A Foundation for Expected Utility in Decision Problems and Games

EPICENTER Working Paper No. 22 (2020)

Andrés Perea*
Maastricht University

June 2020

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, like Gilboa and Schmeidler (2003), 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 it singles out precisely those conditional preference relations that admit an expected utility representation. The key axiom is the *existence of a uniform preference increase*, stating that there must be an alternative conditional preference relation that, for a given choice, uniformly increases the preference for that choice by a constant degree. 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 is equivalent to two easily verifiable conditions: *strong transitivity* and the *line property.*

JEL Classification: C72, D81

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

*Address: EpiCenter and Dept. of Quantitative Economics, 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? That is the key question we wish to explore in this paper. It is often argued that such utilities may be derived from the frameworks developed by Savage (1954) and Anscombe and Aumann (1963). 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 (i.e., opponents' choice combination) some consequence. In that sense, a game can be embedded into the models proposed by Savage and Anscombe and Aumann. Both frameworks provide axiomatic foundations 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. In view of these theorems, 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 or Anscombe and Aumann.

In my view there are at least two problems with this approach. First, both Savage and Anscombe and Aumann assume 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 and Anscombe and Aumann 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 the models of Savage and Anscombe and Aumann, 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 decion 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 and Anscombe and Aumann. 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. This result can thus be viewed as a possible answer to the question what it means, in a decision problem or game, to specify the person's utilities but not his belief. 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.

The first five axioms, *completeness, transitivity, continuity*, *preservation of indifference* and *preservation of strict preference*, which also appear in Gilboa and Schmeidler (2003), may be viewed as basic regularity conditions. The last axiom, *existence of a uniform preference increase,* is new. It states that from a given conditional preference relation, one should always be able to uniformly increase the "degree of preference" for at least one of the choices by a fixed amount. This axiom substitutes, in a sense, the *diversity* axiom in Gilboa and Schmeidler (2003), which states that for every strict ordering of at most four choices, there must be a belief that induces precisely that ordering. Gilboa and Schmeidler show that diversity, together with the regularity axioms, singles out precisely those conditional preference relations that can be represented by a *diversified* utility matrix, that is, a utility matrix where no row is weakly dominated by, or equivalent to, an affine combination of at most three other rows. Our characterization result in Theorem 5.1, in turn, applies to *all* possible utility representations, also those that violate diversity.

We subsequently zoom in on an important special case: the scenario when no choice is weakly dominated by another choice. In fact, every decision problem can be reduced to a problem in this category by eliminating all choices that are weakly dominated. As we may reasonably expect a rational decision maker not to make weakly dominated choices, this class of decision problems may be viewed as "canonical". For this class we show that the existence of a uniform preference increase is equivalent to checking two easily verifiable conditions: *strong transitivity* and the *line property.* The first condition 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. See Figure 9 for an illustration. 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.

This paper is organized as follows. In Section 2 we present the mathematical definitions and results that are required for our analysis. 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 sufficient to characterize the conditional preference relations that admit an expected utility representation. In Section 4 we show that the regularity axioms will no longer suffice if we have more than two choices, and lay out the key axiom of this paper: existence of a uniform preference increase. In Section 5 we present, and prove, the main result in this paper, stating that the new axiom, together with the regularity axioms, characterizes precisely those conditional preference relations that have an expected utility representation. In that section we also describe a procedure that can be used to construct a utility function for a given conditional preference relation. In Section 6 we investigate the special case when there are no weakly dominated choices. We conclude by a discussion in Section 7. All technical and longer proofs can be found in the appendix.

2 Mathematical Tools

In this section we introduce the mathematical definitions and tools needed for this paper, mainly based on well-known definitions and results from linear algebra.

2.1 Linear Spaces

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

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

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 $\dim(V) = 0$.

2.2 Affine Spaces

A subset $V \subseteq \mathbf{R}^X$ is called an *affine subspace* of $\mathbf{R}^X$ if for every $v, w \in V$ and every $\alpha \in \mathbf{R}$ we have that $(1-\alpha)v + \alpha w \in V$. It is well-known that for every affine subspace V there is a linear subspace V' such that $V = V' + \{v\}$ for every $v \in V$. We call V a *hyperplane* if $\dim(V') = |X| - 1$, and we call V a *line* if $\dim(V') = 1$. It is well-known that for every hyperplane V there is a vector $n \neq \underline{0}$ and a number α such that $V = \{v \in \mathbf{R}^X \mid v \cdot n = \alpha\}$. Two hyperplanes V and W are called *parallel* if there is a linear subspace V' and vectors v, w such that $V = V' + \{v\}$ and $W = V' + \{w\}$.

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

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

the set of all (finite) *affine combinations* of elements in V, and call it the *affine span* of V. Here, $\sum_{k=1}^K \alpha_k v_k$ with $\sum_{k=1}^K \alpha_k = 1$ is called an *affine combination* of the vectors $v_1, ..., v_K$. The affine span $\langle V\rangle_a$ is always an affine subspace, and if V itself is an affine subspace then $\langle V\rangle_a = V$. Vectors $v_1, ..., v_K \in \mathbf{R}^X$ are called *affinely independent* if none of the vectors is an affine combination of the other vectors. Geometrically, three vectors are affinely independent if they are not all situated on the same line, whereas four vectors are affinely independent if they are not all situated on the same plane. Consider an affine subspace V of $\mathbf{R}^X$, and vectors $v_1, ..., v_K \in V$. The set of vectors $\{v_1, ..., v_K\}$ is an *affine basis* for V if $v_1, ..., v_K$ are affinely independent, and $\langle\{v_1, ..., v_K\}\rangle_a = V$. It is well-known that if $V = V' + \{v\}$, where V' is a linear subspace, then every affine basis of V will contain $\dim(V') + 1$ elements.

A mapping $f : \mathbf{R}^X \to \mathbf{R}$ is called *affine* if for every $v, w \in \mathbf{R}^X$ and every $\alpha \in \mathbf{R}$ it holds that $f((1-\alpha)v + \alpha w) = (1-\alpha)f(v) + \alpha f(w)$. For every affine mapping f there are numbers β_0 and β_x for every $x \in X$ such that $f(v) = \beta_0 + \sum_{x \in X} v(x)\beta_x$ for all $v \in \mathbf{R}^X$. The following three results will play an important role in our analysis later on.

Lemma 2.1 (Hyperplanes and affine mappings) *(a) For every vector $v \in \mathbf{R}^X$, and every two different, parallel hyperplanes V and W, there is a unique number λ such that $v \in (1-\lambda)V + \lambda W$.*

(b) For every two different, parallel hyperplanes V and W we can find vectors $v_1, ..., v_{|X|} \in V$ and a vector $v_{|X|+1} \in W$ such that $v_1, ..., v_{|X|}, v_{|X|+1}$ are affinely independent.

(c) For all affinely independent vectors $v_1, ..., v_{|X|+1} \in \mathbf{R}^X$ and all numbers $\alpha_1, ..., \alpha_{|X|+1}$ there is a unique affine mapping $f : \mathbf{R}^X \to \mathbf{R}$ such that $f(v_k) = \alpha_k$ for all $k \in \{1, ..., |X|+1\}$.

Property (a) thus states that every vector in $\mathbf{R}^X$ can be written as the affine combination of a vector in V and a vector in W. Property (b), in turn, guarantees that every affine basis of V can be extended to an affine basis of $\mathbf{R}^X$ by selecting an additional element from W. If X contains n elements then, geometrically speaking, property (c) states that for every $n+1$ points in $\mathbf{R}^{n+1}$ there is a unique hyperplane in $\mathbf{R}^{n+1}$ that passes through these points.

2.3 Probability Distributions

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 $y \in X \backslash \{x\}$. By $\Delta^+(X) = \{p \in \Delta(X) \mid p(x) > 0 \text{ for all } x \in X\}$ we denote the set of *full support* probability distributions.

3 Characterization for 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. Consider a decision maker (DM) who must choose from a finite set of choices C. The final outcome depends not only on the choice $c \in C$, but also on the realization of a state $x \in X$ from a finite set of states X. We assume that the decision maker first forms a probabilistic belief p on X, which then induces a preference relation $\succsim_p$ on C. Formally, a *preference relation* $\succsim_p$ on C is a binary relation $\succsim_p \subseteq C \times C$. If $(a, b) \in \succsim_p$ we write $a \succsim_p b$, and the interpretation is that the DM weakly prefers choice a to choice b if his belief is p.

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

In a game, the DM would be one of the players, C would be his set of actions in the game, and X the set of opponents' choice combinations. For two choices 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 choices $a, b \in C$ we define the sets of beliefs $P_{a \sim b} := \{p \in \Delta(X) \mid a \sim_p b\}$, $P_{a \succ b} := \{p \in \Delta(X) \mid a \succ_p b\}$ and $P_{a \succsim b} := \{p \in \Delta(X) \mid a \succsim_p b\}$. Similarly, we define the sets of states $X_{a \sim b} := \{x \in X \mid a \sim_{[x]} b\}$, $X_{a \succ b} := \{x \in X \mid a \succ_{[x]} b\}$ and $X_{a \succsim b} := \{x \in X \mid a \succsim_{[x]} b\}$. We say that (a) *a strictly dominates* b under

$\succcurlyeq$ if $a \succ_p b$ for all $p \in \Delta(X)$; (b) a *weakly dominates* b under $\succcurlyeq$ if $a \succcurlyeq_p b$ for all $p \in \Delta(X)$, and $a \succ_p b$ for at least one $p \in \Delta(X)$; (c) a is *equivalent* to b under $\succcurlyeq$ if $a \sim_p b$ for all $p \in \Delta(X)$; and (d) $\succcurlyeq$ 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, $\succcurlyeq$ either exhibits weak dominance, equivalence, or preference reversals on $\{a, b\}$.

