utility function u with $u(a, s) = u(b, s)$ for all $s \in S_{a\sim b}$, and $u(a, s) > u(b, s)$ for all $s \in S_{a\succ b}$. Then, u represents $\succsim$.

Case 2. Suppose there are preference reversals between a and b. Then, by Lemma 8.1 (b), $\langle P_{a\sim b}\rangle$ is a hyperplane, and there are $|S| - 1$ linearly independent beliefs $p_2, ..., p_{|S|} \in P_{a\sim b}$. Choose some $p_1 \in P_{a\succ b}$. Similarly to the utility design procedure, we can find a utility function u with $u(a, p_1) > u(b, p_1)$ and $u(a, p_k) = u(b, p_k)$ for all $k \in \{2, ..., |S|\}$. Then, in the same way as in the proof of Theorem 5.1, it can be shown that u represents $\succsim$. This completes the proof. ∎

Proof of Theorem 4.1. **(a)** Suppose that $\succsim$ is represented by a utility function u. Then, by Theorem 5.1, $\succsim$ is regular and satisfies the existence of a uniform preference increase.

(b) Suppose that $\succsim$ is regular and satisfies the existence of a uniform preference increase. Hence, there is an act a and a conditional preference relation $\succsim^a$ that uniformly increases the preference for a relative to $\succsim$. As there are preference reversals on all pairs of acts, $(a, b) \in DG[\succsim]$ for every act $b \neq a$. Hence, $DG[\succsim]$ only has one undirected connected component, and there is a spanning tree T for $DG[\succsim]$ that consists of the edges (a, b) for every $b \neq a$. Note that $\succsim^a$ is sufficient to implement the utility design procedure with respect to this spanning tree. Let u be a utility function generated by the utility design procedure, for a specific choice of $\alpha > 0$. Then, it follows from the proof of Theorem 5.1 that u represents $\succsim$. ∎

To prove Theorem 4.2, we need the following result.

Lemma 8.6 (Indifference sets in absence of weak dominance) *Let $\succsim$ be a regular conditional preference relation where there are at least three acts, no acts are weakly dominated, and such that not all indifference sets $P_{a\sim b}$ coincide. Then, for every act a there are acts $b^*, c^* \neq a$ such that (a) $P_{a\sim b^*} \neq P_{a\sim c^*}$, (b) for every $d \neq a, b^*, c^*$ either $P_{a\sim d} \neq P_{b^*\sim d}$ or $P_{a\sim d} \neq P_{c^*\sim d}$, and (c) for every $d, e \neq a, b^*, c^*$ either $P_{a\sim d} \neq P_{a\sim e}$, or $P_{b^*\sim d} \neq P_{b^*\sim e}$, or $P_{c^*\sim d} \neq P_{c^*\sim e}$.*

Proof. We start with a general observation. Suppose that $P_{b\sim c} = P_{b\sim d}$. Then, by transitivity, $P_{b\sim c} \subseteq P_{c\sim d}$ and hence $\langle P_{b\sim c}\rangle \subseteq \langle P_{c\sim d}\rangle$. As, by Lemma 8.1 (b), $\langle P_{b\sim c}\rangle$ and $\langle P_{c\sim d}\rangle$ are hyperplanes, it must be that $\langle P_{b\sim c}\rangle = \langle P_{c\sim d}\rangle$ and hence, by Lemma 8.1 (a), $P_{b\sim c} = P_{c\sim d}$. We thus conclude that

$$P_{b\sim c} = P_{b\sim d} \text{ implies } P_{b\sim c} = P_{c\sim d}. \tag{8.27}$$

(a) Suppose that (a) would not hold. Then, for a given act a^*, and every two acts $b, c \neq a^*$, we must have that $P_{a^*\sim b} = P_{a^*\sim c}$. Hence, in view of (8.27), $P_{a^*\sim b} = P_{a^*\sim c} = P_{b\sim c}$ for every $b, c \neq a^*$. But then, it follows that all indifference sets would coincide, which contradicts the assumption in the lemma. Hence, (a) holds.

(b) Take some $d \neq a, b^*, c^*$ and assume, contrary to what we want to show, that $P_{a\sim d} = P_{b^*\sim d}$ and $P_{a\sim d} = P_{c^*\sim d}$. Then, by (8.27), $P_{a\sim b^*} = P_{a\sim d} = P_{a\sim c^*}$, which contradicts (a). Hence, (b) holds.

(c) Take some $d, e \neq a, b^*, c^*$ and assume, contrary to what we want to show, that $P_{a\sim d} = P_{a\sim e}, P_{b^*\sim d} = P_{b^*\sim e}$ and $P_{c^*\sim d} = P_{c^*\sim e}$. Then, by (8.27), $P_{a\sim d} = P_{d\sim e} = P_{b^*\sim d}$ and $P_{a\sim d} = P_{d\sim e} = P_{c^*\sim d}$, which contradicts (b) Hence, (c) holds, and the proof is complete. ■

Proof of Theorem 4.2. Let u, v be two different utility representations for $\succsim$. Select the acts a, b^*, c^* as in Lemma 8.6. Since there are preference reversals on $\{a, b^*\}$, there is some $p^* \in P_{a\succ b^*}$. Define $\alpha := (u(a,p^*) - u(b^*,p^*))/(v(a,p^*) - v(b^*,p^*))$. We show that

$$u(d,p) - u(e,p) = \alpha \cdot (v(d,p) - v(e,p)) \text{ for all } p \in \Delta(S) \text{ and all } d, e \in A. \tag{8.28}$$

To this purpose, consider a sequence of pairwise different acts $(a^1, a^2, ..., a^{|A|})$ that covers the whole set A, and such that $a^1 = a$, $a^2 = b^*$ and $a^3 = c^*$. We show, by induction on m, that (8.28) holds for all $d, e \in \{a^1, ..., a^m\}$.

If $m = 2$ we have that $\{a^1, a^2\} = \{a, b^*\}$. As there are preference reversals on $\{a, b^*\}$, it follows by Lemma 8.1 (b) that there are $|S| - 1$ linearly independent beliefs $p_1, ..., p_{|S|-1}$ in $P_{a\sim b^*}$. Moreover, $p^* \notin \langle P_{a\sim b^*}\rangle$ as $P_{a\sim b^*} = \langle P_{a\sim b^*}\rangle \cap \Delta(S)$ by Lemma 8.1 (a). Hence, $\{p_1, ..., p_{|S|-1}, p^*\}$ are linearly independent, and thus form a basis for $\mathbf{R}^S$. As, by construction, $u(a,p_k) - u(b^*,p_k) = \alpha \cdot (v(a,p_k) - v(b^*,p_k)) = 0$ for all $k \in \{1, ..., |S|-1\}$, and $u(a,p^*) - u(b^*,p^*) = \alpha \cdot (v(a,p^*) - v(b^*,p^*))$, it follows that (8.28) holds for a, b^* and every $p \in \Delta(S)$.

Let $m \in \{3, ..., |A|\}$ and suppose that (8.28) holds for every $d, e \in \{a^1, ..., a^{m-1}\}$. By Lemma 8.6 (a) there are $d, e \in \{a^1, ..., a^{m-1}\}$ such that $P_{d\sim a^m} \neq P_{e\sim a^m}$. Hence, there is some $q \in P_{e\sim a^m} \backslash P_{d\sim a^m}$. By the induction assumption, we know that

$$u(d,q) - u(e,q) = \alpha \cdot (v(d,q) - v(e,q)). \tag{8.29}$$

Moreover, as $q \in P_{e\sim a^m}$, it must be that

$$u(e,q) - u(a^m,q) = \alpha \cdot (v(e,q) - v(a^m,q)) = 0,$$

which, together with (8.29) implies that

$$u(d,q) - u(a^m,q) = \alpha \cdot (v(d,q) - v(a^m,q)). \tag{8.30}$$

By Lemma 8.1 (b) there are $|S|-1$ linearly independent beliefs $p_1, ..., p_{|S|-1}$ in $P_{d\sim a^m}$. By construction,

$$u(d,p_k) - u(a^m,p_k) = \alpha \cdot (v(d,p_k) - v(a^m,p_k)) = 0 \text{ for all } k \in \{1,...,|S|-1\}. \tag{8.31}$$

Since $q \notin P_{d\sim a^m}$ we can conclude, similarly as above, that $\{p_1, ..., p_{|S|-1}, q\}$ is a basis for $\mathbf{R}^S$. Hence, by (8.30) and (8.31) we conclude that (8.28) holds for d and a^m, and hence

$$u(d,p) - u(a^m,p) = \alpha \cdot (v(d,p) - v(a^m,p)) \text{ for all } p \in \Delta(S). \tag{8.32}$$

Now, select an arbitrary $f \in \{a^1, ..., a^{m-1}\}$, different from d. By the induction assumption, we know that (8.28) holds for d and f,and hence

$$u(f,p) - u(d,p) = \alpha \cdot (v(f,p) - v(d,p)) \text{ for all } p \in \Delta(S). \tag{8.33}$$

By (8.32) and (8.33) it follows that

$$u(f,p) - u(a^m,p) = \alpha \cdot (v(f,p) - v(a^m,p)) \text{ for all } p \in \Delta(S).$$

Hence, (8.28) holds for all pairs of acts in $\{a^1, ..., a^m\}$. By induction on m, (8.28) holds for all $d, e \in A$. This completes the proof. ■

8.3 Appendix C: Proof of Section 6

To prove Theorem 6.1 we need the following preparatory result.

Lemma 8.7 (Implication of strong transitivity and the line property) *Let $\succsim$ be a regular conditional preference relation that satisfies strong transitivity and such that no act weakly dominates another act. Moreover, assume that not all indifference sets $P_{a\sim b}$ coincide. Let $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ be a line that intersects every set $\langle P_{a\sim b}\rangle$ at a unique point $v_{ab} = v + \lambda_{ab} w$, such that $\lambda_{ab} \neq \lambda_{ac}$ whenever $P_{a\sim b} \neq P_{a\sim c}$, and*

$$(\lambda_{ab} - \lambda_{bd})(\lambda_{ac} - \lambda_{bc})(\lambda_{ad} - \lambda_{cd}) = (\lambda_{ab} - \lambda_{bc})(\lambda_{ac} - \lambda_{cd})(\lambda_{ad} - \lambda_{bd})$$

for all $a, b, c, d \in A$. Fix an arbitrary act a. Then, for every $b \neq a$ there is a vector $v'_{ab} = v + \lambda'_{ab} w$ on the line L, different from v_{ab}, such that

(a) for every $b, c \neq a$ there is some $\mu \in \mathbf{R}$ with $v_{bc} = (1-\mu)v_{ab} + \mu v'_{ab}$ and $v_{bc} = (1-\mu)v_{ac} + \mu v'_{ac}$,

(b) for every $b \neq a$ the set $H_{ab} = \langle P_{a\sim b}\rangle + \{v'_{ab} - v_{ab}\}$ has a non-empty intersection with $\Delta(S)$, and

(c) for every $b, c \neq a$ and every $p_{bc} \in P_{b\sim c}$ there is some $\mu \in \mathbf{R}$ with

$$p_{bc} \in (1-\mu)\langle P_{a\sim b}\rangle + \mu H_{ab} \text{ and } p_{bc} \in (1-\mu)\langle P_{a\sim c}\rangle + \mu H_{ac}.$$

Proof. Fix an arbitrary act a. Select the acts $b^*, c^* \neq a$ as in Lemma 8.6. Let the line L have the properties stated above. For every $d, e \neq a$ with $\lambda_{ae} \neq \lambda_{de}$, define the number

$$D_{ade} := \frac{\lambda_{ad} - \lambda_{de}}{\lambda_{ae} - \lambda_{de}}. \tag{8.34}$$

As $\lambda_{ae} \neq \lambda_{de}$, it follows by transitivity that $\lambda_{ad}, \lambda_{ae}$ and λ_{de} are pairwise different, and hence $D_{ade} \neq 0$. By the line property we have that

$$D_{adf} = D_{ade} \cdot D_{aef} \tag{8.35}$$

for all $d, e, f \neq a$ with $\lambda_{af} \neq \lambda_{df}, \lambda_{ae} \neq \lambda_{de}$ and $\lambda_{af} \neq \lambda_{ef}$.