3.2 Regularity Axioms

We will now impose some very basic axioms on conditional preference relations, to which we refer as *regularity axioms.* Later, we will show that these axioms are sufficient to characterize, for any decision problem with *two choices,* those conditional preference relations that have an expected utility representation. However, as we will see in the next section, these are not sufficient for scenarios with more than two choices.

Consider a conditional preference relation $\succcurlyeq$ on (C, X). The first axiom, *completeness,* states that any two choices can always be ranked for every possible belief.

Axiom 3.1 (Completeness) *For every belief p and any two choices $a, b \in C$, either $a \succcurlyeq_p b$ or $b \succcurlyeq_p a$.*

The second axiom, *transitivity,* states that for every three choices a, b and c, the preference between a and c should not contradict the preference between a and b and the preference between b and c.

Axiom 3.2 (Transitivity) *For every belief $p \in \Delta(X)$ and every three choices $a, b, c \in C$ with $a \succcurlyeq_p b$ and $b \succcurlyeq_p c$, it holds that $a \succcurlyeq_p c$..*

Of course, this axiom only imposes restrictions if there are more than two choices, and hence it will be redundant for our characterization theorem for two choices. Nevertheless, we state the axiom already as we interpret it as a basic regularity axiom. Completeness and transitivity together resemble the *ranking* axiom in Gilboa and Scmeidler (2003).

The following three axioms, which also appear in Gilboa and Schmeidler (2003), are all based on the informal principle that if we move from a belief p to another belief q on a line, then the "degree" by which the DM prefers a to b (or b to a) will change linearly. Suppose first that $p \in P_{a \succ b}$ and $q \in P_{b \succ a}$. Then, the degree by which the DM prefers a to b will linearly decrease when moving on a line from p to q. As such, there must be some belief r on this line where the DM is indifferent between a and b. This property is called *continuity.*

Axiom 3.3 (Continuity) *For every two different choices $a, b \in C$ 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 $(1 - \lambda)p + \lambda q \in P_{a \sim b}$.*

Suppose now that $p \in P_{a \sim b}$ and $q \in P_{a \sim b}$. Then, by the principle above, the DM must remain indifferent between a and b when moving from p to q on a line. This property is called *preservation of indifference.*

Axiom 3.4 (Preservation of indifference) *For every two different choices $a, b \in C$ 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 $\lambda \in (0, 1)$.*

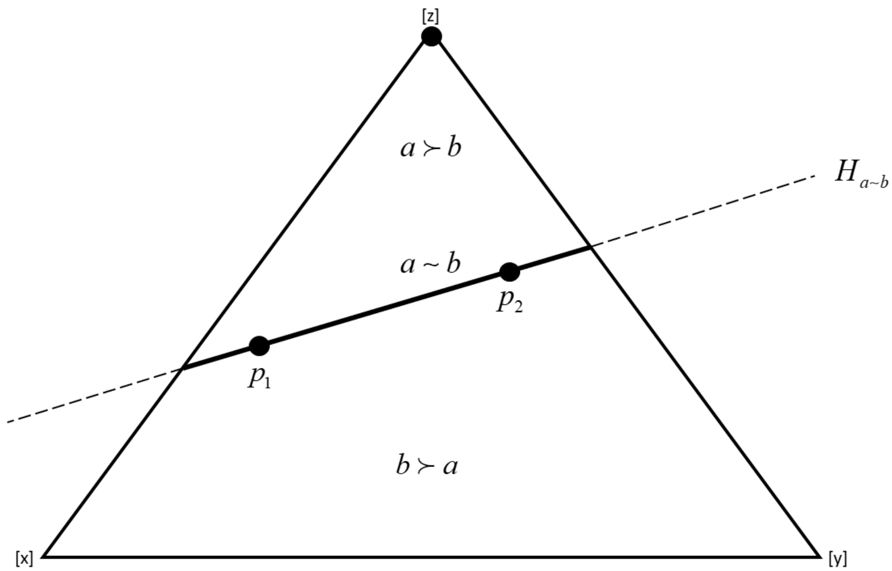

Figure 1: A typical regular conditional preference relation

Assume next that $p \in P_{a \succcurlyeq b}$ and $q \in P_{a \succ b}$. Then, there are two possibilities, $p \in P_{a \sim b}$ or $p \in P_{a \succ b}$. In the first case, the degree by the DM prefers a to b will strictly increase when moving from p to q on a line. In the second case, this degree will either linearly increase or decrease, yet the DM will always strictly prefer a to b when moving from p to q on a line. This property is called *preservation of strict preference.*

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

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 $\succcurlyeq$ with two choices a and b, and three states x, y and z. The area within the triangle represents the set $\Delta(X)$ of all probabilistic beliefs on $X = \{x, y, z\}$, with the probability 1 beliefs $[x], [y]$ and $[z]$ as the extreme points. The two-dimensional plane represents all the vectors in $\mathbf{R}^X$ where the sum of the coordinates is 1, containing the belief simplex $\Delta(X)$ 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 $\succcurlyeq$ satisfies all the regularity axioms.

3.3 Characterization Theorem for Two Choices

Before we state our characterization result, we first formally define what it means for a conditional preference relation to have an *expected utility representation.*

Definition 3.2 (Expected utility representation) *Consider a finite set of choices C and a finite set of states X. A conditional preference relation $\succcurlyeq$ on (C, X) has an expected utility representation if there is a utility function $u : C \times X \to \mathbf{R}$ such that for every belief $p \in \Delta(X)$ and every two choices $a, b \in C$,*

$$a \succcurlyeq_p b \text{ if and only if } \sum_{x \in X} p(x) \cdot u(a, x) \geq \sum_{x \in X} p(x) \cdot u(b, x).$$

In this case, we say that the conditional preference relation $\succcurlyeq$ is *represented* by the utility function u. For a given vector $v \in \mathbf{R}^X$ we use the notation $u(a, v) := \sum_{x \in X} v(x) \cdot u(a, x)$. Hence, the condition above can be written as $a \succcurlyeq_p b$ if and only if $u(a, p) \geq u(b, p)$ for all $a, b \in C$. The following theorem shows that a conditional preference relation on *two choices* has an expected utility representation precisely when it satisfies the regularity axioms.

Theorem 3.1 (Expected utility representation for two choices) *Consider a set C consisting of two choices, a finite set of states X, and a conditional preference relation $\succcurlyeq$ on (C, X). Then, $\succcurlyeq$ has an expected utility representation, if and only if, it satisfies completeness, transitivity, continuity, preservation of indifference and preservation of strict preference.*

In particular, the conditional preference relation $\succcurlyeq$ in Figure 1 has an expected utility representation. One way to generate a utility function u that represents $\succcurlyeq$ 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 $u(b, z) < u(a, z)$. In the following subsection, when proving Theorem 3.1, we provide a general method for constructing a utility function that represents a regular conditional preference relation.

3.4 Proof

For proving Theorem 3.1 we need the following property which states that, under the regularity axioms, the indifference sets $P_{a \sim b}$ have a linear structure.

Lemma 3.1 (Linear structure of indifference sets) *Suppose that the conditional preference relation $\succcurlyeq$ is regular. Then, the following holds:*

(a) for every pair of choices a, b there is a hyperplane $H_{a \sim b}$ such that $P_{a \sim b} = H_{a \sim b} \cap \Delta(X)$;

(b) if $\succcurlyeq$ has preference reversals on $\{a, b\}$, then $\langle P_{a \sim b} \rangle$ is a hyperplane and $P_{a \sim b} = \langle P_{a \sim b} \rangle \cap \Delta(X)$.

If there are three states and $\succcurlyeq$ has preference reversals on $\{a, b\}$, then it follows that the indifference set $P_{a\sim b}$ must be a line segment. See Figure 1, where $P_{a\sim b}$ is the intersection of the hyperplane $H_{a\sim b}$ with the belief simplex $\Delta(X)$, resulting in a line segment.

Proof of Theorem 3.1. (a) Suppose that $C = \{a, b\}$ and that $\succcurlyeq$ is represented by a utility function u. We show that $\succcurlyeq$ satisfies the five axioms listed above. Clearly, $\succcurlyeq$ satisfies completeness and transitivity. To show continuity, take $p \in P_{a\succ b}$ and $q \in P_{b\succ a}$. Then, $u(a, p) > u(b, p)$ and $u(b, q) > u(a, q)$, and hence there is some $\lambda \in (0, 1)$ such that $u(a, (1-\lambda)p + \lambda q) = u(b, (1-\lambda)p + \lambda q)$. Thus, $(1-\lambda)p + \lambda q \in P_{a\sim b}$, and we see that $\succcurlyeq$ satisfies continuity. To show preservation of indifference, take $p, q \in P_{a\sim b}$ and some $\lambda \in (0, 1)$. Then, $u(a, p) = u(b, p)$ and $u(a, q) = u(b, q)$, and hence $u(a, (1-\lambda)p + \lambda q) = u(b, (1-\lambda)p + \lambda q)$. Therefore, $(1-\lambda)p + \lambda q \in P_{a\sim b}$, and we conclude that $\succcurlyeq$ satisfies preservation of indifference. To show preservation of strict preference, take $p \in P_{a\succcurlyeq b}$, $q \in P_{a\succ b}$ and some $\lambda \in (0, 1)$. Then, $u(a, p) \geq u(b, p)$ and $u(a, q) > u(b, q)$, and hence $u(a, (1-\lambda)p + \lambda q) > u(b, (1-\lambda)p + \lambda q)$. Therefore, $(1-\lambda)p + \lambda q \in P_{a\succ b}$, and we conclude that $\succcurlyeq$ satisfies preservation of strict preference.