We will now define, for every act $d \neq a$, a number λ'_{ad} as follows. Recall that we selected the acts a, b^*, c^* as in Lemma 8.6. Hence, in particular, $P_{a\sim b^*} \neq P_{a\sim c^*}$. As $\succsim$ has preference reversals on $\{a, b^*\}$, there is some belief $p'_{ab^*} \in P_{b^*\succ a}$. Let $p_{ab^*} \in P_{a\sim b^*}$, and define the set $H_{ab^*} := \langle P_{a\sim b^*}\rangle + \{p'_{ab^*} - p_{ab^*}\}$. Since, by Lemma 8.1 (b), $\langle P_{a\sim b^*}\rangle$ is a hyperplane, it follows that H_{ab^*} is a hyperplane as well. As L intersects $\langle P_{a\sim b^*}\rangle$ at a single point, it will intersect H_{ab^*} at a single point as well, say at $v'_{ab^*} = v + \lambda'_{ab^*} w$. Note that $\lambda'_{ab^*} \neq \lambda_{ab^*}$.

As $P_{a\sim b^*} \neq P_{a\sim c^*}$, we conclude that $\lambda_{ab^*} \neq \lambda_{ac^*}$ by the way we have chosen the line L. By transitivity it then follows that $\lambda_{ab^*}, \lambda_{ac^*}$ and $\lambda_{b^*c^*}$ are pairwise different. Let λ'_{ac^*} be the unique number such that

$$\lambda'_{ab^*} - \lambda_{ab^*} = D_{ab^*c^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}). \tag{8.36}$$

Note that $D_{ab^*c^*}$ is not equal to 0 as $\lambda_{ab^*}, \lambda_{ac^*}$ and $\lambda_{b^*c^*}$ are pairwise different.

Next, take an act $d \neq a, b^*, c^*$. Then, by Lemma 8.6 (b), either $\lambda_{ad} \neq \lambda_{ab^*}$ or $\lambda_{ad} \neq \lambda_{ac^*}$. Assume first that $\lambda_{ad} \neq \lambda_{ab^*}$. Then, $\lambda_{ad}, \lambda_{ab^*}$ and λ_{db^*} are pairwise different. Let λ'_{ad} be the unique number such that

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot (\lambda'_{ab^*} - \lambda_{ab^*}). \tag{8.37}$$

If $\lambda_{ad} = \lambda_{ab^*}$ then it must be $\lambda_{ad} \neq \lambda_{ac^*}$, and hence $\lambda_{ad}, \lambda_{ac^*}$ and λ_{dc^*} are pairwise different. Let λ'_{ad} be the unique number such that

$$\lambda'_{ad} - \lambda_{ad} = D_{adc^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}). \tag{8.38}$$

The construction of the numbers λ'_{ad}, for $d \neq a$, is hereby complete.

We now show that, for all $d, e \neq a$ with $\lambda_{ad} \neq \lambda_{ae}$,

$$\lambda'_{ad} - \lambda_{ad} = D_{ade} \cdot (\lambda'_{ae} - \lambda_{ae}). \tag{8.39}$$

In view of (8.36), (8.37) and (8.38) it only remains to show (8.39) for the case where $e = c^*$ and $\lambda_{ad} \neq \lambda_{ab^*}$, and for the case where $d, e \neq b^*, c^*$.

Consider first the case where $e = c^*$ and $\lambda_{ad} \neq \lambda_{ab^*}$. Then we have, by (8.36) and (8.37), that

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot (\lambda'_{ab^*} - \lambda_{ab^*}) \text{ and } \lambda'_{ab^*} - \lambda_{ab^*} = D_{ab^*c^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}),$$

which implies that $\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot D_{ab^*c^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*})$. As, by (8.35), $D_{adb^*} \cdot D_{ab^*c^*} = D_{adc^*}$, we obtain that $\lambda'_{ad} - \lambda_{ad} = D_{adc^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*})$, which was to show.

Suppose next that $d, e \neq b^*, c^*$. If $\lambda_{ad} \neq \lambda_{ab^*}$ and $\lambda_{ae} \neq \lambda_{ab^*}$, then it follows from (8.37) that

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot (\lambda'_{ab^*} - \lambda_{ab^*}) \text{ and } \lambda'_{ae} - \lambda_{ae} = D_{aeb^*} \cdot (\lambda'_{ab^*} - \lambda_{ab^*})$$

and hence $\lambda'_{ad} - \lambda_{ad} = (D_{adb^*}/D_{aeb^*}) \cdot (\lambda'_{ae} - \lambda_{ae})$. As, by definition, $D_{ab^*e} = D^{-1}_{aeb^*}$, it follows that

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot D_{ab^*e} \cdot (\lambda'_{ae} - \lambda_{ae}) = D_{ade} \cdot (\lambda'_{ae} - \lambda_{ae})$$

since, by (8.35), $D_{adb^*} \cdot D_{ab^*e} = D_{ade}$.

If $\lambda_{ad} \neq \lambda_{ab^*}$ and $\lambda_{ae} = \lambda_{ab^*}$, then it follows from (8.37) and (8.38) that

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot (\lambda'_{ab^*} - \lambda_{ab^*}) \text{ and } \lambda'_{ae} - \lambda_{ae} = D_{aec^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}).$$

Combined with (8.36) we get

$$\lambda'_{ad} - \lambda_{ad} = D_{adb^*} \cdot D_{ab^*c^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}) \text{ and } \lambda'_{ae} - \lambda_{ae} = D_{aec^*} \cdot (\lambda'_{ac^*} - \lambda_{ac^*}),$$

and hence $\lambda'_{ad} - \lambda_{ad} = (D_{adb^*} D_{ab^*c^*}/D_{aec^*})(\lambda'_{ae} - \lambda_{ae})$. As, by (8.35), $D_{adb^*} \cdot D_{ab^*c^*} = D_{adc^*}$, and $D_{adc^*}/D_{aec^*} = D_{adc^*} \cdot D_{ac^*e} = D_{ade}$, it follows that $\lambda'_{ad} - \lambda_{ad} = D_{ade} \cdot (\lambda'_{ae} - \lambda_{ae})$.

The case where $\lambda_{ad} = \lambda_{ab^*}$ and $\lambda_{ae} \neq \lambda_{ab^*}$, and the case where $\lambda_{ad} = \lambda_{ab^*}$ and $\lambda_{ae} = \lambda_{ab^*}$ can be shown in a similar fashion as above. We have thus established (8.39) for every $d, e \neq a$ with $\lambda_{ad} \neq \lambda_{ae}$.

For every $b \neq a$, define the vector $v'_{ab} := v + \lambda_{ab} w$. We will prove properties (a), (b) and (c).

(a) We will show that for every $b, c \neq a$ there is some μ with

$$v_{bc} = (1-\mu)v_{ab} + \mu v'_{ab} \text{ and } v_{bc} = (1-\mu)v_{ac} + \mu v'_{ac}. \tag{8.40}$$

To prove this, assume first that $\lambda_{ab} \neq \lambda_{ac}$. Let μ be such that $v_{bc} = (1-\mu)v_{ab} + \mu v'_{ab}$. Then, $\lambda_{bc} = (1-\mu)\lambda_{ab} + \mu \lambda'_{ab}$, which implies that $\lambda_{ab} - \lambda_{bc} = \mu(\lambda_{ab} - \lambda'_{ab})$. Together with (8.34) and (8.39), it follows that

$$D_{abc}(\lambda_{ac} - \lambda_{bc}) = \lambda_{ab} - \lambda_{bc} = \mu(\lambda_{ab} - \lambda'_{ab}) = \mu D_{abc}(\lambda_{ac} - \lambda'_{ac}).$$

As $D_{abc} \neq 0$, this implies that $\lambda_{ac} - \lambda_{bc} = \mu(\lambda_{ac} - \lambda'_{ac})$, which yields $v_{bc} = (1-\mu)v_{ac} + \mu v'_{ac}$. Hence, (8.40) is established.

If $\lambda_{ab} = \lambda_{ac}$ then, by the properties of line L, we have that $P_{a\sim b} = P_{a\sim c}$, which implies that $P_{a\sim b} = P_{a\sim c} = P_{b\sim c}$. Hence, $v_{ab} = v_{ac} = v_{bc}$, and thus (8.40) trivially holds for $\mu = 0$. The proof for (a) is hereby complete.

(b) Note that by choosing p'_{ab^*} closer to $P_{a\sim b^*}$, we move the vector v'_{ab} closer to v_{ab} for every $b \neq a$. As, by Lemma 8.1 (a) and (b), $\langle P_{a\sim b}\rangle \cap \Delta(S) = P_{a\sim b}$ and $P_{a\sim b}$ contains a full support belief, we can choose p'_{ab^*} close enough to $P_{a\sim b^*}$ such that $H_{ab} = \langle P_{a\sim b}\rangle + \{v'_{ab} - v_{ab}\}$ has a non-empty intersection with $\Delta(S)$ for all $b \neq a$. This establishes (b).

(c) Since, by Lemma 8.1 (b), $\langle P_{a\sim b}\rangle$ is a hyperplane, it follows that $H_{ab} = \langle P_{a\sim b}\rangle + \{v'_{ab} - v_{ab}\}$ is a hyperplane as well. As the hyperplanes $\langle P_{a\sim b}\rangle$ and H_{ab} are parallel, there is for every $v \in \mathbf{R}^S$ and for every $b \neq a$ a unique number $\mu^{ab}(v)$ such that $v \in (1-\mu^{ab}(v)) \langle P_{a\sim b}\rangle + \mu^{ab}(v) H_{ab}$. Moreover, similarly to Lemma 8.3, the coordinate $\mu^{ab}(v)$ is linear in v.

We show that for all $b, c \neq a$,

$$\mu^{ab}(p) = \mu^{ac}(p) \text{ for every } p \in P_{b\sim c}. \tag{8.41}$$

Suppose first that $\lambda_{ab} = \lambda_{ac}$. By construction of the line L it must then be that $P_{a\sim b} = P_{a\sim c}$, and hence, by transitivity, $P_{a\sim b} = P_{a\sim c} = P_{b\sim c}$. Thus, $\mu^{ab}(p) = \mu^{ac}(p) = 0$ for every $p \in P_{b\sim c}$ and (8.41) follows.

Assume now that $\lambda_{ab} \neq \lambda_{ac}$. Then, by transitivity, $\lambda_{ab}, \lambda_{ac}$ and λ_{bc} are pairwise different, and thus $P_{a\sim b}, P_{a\sim c}$ and $P_{b\sim c}$ are pairwise different. Moreover, $v_{bc} \notin \langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$ by the choice of the line L.

As $P_{a\sim b}$ and $P_{a\sim c}$ are different, it follows by Lemma 8.1 (a) that $\langle P_{a\sim b}\rangle$ and $\langle P_{a\sim c}\rangle$ are different. As, by Lemma 8.1 (b), $\langle P_{a\sim b}\rangle$ and $\langle P_{a\sim c}\rangle$ both have dimension $|S|-1$, the linear space $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$ has dimension $|S|-2$. Hence, there is a basis $\{v_1, ..., v_{|S|-2}\}$ for $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$. By strong transitivity, $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle \subseteq \langle P_{b\sim c}\rangle$, and hence $\{v_1, ..., v_{|S|-2}\} \subseteq \langle P_{b\sim c}\rangle$. Since we have seen that $v_{bc} \notin \langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$ and $v_{bc} \in \langle P_{b\sim c}\rangle$, it follows that $\{v_1, ..., v_{|S|-2}, v_{bc}\}$ is a basis for $\langle P_{b\sim c}\rangle$.

We will now show that $\mu^{ab}(v) = \mu^{ac}(v)$ for every v in this basis. By (8.40) and the fact that $v_{ab} \in \langle P_{a\sim b}\rangle, v'_{ab} \in H_{ab}, v_{ac} \in \langle P_{a\sim c}\rangle, v'_{ac} \in H_{ac}$, it follows that $\mu^{ab}(v_{bc}) = \mu^{ac}(v_{bc})$. Moreover, for every $v \in \{v_1, ..., v_{|S|-2}\} \subseteq \langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$ it holds that $\mu^{ab}(v) = \mu^{ac}(v) = 0$. Thus, $\mu^{ab}(v) = \mu^{ac}(v)$ for every $v \in \{v_1, ..., v_{|S|-2}, v_{bc}\}$. As $\mu^{ab}(v)$ is linear in v, and $\{v_1, ..., v_{|S|-2}, v_{bc}\}$ is a basis for $\langle P_{b\sim c}\rangle$, it follows that $\mu^{ab}(v) = \mu^{ac}(v)$ for every $v \in \langle P_{b\sim c}\rangle$. In particular, (8.41) follows. This completes the proof. ■

Proof of Theorem 6.1. (a) Suppose there is a utility function u that represents $\succsim$. To show strong transitivity, consider three acts a, b and c. As $\succsim$ has preference reversals on every pair of acts, it follows from Lemma 8.1 (a) and (b) that $\langle P_{a\sim b}\rangle$ is a hyperplane and $P_{a\sim b} = \langle P_{a\sim b}\rangle \cap \Delta(S)$, and similarly for $P_{a\sim c}$ and $P_{b\sim c}$. Consequently,

$$\langle P_{a\sim b}\rangle = \{v \in \mathbf{R}^S \mid u(a,v) = u(b,v)\},\ \ \langle P_{a\sim c}\rangle = \{v \in \mathbf{R}^S \mid u(a,v) = u(c,v)\} \text{ and}$$
$$\langle P_{b\sim c}\rangle = \{v \in \mathbf{R}^S \mid u(b,v) = u(c,v)\},$$

which immediately implies that $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle \subseteq \langle P_{b\sim c}\rangle$.