(b) Suppose now that $\succcurlyeq$ satisfies the five axioms. We will construct a utility function u that represents $\succcurlyeq$. If a is equivalent to b then we can choose any utility function u with $u(a, x) = u(b, x)$ for all $x \in X$.

Suppose now that a is not equivalent to b. Then, there is a belief $q \in \Delta(X)$ with $a \nsim_q b$, say $b \succ_q a$. By Lemma 3.1 there is a hyperplane $H_{a\sim b}$ such that $P_{a\sim b} = H_{a\sim b} \cap \Delta(X)$. Let H' be the unique hyperplane that is parallel to $H_{a\sim b}$ and such that $q \in H'$. Then, $H' \neq H_{a\sim b}$. By Lemma 2.1 (b) there are vectors $v_1, ..., v_{|X|} \in H_{a\sim b}$ and a vector $v' \in H'$ such that $v_1, ..., v_{|X|}, v'$ are affinely independent.

Now, fix the utilities $u(a, x)$ arbitrarily for all $x \in X$ and choose some $\alpha > 0$. By Lemma 2.1 (c) there is a unique affine mapping $u_b : \mathbf{R}^X \to \mathbf{R}$ such that

$$u_b(v_k) = u(a, v_k) \text{ for all } k \in \{1, ..., |X|\} \text{ and } u_b(v') = u(a, v') + \alpha. \tag{3.1}$$

Set $u(b, x) := u_b([x])$ for all $x \in X$, which completes the utility function u. We show that u represents $\succcurlyeq$.

As $H_{a\sim b}$ is a hyperplane and $v_1, ..., v_{|X|} \in H_{a\sim b}$ are affinely independent, the vectors $v_1, ..., v_{|X|}$ constitute an affine basis for $H_{a\sim b}$. By (3.1) we then conclude

$$u_b(v) = u(a, v) \text{ for all } v \in H_{a\sim b}. \tag{3.2}$$

We now show that

$$u_b(v) = u(a, v) + \alpha \text{ for all } v \in H'. \tag{3.3}$$

To see this, take some $v \in H'$. Since H' is parallel to $H_{a\sim b}$ and $v' \in H'$, there are some vectors $w, w' \in H_{a\sim b}$ such that $v = v' + w - w'$. As v is an affine combination of v', w and w', and the mapping u_b is affine, it follows that

$$\begin{aligned} u_b(v) &= u_b(v') + u_b(w) - u_b(w') = u(a, v') + \alpha + u(a, w) - u(a, w') \\ &= u(a, v' + w - w') + \alpha = u(a, v) + \alpha, \end{aligned}$$

where the second equality follows from (3.1) and the third equality from the fact that $u(a, v)$ is linear in v.

Take some $p \in \Delta(X)$. By Lemma 2.1 (a) there is some λ with $p \in (1-\lambda)H_{a\sim b} + \lambda H'$. Using (3.2) and (3.3) it then follows that $u(b,p) = u_b(p) = u(a,p) + \lambda\alpha$. Hence,

$$u(a,p) > u(b,p) \text{ if and only if } \lambda < 0 \text{ and } u(a,p) < u(b,p) \text{ if and only if } \lambda > 0. \tag{3.4}$$

Recall that $H_{a\sim b} \cap \Delta(X) = P_{a\sim b}$. As $\succcurlyeq$ is regular, the hyperplane $H_{a\sim b}$ separates $P_{b\succ a}$ from $P_{a\succ b}$. Assume that $p \in (1-\lambda)H_{a\sim b} + \lambda H'$. Since $q \in P_{b\succ a}$ and $q \in H'$, we conclude that

$$p \in P_{a\succ b} \text{ if and only } \lambda < 0 \text{ and } p \in P_{b\succ a} \text{ if and only } \lambda > 0. \tag{3.5}$$

By (3.4) and (3.5) we see that $p \in P_{a\succ b}$ if and only $u(a,p) > u(b,p)$, and $p \in P_{b\succ a}$ if and only if $u(a,p) < u(b,p)$. Hence, u represents $\succcurlyeq$. ■

4 Existence of Uniform Preference Increase

In this section we start by showing that the regularity axioms are no longer sufficient to guarantee an expected utility representation if we move to more than two choices. To this purpose, we present an example of a conditional preference relation that satisfies the regularity axioms, yet lacks an expected utility representation. The reason for this failure is that, starting from this conditional preference relation, we cannot uniformly increase the preference for any given choice. In the second part of this section we will formalize what we mean by a uniform preference increase, and define a new axiom which states that, starting from the conditional preference relation at hand, we must be able to uniformly increase the preference for at least one choice.

4.1 Why Regularity Axioms Are Not Sufficient

Consider the conditional preference relation $\succcurlyeq$ represented by Figure 2. It may be verified that $\succcurlyeq$ satisfies all the regularity axioms. Yet, there is no expected utility representation for $\succcurlyeq$. To see why, suppose there would be a utility function u that represents $\succcurlyeq$. Then, the induced expected utilities of a and b must be equal on the hyperplane $H_{a\sim b}$, the expected utilities of b and c must be equal on the hyperplane $H_{b\sim c}$ and the expected utilities of a and c must be equal on the hyperplane $H_{a\sim c}$, 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 $H_{a\sim b}$ and $H_{b\sim c}$ intersect, which is impossible since v does not belong to $H_{a\sim c}$.

This raises the question: What is "wrong" with this conditional preference relation? As it turns out, we cannot uniformly increase the preference for choice a by a fixed degree without violating transitivity. To see this, suppose there would be an alternative conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for choice a by a fixed degree, relative to $\succcurlyeq$. Then, the degree of preference between a and b and the degree of preference 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 degree of preference between a and b is equal to his degree of preference 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,

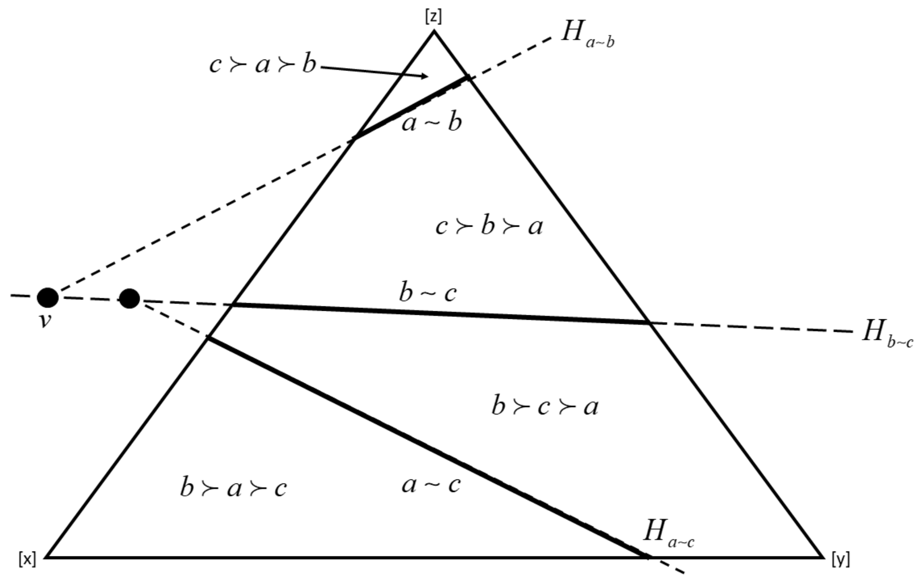

Figure 2: Regularity axioms are not sufficient for expected utility representation

the preference between a and b and the preference 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 s that moves from one point in $H_{b\sim c}$ to another point in $H_{b\sim c}$. 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 choice 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 s. 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 choice b or for choice c in this example.

As we will show in this paper, the absence of a uniform preference increase is precisely what prevents a regular conditional preference relation from having an expected utility representation. 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 choices.

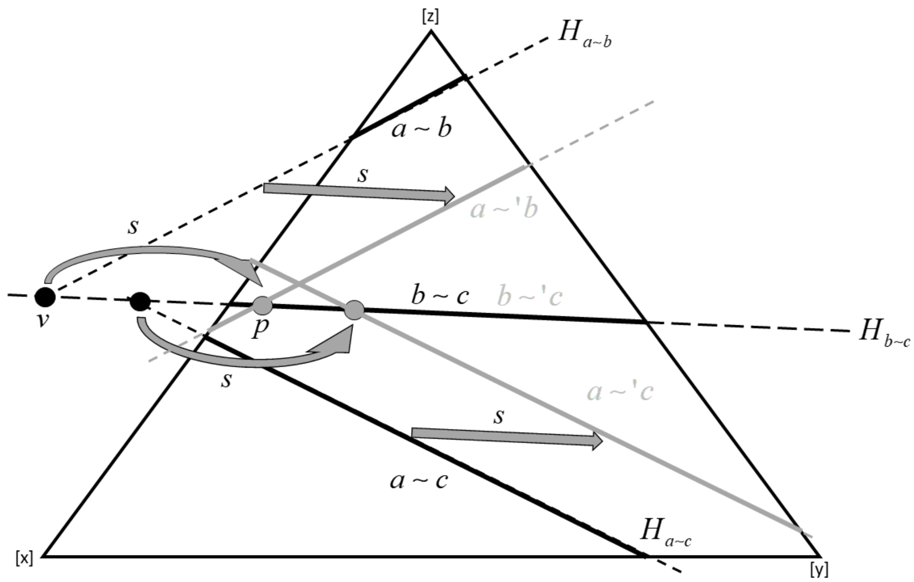

Figure 3: When there is no uniform preference increase

4.2 The Axiom "Existence of Uniform Preference Increase"

Imagine the DM holds a conditional preference relation $\succcurlyeq$, and decides to uniformly increase his preference for choice a by the degree δ. That is, for every belief p, and relative to every other choice b, the preference for a is increased by the same degree δ. How would the new conditional preference relation $\succcurlyeq'$ compare to $\succcurlyeq$?