To show the line property, note first that we can always find a line $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ that intersects each of the hyperplanes $\langle P_{e\sim f}\rangle$ at a single point $v_{ef} = v + \lambda_{ef} w$, and such that $\lambda_{ef} \neq \lambda_{eg}$ whenever $P_{e\sim f} \neq P_{e\sim g}$. To see this, select a vector w such that $w \notin \langle P_{e\sim f}\rangle$ for every $e, f \in A$. That is, w is not parallel to any of the hyperplanes $\langle P_{e\sim f}\rangle$. Then, the line $\{\lambda w \mid \lambda \in \mathbf{R}\}$ will intersect each of the hyperplanes $\langle P_{e\sim f}\rangle$ exactly once. We can then choose the vector v such that the line $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ intersects each of the hyperplanes $\langle P_{e\sim f}\rangle$ at a single point $v_{ef} = v + \lambda_{ef} w$, and such that $\lambda_{ef} \neq \lambda_{eg}$ whenever $P_{e\sim f} \neq P_{e\sim g}$.

Hence, $u(e, v_{ef}) = u(f, v_{ef})$ for all $e, f \in \{a, b, c, d\}$. We will show that

$$(\lambda_{ab} - \lambda_{bd})(\lambda_{ac} - \lambda_{bc})(\lambda_{ad} - \lambda_{cd}) = (\lambda_{ab} - \lambda_{bc})(\lambda_{ac} - \lambda_{cd})(\lambda_{ad} - \lambda_{bd}). \tag{8.42}$$

Assume first that $\lambda_{ef} = \lambda_{eg}$ for some $e, f, g \in \{a, b, c, d\}$. Then, by transitivity of $\succsim$, we have that $\lambda_{ef} = \lambda_{eg} = \lambda_{fg}$, and (8.42) trivially holds.

Suppose next that $\lambda_{ef} \neq \lambda_{eg}$ for all $e, f, g \in \{a, b, c, d\}$. Define the affine mappings δ_{ab}, δ_{ac} and δ_{ad} from $\mathbf{R}$ to $\mathbf{R}$ by

$$\delta_{ab}(\lambda) := u(a, v + \lambda w) - u(b, v + \lambda w),\ \delta_{ac}(\lambda) := u(a, v + \lambda w) - u(c, v + \lambda w) \text{ and} \\ \delta_{ad}(\lambda) := u(a, v + \lambda w) - u(d, v + \lambda w) \text{ for all } \lambda \in \mathbf{R}. \tag{8.43}$$

Moreover, these mappings are nonconstant as $\delta_{ab}(\lambda_{ab}) = 0$ and $\delta_{ab}(\lambda) \neq 0$ for all $\lambda \neq \lambda_{ab}$, and similarly for δ_{ac} and δ_{ad}. As these mappings are affine and nonconstant, there are nonzero numbers D_{ab}, D_{ac} and D_{ad} such that

$$\delta_{ab}(\lambda) - \delta_{ab}(\mu) = D_{ab} \cdot (\lambda - \mu) \tag{8.44}$$

$$\delta_{ac}(\lambda) - \delta_{ac}(\mu) = D_{ac} \cdot (\lambda - \mu) \text{ and} \tag{8.45}$$

$$\delta_{ad}(\lambda) - \delta_{ad}(\mu) = D_{ad} \cdot (\lambda - \mu) \tag{8.46}$$

for all $\lambda, \mu \in \mathbf{R}$. We will now show that

$$\frac{D_{ab}}{D_{ac}} = \frac{\lambda_{ac} - \lambda_{bc}}{\lambda_{ab} - \lambda_{bc}}. \tag{8.47}$$

By taking $\lambda = \lambda_{ab}$ and $\mu = \lambda_{bc}$, we obtain from (8.44) that

$$D_{ab} = \frac{\delta_{ab}(\lambda_{ab}) - \delta_{ab}(\lambda_{bc})}{\lambda_{ab} - \lambda_{bc}} = -\frac{\delta_{ab}(\lambda_{bc})}{\lambda_{ab} - \lambda_{bc}} \tag{8.48}$$

since $\delta_{ab}(\lambda_{ab}) = 0$. Similarly, by taking $\lambda = \lambda_{ac}$ and $\mu = \lambda_{bc}$, we obtain from (8.45) that

$$D_{ac} = \frac{\delta_{ac}(\lambda_{ac}) - \delta_{ac}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} = -\frac{\delta_{ac}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} = -\frac{\delta_{ab}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} \tag{8.49}$$

since $\delta_{ac}(\lambda_{ac}) = 0$ and $\delta_{ac}(\lambda_{bc}) = \delta_{ab}(\lambda_{bc})$. By combining (8.48) and (8.49) we obtain (8.47).

In a similar fashion it can be shown that

$$\frac{D_{ac}}{D_{ad}} = \frac{\lambda_{ad} - \lambda_{cd}}{\lambda_{ac} - \lambda_{cd}} \text{ and } \frac{D_{ab}}{D_{ad}} = \frac{\lambda_{ad} - \lambda_{bd}}{\lambda_{ab} - \lambda_{bd}}. \tag{8.50}$$

As $D_{ab}/D_{ad} = (D_{ab}/D_{ac}) \cdot (D_{ac}/D_{ad})$, equation (8.42) follows from (8.47) and (8.50). Hence, $\succsim$ satisfies strong transitivity and the line property, which was to show.

(b) Assume now that $\succsim$ is regular, and satisfies strong transitivity and the line property. We construct a utility function u, distinguishing two cases: (1) not all indifference sets $P_{a\sim b}$ are identical, and (2) all indifference sets $P_{a\sim b}$ are identical.

Case 1. Suppose that not all indifference sets $P_{a\sim b}$ are identical. Since $\succsim$ satisfies the line property, there is a line $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ that intersects every set $\langle P_{a\sim b}\rangle$ at a unique point $v_{ab} = v + \lambda_{ab} w$, such that $\lambda_{ab} \neq \lambda_{ac}$ whenever $P_{a\sim b} \neq P_{a\sim c}$, and

$$(\lambda_{ab} - \lambda_{bd})(\lambda_{ac} - \lambda_{bc})(\lambda_{ad} - \lambda_{cd}) = (\lambda_{ab} - \lambda_{bc})(\lambda_{ac} - \lambda_{cd})(\lambda_{ad} - \lambda_{bd})$$

for all $a, b, c, d \in A$.

Fix an arbitrary act a. Then, by Lemma 8.7 there is for every $b \neq a$ a vector $v'_{ab} = v + \lambda'_{ab} w$ on the line L, different from v_{ab}, such that for every $b \neq a$ the set $H_{ab} = \langle P_{a\sim b} \rangle + \{v'_{ab} - v_{ab}\}$ has a non-empty intersection with $\Delta(S)$, and such that for every $b, c \neq a$ and every $p_{bc} \in P_{b\sim c}$ there is some $\mu \in \mathbf{R}$ with

$$p_{bc} \in (1-\mu) \langle P_{a\sim b} \rangle + \mu H_{ab} \text{ and } p_{bc} \in (1-\mu) \langle P_{a\sim c} \rangle + \mu H_{ac}. \tag{8.51}$$

We now define the utility function u as follows. Choose $b^*, c^* \neq a$ as in Lemma 8.6. Set $\alpha > 0$ if $H_{ab^*} \cap \Delta(S) \subseteq P_{b^* \succ a}$, and choose $\alpha < 0$ if $H_{ab^*} \cap \Delta(S) \subseteq P_{a \succ b^*}$. Define $u(a, s)$ arbitrarily for all $s \in S$.

Consider now some $d \neq a$. By Lemma 8.1 (b) there is a basis $\{p_1, ..., p_{|S|-1}\}$ for $\langle P_{a\sim d} \rangle$. As v'_{ad} is on the line L, the line L intersects $\langle P_{a\sim d} \rangle$ only at v_{ad}, and $v'_{ad} \neq v_{ad}$, it follows that $v'_{ad} \notin \langle P_{a\sim d} \rangle$. Hence, $\{p_1, ..., p_{|S|-1}, v'_{ad}\}$ is a basis for $\mathbf{R}^S$. Choose the unique utilities $(u(d, s))_{s \in S}$ such that

$$u(d, p_k) := u(a, p_k) \text{ for all } k \in \{1, ..., |S|-1\}, \text{ and } u(d, v'_{ad}) := u(a, v'_{ad}) + \alpha. \tag{8.52}$$

We will show that u represents $\succsim$.

We assume, without loss of generality, that $H_{ab^*} \cap \Delta(S) \subseteq P_{b^* \succ a}$ and hence $\alpha > 0$. The case where $\alpha < 0$ follows similarly, and is therefore omitted. To show that u represents $\succsim$, we proceed in steps.

Step 1. Show that u represents $\succsim$ on $\{a, b^*\}$.
Proof. By (8.52) and the fact that $H_{ab^*} = \langle P_{a\sim b^*} \rangle + \{v'_{ab^*} - v_{ab^*}\}$,

$$\langle P_{a\sim b^*} \rangle = \{v \in \mathbf{R}^S \mid u(b^*, v) = u(a, v)\} \text{ and } H_{ab^*} = \{v \in \mathbf{R}^S \mid u(b^*, v) = u(a, v) + \alpha\},$$

and hence, in particular, $P_{a\sim b^*} = P_{u(a)=u(b^*)}$. As $\alpha > 0$, there is some $p'_{ab^*} \in H_{ab^*} \cap \Delta(S)$ with $p'_{ab^*} \in P_{b^* \succ a}$. Hence, there is some $p'_{ab^*} \in P_{b^* \succ a}$ with $u(b^*, p'_{ab^*}) > u(a, p'_{ab^*})$. Similarly to the proof of Theorem 5.1 it then follows that $P_{b^* \succ a} = P_{u(b^*) > u(a)}$. Hence, u represents $\succsim$ on $\{a, b^*\}$.

Step 2. Show that u represents $\succsim$ on $\{a, c^*\}$.
Proof. Similarly to the proof of Step 1, we have that

$$\langle P_{a\sim c^*} \rangle = \{v \in \mathbf{R}^S \mid u(c^*, v) = u(a, v)\} \text{ and } H_{ac^*} = \{v \in \mathbf{R}^S \mid u(c^*, v) = u(a, v) + \alpha\},$$

which implies that $P_{a\sim c^*} = P_{u(a)=u(c^*)}$. Since $P_{a\sim b^*}$ and $P_{a\sim c^*}$ are different, it follows by transitivity that $P_{b^*\sim c^*} \neq P_{a\sim b^*}$. Hence, there is some $p \in P_{b^*\sim c^*} \backslash P_{a\sim b^*}$. Assume, without loss of generality, that $p \in P_{b^* \succ a}$. Then, by transitivity, $p \in P_{c^* \succ a}$.

By (8.51) there is some μ with $p \in (1-\mu) \langle P_{a\sim b^*} \rangle + \mu H_{ab^*}$ and $p \in (1-\mu) \langle P_{a\sim c^*} \rangle + \mu H_{ac^*}$. As $p \in P_{b^* \succ a}$ and $H_{ab^*} \cap \Delta(S) \subseteq P_{b^* \succ a}$, it follows that $\mu > 0$. Since $p \in (1-\mu) \langle P_{a\sim c^*} \rangle + \mu H_{ac^*}$ with $\mu > 0$ and $p \in P_{c^* \succ a}$, it must thus be that $H_{ac^*} \cap \Delta(S) \subseteq P_{c^* \succ a}$. Hence, there is some $q \in P_{c^* \succ a}$ with $q \in H_{ac^*}$, and therefore $u(c^*, q) > u(a, q)$. Similarly to the proof of Theorem 5.1 it then follows that $P_{c^* \succ a} = P_{u(c^*) > u(a)}$. Hence, u represents $\succsim$ on $\{a, c^*\}$.