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 degree of preference between a and b will change linearly.

Principle 2: The DM prefers b to c precisely when the degree by which he prefers a to b is less than the degree by which he prefers a to c.

Here, the degree by which the DM prefers a to b may also be negative, which means that he prefers b to a.

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 $\succcurlyeq'$ increases the degree of preference for a by the amount δ, relative to $\succcurlyeq$. Hence, $\deg_{a\succ' b}(p'_{ab}) = 0$ and $\deg_{a\succ' b}(p_{ab}) = \delta$, where $\deg_{a\succ' b}(q)$ informally denotes the degree by which the DM prefers a to b at the belief q under $\succcurlyeq'$. By principle 1 we then conclude that $\deg_{a\succ' b}(p) = \lambda\delta$.

Now consider an alternative belief $q_{ab} \in P_{a\sim b}$ and some belief q' such that $p = (1-\lambda)q' + \lambda q_{ab}$. See

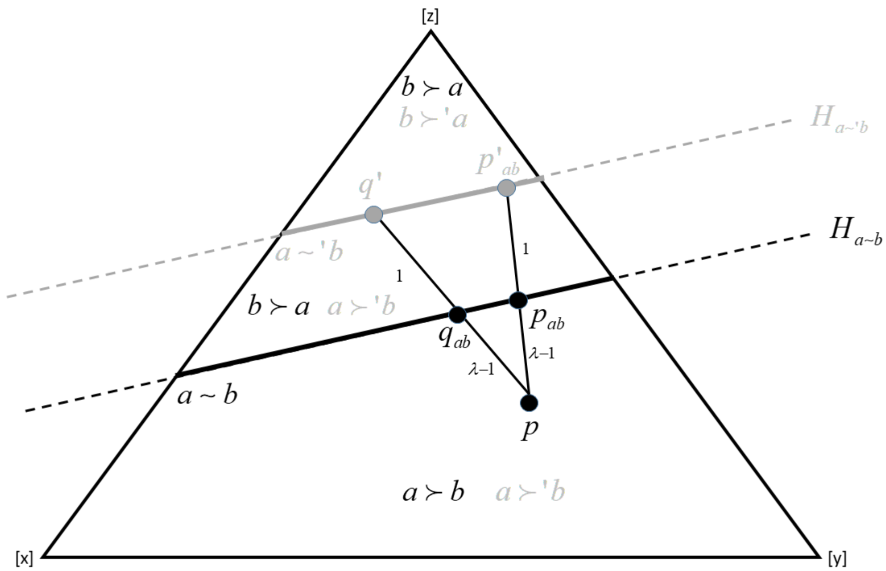

Figure 4: A uniform preference increase generates parallel indifference sets $P_{a\sim b}$ and $P_{a\sim' b}$

Figure 4. As $\deg_{a\succ' b}(q_{ab}) = \delta$, it follows by principle 1 that

$$\lambda\delta = \deg_{a\succ' b}(p) = (1-\lambda)\deg_{a\succ' b}(q') + \lambda\delta,$$

which implies that $\deg_{a\succ' b}(q') = 0$, and hence $q' \in P_{a\sim' b}$. We thus see that, whenever

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

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. Geometrically speaking, this means that the hyperplanes $H_{a\sim b}$ and $H_{a\sim' b}$, which generate the indifference sets $P_{a\sim b}$ and $P_{a\sim' b}$, must be parallel. See Figure 4.

Moreover, it can be seen from the figure that

$$\begin{aligned} P_{a\succsim b} &= \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\} \text{ and} \\ P_{a\succsim' b} &= \{p \in \Delta(X) \mid \text{there is } \lambda \geq 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}. \end{aligned}$$

That is, the conditional preference relations between a and b in $\succsim$ and $\succsim'$ are completely determined by the hyperplanes $H_{a\sim b}$ and $H_{a\sim' b}$.

Consider now a third choice c. Choose the beliefs $p \in \Delta(X), p_{ab} \in P_{a\sim b}, p'_{ab} \in P_{a\sim' b}, p_{ac} \in P_{a\sim c}, p'_{ac} \in P_{a\sim' c}$ such that

$$p = (1-\lambda)p'_{ab} + \lambda p_{ab} \text{ and } p = (1-\mu)p'_{ac} + \mu p_{ac}.$$

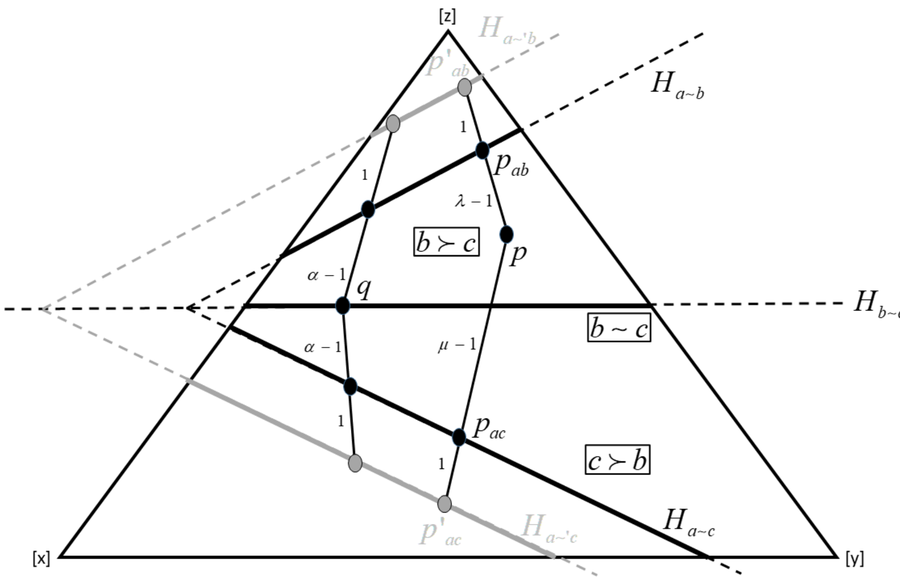

Figure 5: A uniform preference increase for a with three choices

See Figure 5 for an illustration where $\lambda, \mu > 1$ but $\lambda < \mu$. By the same argument as above, we then conclude that $\deg_{a \succ' b}(p) = \lambda\delta$ and $\deg_{a \succ' c}(p) = \mu\delta$. By principle 2 it then follows that the DM weakly prefers b to c precisely when $\lambda \leq \mu$. Note that $\succcurlyeq$ and $\succcurlyeq'$ hold the same preferences between b and c, as only the preference for a is altered when moving from $\succcurlyeq$ to $\succcurlyeq'$. Thus, we see that

$$P_{b \succcurlyeq c} = P_{b \succcurlyeq' c} = \{p \in \Delta(X) \mid \text{there are } \lambda \leq \mu \text{ with}$$
$$p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b} \text{ and } p \in (1-\mu)H_{a \sim' c} + \mu H_{a \sim c}\}.$$

In particular, the indifference set $P_{b \sim c}$ contains precisely those beliefs q for which there is some α with $q \in (1-\alpha)H_{a \sim' b} + \alpha H_{a \sim b}$ and $q \in (1-\alpha)H_{a \sim' c} + \alpha H_{a \sim c}$. See the belief q in Figure 5.

All these considerations give rise to the following definition.