Step 3. Show that u represents $\succsim$ on $\{a, d\}$ for every $d \neq a, b^*, c^*$.

Proof. Take some $d \neq a, b^*, c^*$. Then, $P_{a\sim d}$ is either different from $P_{a\sim b^*}$ or different from $P_{a\sim c^*}$. Assume, without loss of generality, that $P_{a\sim d} \neq P_{a\sim b^*}$. Then, it can be shown in the same way as for $\{a, c^*\}$ that u represents $\succsim$ on $\{a, d\}$, by exchanging the roles of d and c^*.

Step 4. Show that u represents $\succsim$ on $\{b^*, c^*\}$.
Proof. Take some $p \in P_{b^*\sim c^*}$. Then, by (8.51), there is some μ with $p \in (1-\mu)\langle P_{a\sim b^*}\rangle + \mu H_{ab^*}$ and $p \in (1-\mu)\langle P_{a\sim c^*}\rangle + \mu H_{ac^*}$. Hence, by (8.52), $u(b^*, p) - u(a, p) = \mu\alpha$ and $u(c^*, p) - u(a, p) = \mu\alpha$, which implies that $u(b^*, p) = u(c^*, p)$. That is, $P_{b^*\sim c^*} \subseteq P_{u(b^*)=u(c^*)}$.

We show, in fact, that $P_{b^*\sim c^*} = P_{u(b^*)=u(c^*)}$. To this purpose, we first prove that $P_{u(b^*)=u(c^*)} \neq \Delta(S)$. Since $P_{a\sim b^*} \neq P_{a\sim c^*}$ we can find some $p \in P_{a\sim b^*} \backslash P_{a\sim c^*}$. As u represents $\succsim$ on $\{a, b^*\}$ and $\{a, c^*\}$ we must have that $u(a, p) = u(b^*, p)$ and $u(a, p) \neq u(c^*, p)$, which implies that $u(b^*, p) \neq u(c^*, p)$. Thus, $P_{u(b^*)=u(c^*)} \neq \Delta(S)$. As $P_{u(b^*)=u(c^*)} = \left\langle P_{u(b^*)=u(c^*)}\right\rangle \cap \Delta(S)$, it follows that $\left\langle P_{u(b^*)=u(c^*)}\right\rangle \neq \mathbf{R}^S$. Hence, $\left\langle P_{u(b^*)=u(c^*)}\right\rangle$ has dimension at most $|S|-1$. Since $P_{b^*\sim c^*} \subseteq P_{u(b^*)=u(c^*)}$ and, by Lemma 8.1 (b), $\langle P_{b^*\sim c^*}\rangle$ has dimension $|S|-1$, it must be that $\langle P_{b^*\sim c^*}\rangle = \left\langle P_{u(b^*)=u(c^*)}\right\rangle$. Since, by Lemma 8.1 (a), $P_{b^*\sim c^*} = \langle P_{b^*\sim c^*}\rangle \cap \Delta(S)$ and we have seen that $P_{u(b^*)=u(c^*)} = \left\langle P_{u(b^*)=u(c^*)}\right\rangle \cap \Delta(S)$, we conclude that $P_{b^*\sim c^*} = P_{u(b^*)=u(c^*)}$.

We have seen that there is some $p \in P_{a\sim b^*} \backslash P_{a\sim c^*}$. Assume, without loss of generality, that $p \in P_{c^*\succ a}$. Then, by transitivity, $p \in P_{c^*\succ b^*}$. Since u represents $\succsim$ on $\{a, b^*\}$ and $\{a, c^*\}$ it must be that $u(a, p) = u(b^*, p)$ and $u(c^*, p) > u(a, p)$, which implies that $u(c^*, p) > u(b^*, p)$. We thus have found some $p \in P_{c^*\succ b^*}$ with $u(c^*, p) > u(b^*, p)$. Similarly to the proof of Theorem 5.1 it then follows that $P_{c^*\succ b^*} = P_{u(c^*)>u(b^*)}$. Hence, u represents $\succsim$ on $\{b^*, c^*\}$.

Step 5. Show that u represents $\succsim$ on $\{b^*, d\}$ and $\{c^*, d\}$ for every $d \neq a, b^*, c^*$.
Proof. Take some $d \neq a, b^*, c^*$. As $P_{a\sim b^*} \neq P_{a\sim c^*}$, we must have that either $P_{a\sim d} \neq P_{a\sim b^*}$ or $P_{a\sim d} \neq P_{a\sim c^*}$. Assume, without loss of generality, that $P_{a\sim d} \neq P_{a\sim b^*}$. As, by Steps 1 and 3, u represents $\succsim$ on $\{a, b^*\}$ and $\{a, d\}$, we can show in the same way as in Step 4 that u represents $\succsim$ on $\{b^*, d\}$.

We now show that u also represents $\succsim$ on $\{c^*, d\}$. If $P_{a\sim d} \neq P_{a\sim c^*}$, then it can be shown in the same was as in Step 4 that u represents $\succsim$ on $\{c^*, d\}$. Assume now that $P_{a\sim d} = P_{a\sim c^*}$. Then, $P_{c^*\sim d} = P_{a\sim d}$. Since $P_{a\sim d} \neq P_{a\sim b^*}$, it follows that $P_{b^*\sim d} \neq P_{a\sim d} = P_{c^*\sim d}$. Hence, $P_{b^*\sim c^*} \neq P_{b^*\sim d}$.

In the same way as in Step 4 it can be shown that $P_{c^*\sim d} \subseteq P_{u(c^*)=u(d)}$. To prove that $P_{u(c^*)=u(d)} \neq \Delta(S)$, recall that $P_{b^*\sim c^*} \neq P_{b^*\sim d}$, and hence there is some $p \in P_{b^*\sim c^*} \backslash P_{b^*\sim d}$. As we have seen that u represents $\succsim$ on $\{b^*, c^*\}$ and $\{b^*, d\}$, it follows in the same way as in Step 4 that $u(c^*, p) \neq u(d, p)$, and hence $P_{u(c^*)=u(d)} \neq \Delta(S)$. Analogously to Step 4, this implies that $P_{c^*\sim d} = P_{u(c^*)=u(d)}$. Moreover, in a similar way as in Step 4, it can then be shown that $P_{c^*\succ d} = P_{u(c^*)>u(d)}$, which implies that u represents $\succsim$ on $\{c^*, d\}$.