Definition 4.1 (Uniform preference increase) *Consider a conditional preference relation $\succcurlyeq$, a choice a, and an alternative conditional preference relation $\succcurlyeq'$. Then, $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$ if $\succcurlyeq'$ is regular, and if for every choice $b \neq a$ there are different, parallel hyperplanes $H_{a \sim b}, H_{a \sim' b}$ such that*

(a) for every choice $b \neq a$

$$\begin{aligned} P_{a \succcurlyeq b} &= \{p \in \Delta(X) \mid \textit{there is } \lambda \geq 1 \textit{ with } p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}\} \textit{ and} \\ P_{a \succcurlyeq' b} &= \{p \in \Delta(X) \mid \textit{there is } \lambda \geq 0 \textit{ with } p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}\}; \end{aligned}$$

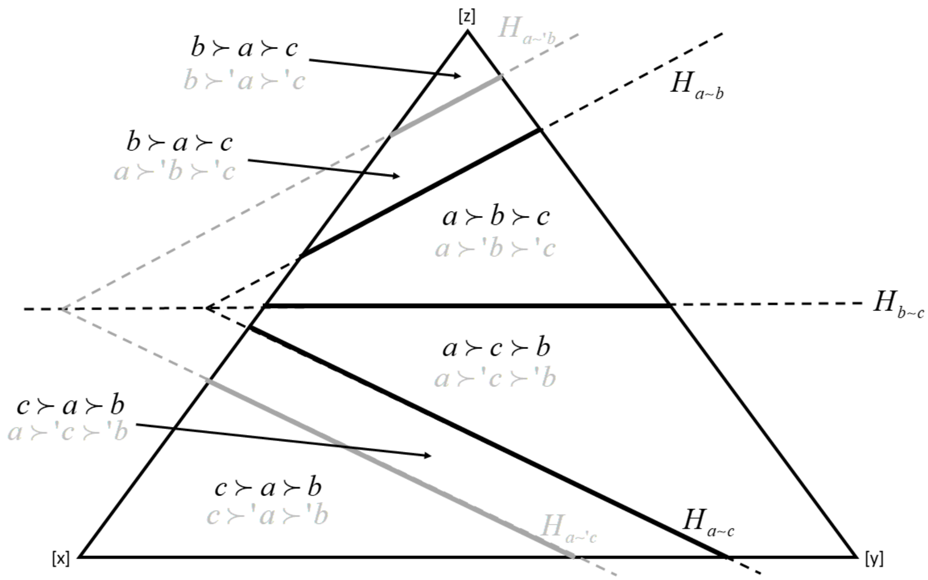

Figure 6: A typical uniform preference increase for a

(b) for every two choices $b, c \neq a$

$$P_{b \succcurlyeq c} = P_{b \succcurlyeq' c} = \{p \in \Delta(X) \mid \textit{there are } \lambda \leq \mu \textit{ with}$$
$$p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b} \textit{ and } p \in (1-\mu) H_{a \sim' c} + \mu H_{a \sim c}\}.$$

Figure 6 depicts a typical uniform preference increase for a when there are three choices.

The following axiom, which plays a central role in this paper, states that one should always be able to find a new conditional preference relation that uniformly increases the preference for one of the choices.

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

As we will see, this axiom opens the door towards a characterization of those conditional preference relations that admit an expected utility representation.

5 Characterization for More than Two Choices

In the previous section we have seen that the regularity axioms are no longer sufficient to guarantee an expected utility representation if there are more than two choices. We will now show that the new axiom "existence of a uniform preference increase", in combination with the regularity axioms, is sufficient to close this gap.

Theorem 5.1 (Expected utility representation for more than two choices) *Consider a finite set of choices C, a finite set of states X, and a conditional preference relation $\succcurlyeq$ on (C, X). Then, $\succcurlyeq$ 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.*

Before we prove this result, we first present a procedure that can be used to generate a utility function for a given conditional preference relation that satisfies our axioms. This procedure will be illustrated by means of a numerical example. Subsequently, we show in the proof of Theorem 5.1 that the utility function so obtained will always represent the conditional preference relation at hand, provided the latter satisfies all the axioms above.

5.1 Procedure for Constructing a Utility Function

We now present a procedure that generates a utility function u for a given conditional preference relation $\succcurlyeq$ satisfying our axioms. Importantly, this procedure explicitly uses a conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for some choice a relative to $\succcurlyeq$. Later we will show that the utility function u so constructed represents $\succcurlyeq$.

Definition 5.1 (Utility design procedure) *Consider a regular conditional preference relation $\succcurlyeq$ that satisfies existence of a uniform preference increase. Select a choice a and a conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for a relative to $\succcurlyeq$. For every choice $b \neq a$, let $H_{a\sim b}, H_{a\sim' b}$ be some different, parallel hyperplanes satisfying properties (a) and (b) in Definition 4.1.*

Start by choosing the utilities $u(a, x)$ arbitrarily for every $x \in X$, and by selecting a number $\alpha > 0$.

For every choice $b \neq a$, find vectors $v_1, ..., v_{|X|} \in H_{a\sim b}$ and a vector $v' \in H_{a\sim' b}$ such that $v_1, ..., v_{|X|}, v'$ are affinely independent. Find the unique affine mapping $u_b : \mathbf{R}^X \to \mathbf{R}$ such that $u_b(v_k) = u(a, v_k)$ for all $k \in \{1, ..., |X|\}$ and $u_b(v') = u(a, v') + \alpha$. Set $u(b, x) := u_b([x])$ for every $x \in X$.

Note that by Lemma 2.1 (b) we can always find vectors $v_1, ..., v_{|X|}$ and a vector $v' \in H_{a\sim' b}$ that are affinely independent. Moreover, Lemma 2.1 (c) guarantees that there is a unique affine mapping $u_b : \mathbf{R}^X \to \mathbf{R}$ such that $u_b(v_k) = u(a, v_k)$ for all $k \in \{1, ..., |X|\}$ and $u_b(v') = u(a, v') + \alpha$. Finding this affine mapping amounts to solving a system of linear equations.

The utilities $u(a, x)$, which can be chosen freely, may be viewed as the "baseline utilities", whereas the number $\alpha > 0$ can be interpreted as a *numeraire* that determines the utility increase for a associated with the uniform preference increase. We will now illustrate the procedure by a numerical example.

Example. Consider the conditional preference relation $\succcurlyeq$ represented by Figure 7, with three choices a, b, c and three states x, y and z. Note that a is weakly (but not strictly) dominated by c. Here, $(\frac{2}{5}, 0, \frac{3}{5})$ refers to the vector v with $v(x) = \frac{2}{5}, v(y) = 0$ and $v(z) = \frac{3}{5}$, and similarly for the other points. To show that there is a uniform preference increase for a, consider Figure 8. Let $H_{a\sim b} := \langle P_{a\sim b} \rangle$ be the unique hyperplane that passes through the points $(0, 0, 0), (\frac{2}{5}, 0, \frac{3}{5})$ and $(0, \frac{1}{4}, \frac{3}{4})$, let $H_{a\sim' b}$ be the unique hyperplane that is parallel

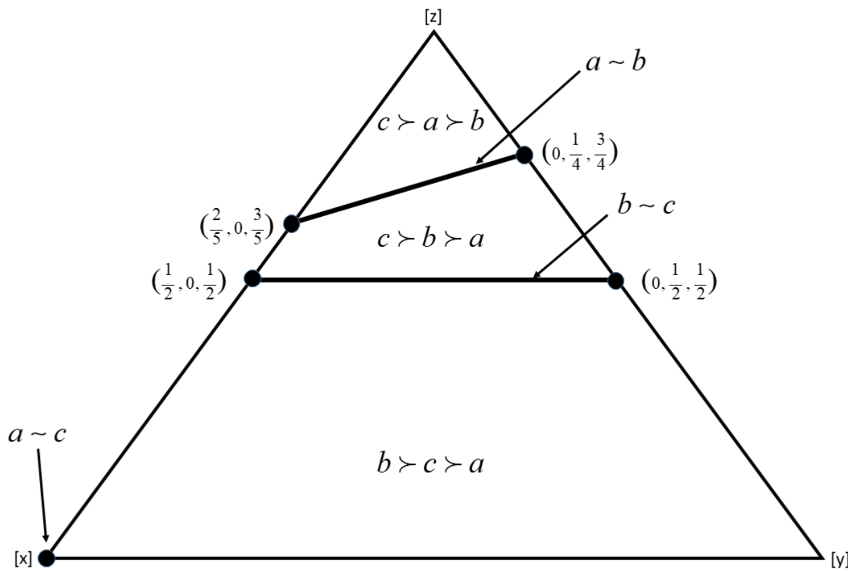

Figure 7: A numerical example for the utility design procedure

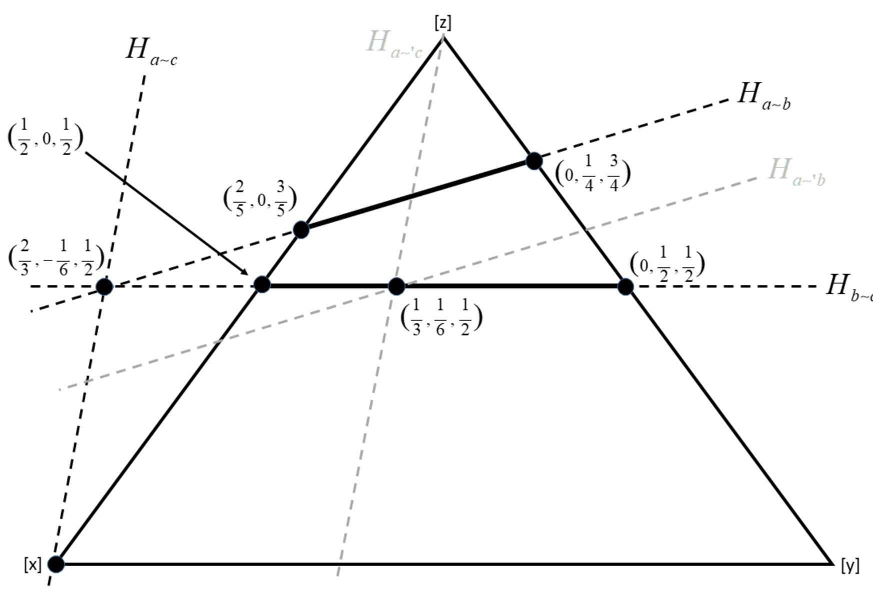

Figure 8: Constructing a utility function for a conditional preference relation

to $H_{a\sim b}$ and passes through the point $(\frac{1}{3}, \frac{1}{6}, \frac{1}{2})$, let $H_{a\sim c}$ be the unique hyperplane that passes through the points $(0,0,0), (\frac{2}{3}, -\frac{1}{6}, \frac{1}{2})$ and $(1,0,0) = [x]$, and let $H_{a\sim' c}$ be the unique hyperplane that is parallel to $H_{a\sim c}$ and passes through the point $(\frac{1}{3}, \frac{1}{6}, \frac{1}{2})$. Note that $H_{a\sim c} \cap \Delta(X) = \{[x]\} = P_{a\sim c}$.