Step 6. Show that u represents $\succsim$ on $\{d, e\}$ for every $d, e \neq a, b^*, c^*$.
Proof. Take some $d, e \neq a, b^*, c^*$. Then, by Lemma 8.6 (c), there are three possible cases: (i) $P_{a\sim d} \neq P_{a\sim e}$, (ii) $P_{b^*\sim d} \neq P_{b^*\sim e}$, and (iii) $P_{c^*\sim d} \neq P_{c^*\sim e}$.

Case (i). Suppose that $P_{a\sim d} \neq P_{a\sim e}$. Then, it can be shown in a similar way as in Step 4 that u represents $\succsim$ on $\{d, e\}$.

Case (ii). Suppose that $P_{b^*\sim d} \neq P_{b^*\sim e}$. As we have seen that u represents $\succsim$ on $\{b^*, d\}$ and $\{b^*, e\}$, it can be shown in a similar way as in Step 4 that u represents $\succsim$ on $\{d, e\}$.

Case (iii). Suppose that $P_{c^*\sim d} \neq P_{c^*\sim e}$. As we have seen that u represents $\succsim$ on $\{c^*, d\}$ and $\{c^*, e\}$, it can be shown in a similar way as in Step 4 that u represents $\succsim$ on $\{d, e\}$.

By Steps 1–6, we conclude that u represents $\succsim$ *on* $\{d, e\}$ for every $d, e \in A$, and hence u represents $\succsim$. This completes the proof of Case 1.

Case 2. Suppose that all indifference sets $P_{a\sim b}$ are equal. That is, there is some linear space H with dimension $|S| - 1$ such that $P_{a\sim b} = H \cap \Delta(S)$ for every $a, b \in A$. Let $\{v_1, ..., v_{|S|-1}\}$ be a basis for H. Take some $p^* \in \Delta(S)\backslash H$. Then, $\{v_1, ..., v_{|S|-1}, p^*\}$ is a basis for $\mathbf{R}^S$. As $p^* \notin H$, there must be a strict ordering of the acts at p^*. Let $c^1, c^2, ..., c^{|A|}$ be a numbering of the acts such that

$$c^1 \succ_{p^*} c^2 \succ_{p^*} ... \succ_{p^*} c^{|A|}. \tag{8.53}$$

Construct a utility function u such that

$$u(a, v_k) = u(b, v_k) \text{ for all } k \in \{1, ..., |S| - 1\}, \text{ and } u(c^1, p^*) > u(c^2, p^*) > ... > u(c^{|A|}, p^*). \tag{8.54}$$

To show that u represents $\succsim$, take two acts a, b and assume, without loss of generality, that $p^* \in P_{a\succ b}$. Then, by (8.54), $P_{u(a)=u(b)} = H \cap \Delta(S) = P_{a\sim b}$. Moreover, $p^* \in P_{u(a)>u(b)}$. As $p^* \in P_{a\succ b}$ it then follows, similarly to the proof of Theorem 5.1, that $P_{u(a)>u(b)} = P_{a\succ b}$. Thus, u represents $\succsim$ on $\{a, b\}$. As a, b were arbitrary, we conclude that u represents $\succsim$. This completes the proof. ■

References

[1] Anscombe, F.J. and R.J. Aumann (1963), A definition of subjective probability, *Annals of Mathematical Statistics* **34,** 199–205.

[2] Aumann, R.J. and A. Brandenburger (1995), Epistemic conditions for Nash equilibrium, *Econometrica* **63,** 1161–1180.

[3] Aumann, R.J. and J.H. Drèze (2002), Bayes rationality in games, Working paper.

[4] Aumann, R.J. and J.H. Drèze (2005), Assessing strategic risk, Discussion paper 2005-20, Département des Sciences Économiques de l'Université Catholique de Louvain.

[5] Baccelli, J. and P. Mongin (2016), Choice-based cardinal utility: A tribute to Patrick Suppes, *Journal of Economic Methodology* **23,** 268–288.

[6] Baumol, W.J. (1958), The cardinal utility which is ordinal, *The Economic Journal* **68,** 665–672.

[7] Bernheim, B.D. (1984), Rationalizable strategic behavior, *Econometrica* **52,** 1007–1028.

[8] Brandenburger, A. and E. Dekel (1987), Rationalizability and correlated equilibria, *Econometrica* **55,** 1391–1402.

[9] Burghart, D.R. (2020), The two faces of independence: betweenness and homotheticity, *Theory and Decision* **88,** 567–593.

[10] Fishburn, P.C. (1976), Axioms for expected utility in n-person games, *International Journal of Game Theory* **5,** 137–149.

[11] Fishburn, P.C. and F.S. Roberts (1978), Mixture axioms in linear and multilinear utility theories, *Theory and Decision* **9,** 161–171.

[12] Gilboa, I. and D. Schmeidler (1995), Case-based decision theory, *The Quarterly Journal of Economics* **110,** 605–639.

[13] Gilboa, I. and D. Schmeidler (2003), A derivation of expected utility maximization in the context of a game, *Games and Economic Behavior* **44,** 184–194.

[14] Karni, E., Schmeidler, D. and K. Vind (1983), On state dependent preferences and subjective probabilities, *Econometrica* **51,** 1021–1031.

[15] Mariotti, M. (1995), Is Bayesian rationality compatible with strategic rationality?, *The Economic Journal* **105,** 1099–1109.

[16] Moscati, I. (2018), *Measuring Utility: From the Marginal Revolution to Behavioral Economics,* Oxford University Press, Oxford, UK.

[17] Nash, J.F. (1950), Equilibrium points in N-person games, *Proceedings of the National Academy of Sciences of the United States of America* **36,** 48–49.

[18] Nash, J.F. (1951), Non-cooperative games, *Annals of Mathematics* **54**, 286–295.

[19] Pearce, D.G. (1984), Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52,** 1029–1050.

[20] Savage, L.J. (1954), *The Foundation of Statistics,* Wiley, New York.

[21] von Neumann, J. and O. Morgenstern (1947), *Theory of Games and Economic Behavior*, 2nd edition, Princeton University Press, Princeton, NJ.

[22] Wakker, P.P. (1987), Subjective probabilities for state dependent continuous utility, *Mathematical Social Sciences* **14,** 289–298.

[23] Wakker, P.P. (1989), *Additive Representations of Preferences*, Kluwer Academic Publishers, Dordrecht.