Take the conditional preference relation $\succcurlyeq'$ that coincides with $\succcurlyeq$ on $\{b, c\}$, and that is otherwise induced by the hyperplanes $H_{a\sim b}, H_{a\sim' b}, H_{a\sim c}, H_{a\sim' c}$ through property (a) in Definition 4.1. It may be verified that $\succcurlyeq'$ is regular, and that properties (a) and (b) in Definition 4.1 are satisfied. As such, $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$.

We will now use the utility design procedure to construct a utility function u that represents $\succcurlyeq$, relying on $\succcurlyeq'$. We start by choosing $u(a,x) = u(b,x) = u(c,x) = 0$ and $\alpha = 1$.

To construct the utilities for b, recall that $H_{a\sim b}$ is the unique hyperplane that passes through the points $v_1 = (0,0,0), v_2 = (\frac{2}{5}, 0, \frac{3}{5})$ and $v_3 = (0, \frac{1}{4}, \frac{3}{4})$, and that $H_{a\sim' b}$ is the unique hyperplane that is parallel to $H_{a\sim b}$ and passes through $v' = (\frac{1}{3}, \frac{1}{6}, \frac{1}{2})$. It may be verified that v_1, v_2, v_3 and v' are affinely independent. Consider a general affine mapping $u_b : \mathbf{R}^X \to \mathbf{R}$ given by $u_b(v) = u_{b0} + v(x)u_{bx} + v(y)u_{by} + v(z)u_{bz}$. Then, the conditions $u_b(v_1) = u(a, v_1), u_b(v_2) = u(a, v_2), u_b(v_3) = u(a, v_3)$ and $u_b(v') = u(a, v') + \alpha$ give rise to the system of linear equations

$$u_{b0} = 0,\ u_{b0} + \tfrac{2}{5}u_{bx} + \tfrac{3}{5}u_{bz} = 0,\ u_{b0} + \tfrac{1}{4}u_{by} + \tfrac{3}{4}u_{bz} = 0,\ u_{b0} + \tfrac{1}{3}u_{bx} + \tfrac{1}{6}u_{by} + \tfrac{1}{2}u_{bz} = 1$$

which has the unique solution $u_{b0} = 0$, $u_{bx} = 3$, $u_{by} = 6$ and $u_{bz} = -2$. As such,

$$u(b,x) = u_b([x]) = 3,\ u(b,y) = u_b([y]) = 6 \text{ and } u(b,z) = u_b([z]) = -2.$$

To construct the utilities for c, recall that $H_{a\sim c}$ is the unique hyperplane that passes through the points $v_1 = (0,0,0), v_2 = (\frac{2}{3}, -\frac{1}{6}, \frac{1}{2})$ and $v_3 = (1,0,0)$, and that $H_{a\sim' c}$ is the unique hyperplane that is parallel to $H_{a\sim c}$ and passes through the point $v' = (\frac{1}{3}, \frac{1}{6}, \frac{1}{2})$. It may be verified that v_1, v_2, v_3 and v' are affinely independent. Consider a general affine mapping $u_c : \mathbf{R}^X \to \mathbf{R}$ given by $u_c(v) = u_{c0} + v(x)u_{cx} + v(y)u_{cy} + v(z)u_{cz}$. Then, the conditions $u_c(v_1) = u(a, v_1), u_c(v_2) = u(a, v_2), u_c(v_3) = u(a, v_3)$ and $u_c(v') = u(a, v') + \alpha$ give rise to the system of linear equations

$$u_{c0} = 0,\ u_{c0} + \tfrac{2}{3}u_{cx} - \tfrac{1}{6}u_{cy} + \tfrac{1}{2}u_{bz} = 0,\ u_{c0} + u_{cx} = 0,\ u_{c0} + \tfrac{1}{3}u_{cx} + \tfrac{1}{6}u_{cy} + \tfrac{1}{2}u_{cz} = 1$$

which has the unique solution $u_{c0} = 0$, $u_{cx} = 0$, $u_{cy} = 3$ and $u_{cz} = 1$. As such,

$$u(c,x) = u_c([x]) = 0,\ u(c,y) = u_c([y]) = 3 \text{ and } u(c,z) = u_c([z]) = 1.$$

The utility function u may be summarized by the table

u	x	y	z
a	0	0	0
b	3	6	-2
c	0	3	1

and it may be verified that the utility function u so constructed represents $\succcurlyeq$. The conditional preference relation $\succcurlyeq'$, which uniformly increases the preference for a relative to $\succcurlyeq$, is represented by the utility function u' given by

u'	x	y	z
a	1	1	1
b	3	6	−2
c	0	3	1

where the utility of choice a is uniformly increased by 1.

The utility function u is only one out of many different utility representations for $\succcurlyeq$. In fact, there are four degrees of freedom here. The first three degrees of freedom arise because the baseline utilities $u(a, x), u(b, x)$ and $u(c, x)$ can be chosen completely arbitrarily, whereas the freedom to choose any $\alpha > 0$ leads to a fourth degree of freedom. Also the conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for one of the choices is not unique, but it turns out that for any such $\succcurlyeq'$ the hyperplanes that induce the indifference sets for $\succcurlyeq$ and $\succcurlyeq'$ must all be parallel to the ones we have in Figure 8, and hence no additional degrees of freedom arise from the choice of $\succcurlyeq'$.

5.2 Proof

We are now ready to prove Theorem 5.1.

(a) Suppose that $\succcurlyeq$ is represented by a utility function u. In the proof of Theorem 3.1 we have already shown that $\succcurlyeq$ satisfies the five regularity axioms. To show existence of a uniform preference increase, consider an arbitrary choice $a \in C$ and a number $\alpha > 0$. Let u' be the utility function given by $u'(a, x) := u(a, x) + \alpha$ for every $x \in X$, and $u'(b, x) := u(b, x)$ for every choice $b \neq a$ and every $x \in X$. Let $\succcurlyeq'$ be the conditional preference relation induced by u'. We show that $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$.

By the proof of Theorem 3.1 we know that $\succcurlyeq'$ is regular. We will now construct, for every $b \neq a$, different, parallel hyperplanes $H_{a\sim b}, H_{a\sim' b}$ that satisfy the conditions (a) and (b) in Definition 4.1. We distinguish two cases: (i) b is not equivalent to a under $\succcurlyeq$, and (ii) b is equivalent to a under $\succcurlyeq$.

(i) Suppose first that b is not equivalent to a under $\succcurlyeq$. That is, there is some $x \in X$ with $u(a, x) \neq u(b, x)$. Then define the sets

$$H_{a\sim b} := \{v \in \mathbf{R}^X \mid u(a, v) = u(b, v)\} \text{ and } H_{a\sim' b} := \{v \in \mathbf{R}^X \mid u(a, v) + \alpha = u(b, v)\}. \tag{5.1}$$

To see that $H_{a\sim b}, H_{a\sim' b}$ are different, parallel hyperplanes consider the vector $n \in \mathbf{R}^X$ given by $n(x) := u(a, x) - u(b, x)$ for all $x \in X$. Then, by our assumption in case (i), $n \neq \underline{0}$. Moreover, by construction, $H_{a\sim b} = \{v \in \mathbf{R}^X \mid v \cdot n = 0\}$ and $H_{a\sim' b} = \{v \in \mathbf{R}^X \mid v \cdot n = -\alpha\}$, which implies that $H_{a\sim b}$ and $H_{a\sim' b}$ are different and parallel.

(ii) Suppose next that b is equivalent to a under $\succcurlyeq$. In that case we define

$$H_{a\sim b} := \{v \in \mathbf{R}^X \mid \textstyle\sum_{x \in X} v(x) = 1\}, \tag{5.2}$$

which is clearly a hyperplane, and we choose $H_{a\sim' b}$ equal to some arbitrary hyperplane that is different from, but parallel to, $H_{a\sim b}$.

It remains to show that properties (a) and (b) in Definition 4.1 are satisfied for all $b, c \neq a$. To show property (a), take some $b \neq a$. We distinguish two cases: (i) b is not equivalent to a under $\succcurlyeq$, and (ii) b is equivalent to a under $\succcurlyeq$.

(i) Suppose that b is not equivalent to a under $\succcurlyeq$. Then, by (5.1), Lemma 2.1 (a), the fact that u represents $\succcurlyeq$ and the fact that u' represents $\succcurlyeq'$,

$$\begin{aligned} P_{a\succcurlyeq b} &= \{p \in \Delta(X) \mid u(a,p) \geq u(b,p)\} \\ &= \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}, \end{aligned}$$

and

$$\begin{aligned} P_{a\succcurlyeq' b} &= \{p \in \Delta(X) \mid u'(a,p) \geq u'(b,p)\} = \{p \in \Delta(X) \mid u(a,p) + \alpha \geq u(b,p)\} \\ &= \{p \in \Delta(X) \mid \text{there is } \lambda \geq 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}. \end{aligned}$$

(ii) Suppose next that b is equivalent to a under $\succcurlyeq$. Then, by (5.2) we conclude that $H_{a\sim b} \cap \Delta(X) = \Delta(X)$. Moreover, $P_{a\succcurlyeq' b} = \Delta(X)$ as $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$. Hence,

$$\begin{gathered} \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\} = \Delta(X) = P_{a\succcurlyeq b} \text{ and} \\ \{p \in \Delta(X) \mid \text{there is } \lambda \geq 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\} = \Delta(X) = P_{a\succcurlyeq' b}. \end{gathered}$$

To show property (b) in Definition 4.1 take some $b, c \neq a$. We must prove that

$$\begin{gathered} P_{b\succcurlyeq c} = P_{b\succcurlyeq' c} = \{p \in \Delta(X) \mid \text{there are } \lambda \leq \mu \text{ with} \\ p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}. \end{gathered} \tag{5.3}$$

We distinguish two cases: (i) b and c are not equivalent to a under $\succcurlyeq$, and (ii) b or c is equivalent to a under $\succcurlyeq$.

(i) Suppose first that b and c are not equivalent to a under $\succcurlyeq$. Let A be the set on the righthand side of (5.3). To show that $P_{b\succcurlyeq c} \subseteq A$, take some $p \in P_{b\succcurlyeq c}$. As $H_{a\sim b}, H_{a\sim' b}$ are different but parallel hyperplanes and $H_{a\sim c}, H_{a\sim' c}$ are different but parallel hyperplanes, it follows by Lemma 2.1 (a) that there are numbers λ, μ such that $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$. By (5.1) it then follows that

$$u(a,p) - u(b,p) = (\lambda - 1)\alpha \text{ and } u(a,p) - u(c,p) = (\mu - 1)\alpha. \tag{5.4}$$

As $p \in P_{b\succcurlyeq c}$ and u represents $\succcurlyeq$ we know that $u(b,p) \geq u(c,p)$. We thus conclude from (5.4) that $\lambda \leq \mu$, and hence $p \in A$.

To show that $A \subseteq P_{b\succcurlyeq c}$, take some $p \in A$. Hence, there are numbers λ, μ with $\lambda \leq \mu$ such that $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$. Then, by (5.1),

$$u(a,p) - u(b,p) = (\lambda - 1)\alpha \text{ and } u(a,p) - u(c,p) = (\mu - 1)\alpha$$

which implies that $u(b,p) \geq u(c,p)$, and hence $p \in P_{b \succcurlyeq c}$. We thus conclude that $P_{b \succcurlyeq c} = A$.

(ii) Suppose next that b or c is equivalent to a under $\succcurlyeq$. Say b is equivalent to a under $\succcurlyeq$. Then, by (5.2), $H_{a\sim b} \cap \Delta(X) = \Delta(X)$, and $(1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ has an empty intersection with $\Delta(X)$ unless $\lambda = 1$. Hence,

$$A = \{p \in \Delta(X) \mid \text{there is } \mu \geq 1 \text{ with } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\} = P_{a \succcurlyeq c} = P_{b \succcurlyeq c},$$

where the second equality follows from property (a) and the last equality from the facts that b is equivalent to a under $\succcurlyeq$ and $\succcurlyeq$ is transitive.

We thus see that $P_{b \succcurlyeq c} = A$ for both cases (i) and (ii). As, by construction, $P_{b \succcurlyeq' c} = P_{b \succcurlyeq c}$, property (b) in Definition 4.1 is satisfied.

Summarizing, we conclude that $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$.

(b) Let $\succcurlyeq$ be a conditional preference relation on (C, X) that satisfies completeness, transitivity, continuity, preservation of indifference, preservation of strict preference and existence of a uniform preference increase. Let u be a utility function that is obtained by the utility design procedure in Definition 5.1, based on the hyperplanes $H_{a\sim b}, H_{a\sim' b}$ for every choice $b \neq a$. We show that u represents $\succcurlyeq$.

Take some choice $b \neq a$. Then, by construction of the utility design procedure, there are vectors $v_1, ..., v_{|X|} \in H_{a\sim b}$, a vector $v' \in H_{a\sim' b}$ such that $v_1, ..., v_{|X|}, v'$ are affinely independent, and an affine mapping $u_b : \mathbf{R}^X \to \mathbf{R}$ such that

$$u_b(v_k) = u(a, v_k) \text{ for all } k \in \{1, ..., |X|\},\ u_b(v') = u(a, v') + \alpha \text{ and} \tag{5.5}$$

$$u(b, x) = u_b([x]) \text{ for all } x \in X. \tag{5.6}$$

As $v_1, ..., v_{|X|} \in H_{a\sim b}$ are affinely independent, it follows that $\{v_1, ..., v_{|X|}\}$ is an affine basis for $H_{a\sim b}$. As $u_b(v_k) = u(a, v_k)$ for all $k \in \{1, ..., |X|\}$, it follows that

$$u_b(v) = u(a, v) \text{ for all } v \in H_{a\sim b}. \tag{5.7}$$

In a similar way as in the proof of Theorem 3.1 it can be shown that

$$u_b(v) = u(a, v) + \alpha \text{ for all } v \in H_{a\sim' b}. \tag{5.8}$$

We now show that u represents $\succcurlyeq$ on $\{a, b\}$. Take some $p \in \Delta(X)$. By Lemma 2.1 (a) there is some number λ with $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$. By (5.6), (5.7) and (5.8),

$$u(b,p) - u(a,p) = u_b(p) - u(a,p) = (1-\lambda)\alpha. \tag{5.9}$$

Moreover, by regularity of $\succcurlyeq$ and property (a) in Definition 4.1, we see that

$$p \in P_{a \succ b} \text{ if and only if } \lambda > 1, \text{ and } p \in P_{a\sim b} \text{ if and only if } \lambda = 1. \tag{5.10}$$

By combining (5.9) and (5.10) we conclude that $p \in P_{a \succ b}$ if and only if $u(a,p) > u(b,p)$ and $p \in P_{a\sim b}$ if and only if $u(a,p) = u(b,p)$. Thus, u represents $\succcurlyeq$ on $\{a, b\}$.

Consider now some choices $b, c \neq a$. We show that u represents $\succcurlyeq$ on $\{a, b\}$. Take some $p \in \Delta(X)$. By Lemma 2.1 (a) there are some numbers λ, μ such that $p \in (1-\lambda)H_{a\sim' b}+\lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c}+\mu H_{a\sim c}$. By (5.6), (5.7) and (5.8),

$$u(b,p) - u(a,p) = u_b(p) - u(a,p) = (1-\lambda)\alpha \text{ and } u(c,p) - u(a,p) = u_c(p) - u(a,p) = (1-\mu)\alpha. \quad (5.11)$$

Moreover, by regularity of $\succcurlyeq$ and property (b) in Definition 4.1, we see that

$$p \in P_{b\succ c} \text{ if and only if } \lambda < \mu, \text{ and } p \in P_{b\sim c} \text{ if and only if } \lambda = \mu. \quad (5.12)$$

By combining (5.11) and (5.12) we conclude that $p \in P_{b\succ c}$ if and only if $u(b,p) > u(c,p)$ and $p \in P_{b\sim c}$ if and only if $u(b,p) = u(c,p)$. Hence, u represents $\succcurlyeq$ on $\{b, c\}$.

Since we have shown that u represents $\succcurlyeq$ on $\{a, b\}$ for every $b \neq a$ and u represents $\succcurlyeq$ on $\{b, c\}$ for every $b, c \neq a$, it follows that u represents $\succcurlyeq$, which was to show. This completes the proof. ■

6 The Case of No Weakly Dominated Choices

In the previous section we have seen that a regular conditional preference relation has an expected utility representation precisely when there is a uniform preference increase for at least one of the choices. 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 for an important special case: the scenario when there are no weakly dominated choices. In that case, the existence of a uniform preference increase is equivalent to two easily verifiable conditions: *strong transitivity* and the *line property.* These conditions can be tested directly by only considering the conditional preference relation $\succcurlyeq$ at hand, without having to search for a uniform preference increase explicitly.

The case of no weakly dominated choices is an important and typical case for a rational DM. The reason is that a rational DM may be expected to never make a weakly dominated choice. Indeed, if choice a is weakly dominated by b then, whatever belief the DM holds, he will always weakly prefer b to a, and sometimes strictly prefer b to a. Hence, there seems to be no good reason to make choice a. But if we eliminate all weakly dominated choices from the decision problem, then we are left with a reduced decision problem in which no remaining choice is weakly dominated.

In this section we start by presenting the result that, for the scenario of no weakly dominated choices, the existence of a uniform preference increase is equivalent to strong transitivity and the line property, and discuss some important consequences of this result. Subsequently, we provide an intuitive sketch of the proof. The full proof can be found in the appendix.

6.1 Characterization for the Case of No Weakly Dominated Choices

Suppose there are no choices that are weakly dominated by, or equivalent to, some other choice. The question whether a uniform preference increase exists or not can then be answered by testing two easily verifiable conditions.

Theorem 6.1 (The case of no weakly dominated choices) *Let $\succcurlyeq$ be a regular conditional preference relation such that no two choices weakly dominate each other, or are equivalent to each other, under $\succcurlyeq$. Then, $\succcurlyeq$ satisfies existence of a uniform preference increase, if and only if,*

(a) (strong transitivity) for every three choices $a, b, c \in C$ we have that $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle \subseteq \langle P_{b\sim c}\rangle$, and

(b) (line property) there is a line $L = \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ that intersects each of the hyperplanes $\langle P_{a\sim b}\rangle$ at a single point $v + \lambda_{ab} w$, and where $\lambda_{ab} \neq \lambda_{ac}$ whenever $\langle P_{a\sim b}\rangle \neq \langle P_{a\sim c}\rangle$, 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 $a, b, c, d \in C$.

Strong transitivity thus states that the linear spans of the indifference sets $P_{a\sim b}$, $P_{a\sim c}$ and $P_{b\sim c}$ must have a common intersection. See, for instance, Figure 6 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 $\Delta(X)$. If $P_{a\sim b}$, $P_{a\sim c}$ and $P_{b\sim c}$ would all meet inside the belief simplex, then property (a) would correspond to the usual transitivity of the indifference relation between choices. In that sense, property (a) can 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 violates strong transitivity. Recall that we argued informally that this conditional preference relation could not have a uniform preference increase. This property now follows formally from Theorem 6.1.

The line property only has bite if there are at least four choices. An important consequence of this property is that on the line L, whenever we know the five points where the linear spans $\langle P_{e\sim f}\rangle$ intersect the line L for all $e, f \in \{a, b, c, d\}$, $\{e, f\} \neq \{a, b\}$, then we also know where $\langle P_{a\sim b}\rangle$ intersects the line L.

If we combine Theorem 6.1 with Theorem 3.1, we obtain the following characterization of conditional preference relations that admit an expected utility representation.

Corollary 6.1 (Characterization for the case of no weakly dominated choices) *Let $\succcurlyeq$ be a regular conditional preference relation such that no two choices weakly dominate each other, or are equivalent to each other, under $\succcurlyeq$. Then, $\succcurlyeq$ has an expected utility representation, if and only if, strong transitivity and the line property hold.*

The practical advantage of this result is that strong transitivity and the line property are easily veriable conditions that only require us to investigate the conditional preference relation $\succcurlyeq$ at hand, without having to look for a uniform preference increase. This result, however, no longer holds if we allow for weakly dominated choices. Consider, for instance, the conditional preference relation $\succcurlyeq$ from Figure 7. It clearly violates strong transitivity as $\langle P_{a\sim b}\rangle \cap \langle P_{b\sim c}\rangle$ is not a subset of $\langle P_{a\sim c}\rangle = \langle \{[x]\}\rangle$. At the same time, $\succcurlyeq$ admits an expected utility representation as we have seen.

Corollary 6.1 also has an interesting consequence for the case of two states, that is, when $X = \{x, y\}$. In that scenario, strong transitivity is equivalent to the usual transitivity of the indifference relation between choices, 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 6.2 (Characterization for the case of two states) *Let $X=\{x,y\}$ and $\succcurlyeq$ a regular conditional preference relation such that no two choices weakly dominate each other, or are equivalent to each other, under $\succcurlyeq$. Let $P_{a\sim b}=\{p_{ab}\}$ for all choices $a,b\in C$. Then, $\succcurlyeq$ has an expected utility representation, if and only if,*

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

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

Hence, for the case of two states and no weak dominance, checking whether an expected utility representation exists is particularly easy, as one only needs to verify the regularity axioms and the formula above for every tuple of choices a,b,c,d. A direct consequence is that for two choices and three states, *every* regular conditional preference relation for which there are no weakly dominated choices will have an expected utility representation. In fact, this property even holds if we would allow for weakly dominated choices.

6.2 Sketch of the Proof

We will now give a sketch of the proof for Theorem 6.1. To start, we explain why the regularity axioms in combination with the existence of a uniform preference change imply strong transitivity. Consider three choices a,b and c. Assume there is a conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for a relative to $\succcurlyeq$, with associated hyperplanes $H_{a\sim'b}$ and $H_{a\sim'c}$. See Figure 9 for an illustration. Choose vectors $v'\in H_{a\sim'b}\cap H_{a\sim'c}$, $v\in\langle P_{a\sim b}\rangle\cap\langle P_{a\sim c}\rangle$ and two different beliefs p and q on the line through v and v', such that $p=(1-\lambda)v'+\lambda v$ and $q=(1-\mu)v'+\mu v$. See Figure 9 for an illustration with $\lambda,\mu>1$. By setting $H_{a\sim b}:=\langle P_{a\sim b}\rangle$ and $H_{a\sim c}:=\langle P_{a\sim c}\rangle$ we obtain, by construction, that

$$p\in(1-\lambda)H_{a\sim'b}+\lambda H_{a\sim b}\text{ and }p\in(1-\lambda)H_{a\sim'c}+\lambda H_{a\sim c}$$

and

$$q\in(1-\mu)H_{a\sim'b}+\mu H_{a\sim b}\text{ and }q\in(1-\mu)H_{a\sim'c}+\mu H_{a\sim c}.$$

Hence, by property (b) in Definition 4.1, p and q are in $P_{b\sim c}$. As a consequence, v is in $\langle P_{b\sim c}\rangle$ and hence $\langle P_{a\sim b}\rangle\cap\langle P_{a\sim c}\rangle\subseteq\langle P_{b\sim c}\rangle$, establishing strong transitivity.

We next intuitively explain why the line property follows from the existence of a uniform preference increase in combination with the regularity axioms. Suppose that the conditional preference relation $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$ with associated hyperplanes $H_{a\sim'b}$ for every $b\neq a$. Consider four choices a,b,c,d and a line $L=\{v+\lambda w\mid\lambda\in\mathbf{R}\}$ which intersects each of the sets $\langle P_{e\sim f}\rangle$ at

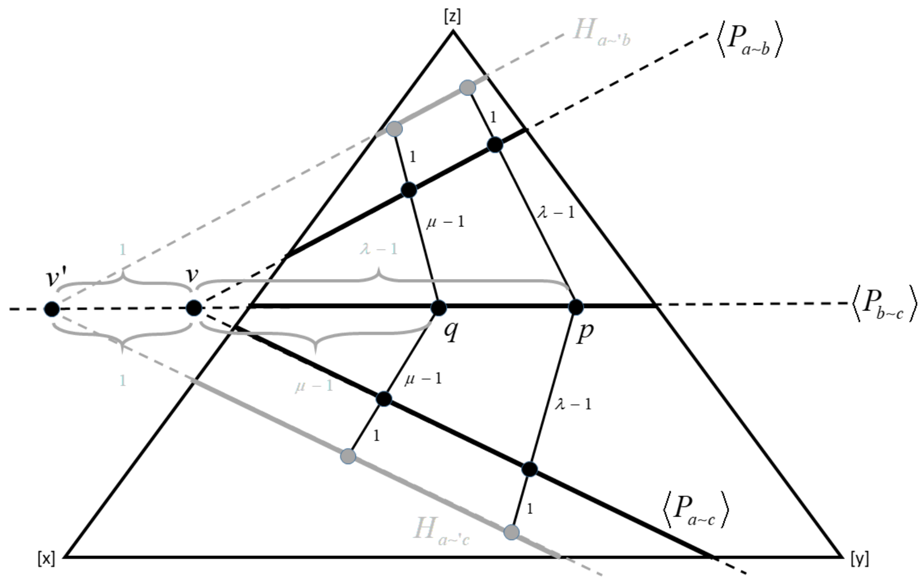

Figure 9: Why the existence of a uniform preference change implies strong transitivity

a unique point $v_{ef} = v + \lambda_{ef} w$ for all $e, f \in \{a, b, c, d\}$, and intersects each of the hyperplanes $H_{a\sim' e}$ at a unique point $v'_{ae} = v + \lambda'_{ae} w$ for all $e \in \{b, c, d\}$. Suppose, for the purpose of our argument, that all these points are different.

Concentrate first on the choices a, b and c. As $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$ there must, intuitively, be some number $\alpha > 0$ such that $\deg_{a\succ' b}(p) = \deg_{a\succ b}(p) + \alpha$ and $\deg_{a\succ' c}(p) = \deg_{a\succ c}(p) + \alpha$ for all beliefs p. Here, $\deg_{a\succ' b}(p)$ denotes the degree by which the DM prefers a to b under $\succcurlyeq'$ at the belief p, and similarly for $\deg_{a\succ b}(p)$. Moreover, $\deg_{a\succ' b}(p) = \deg_{a\succ' c}(p)$ for all $p \in P_{b\sim c}$. As the hyperplanes $\langle P_{a\sim b}\rangle, H_{a\sim' b}, \langle P_{a\sim c}\rangle, H_{a\sim' c}$ and $\langle P_{b\sim c}\rangle$ generate the sets $P_{a\sim b}, P_{a\sim' b}, P_{a\sim c}, P_{a\sim' c}$ and $P_{b\sim c}$, respectively, we must have that

$$\deg_{a\succ' b}(v'_{ab}) = 0,\ \ \deg_{a\succ' b}(v_{ab}) = \alpha,\ \ \deg_{a\succ' c}(v'_{ac}) = 0,\ \ \deg_{a\succ' c}(v_{ac}) = \alpha \text{ and}$$
$$\deg_{a\succ' b}(v_{bc}) = \deg_{a\succ' c}(v_{bc}).$$

See Figure 10 for an illustration. It thus follows that

$$\frac{\Delta \deg_{a\succ' b}}{\Delta \deg_{a\succ' c}} = \frac{\lambda_{ac} - \lambda_{bc}}{\lambda_{ab} - \lambda_{bc}}.$$

In a similar way we conclude that

$$\frac{\Delta \deg_{a\succ' c}}{\Delta \deg_{a\succ' d}} = \frac{\lambda_{ad} - \lambda_{cd}}{\lambda_{ac} - \lambda_{cd}} \text{ and } \frac{\Delta \deg_{a\succ' b}}{\Delta \deg_{a\succ' d}} = \frac{\lambda_{ad} - \lambda_{bd}}{\lambda_{ab} - \lambda_{bd}}.$$

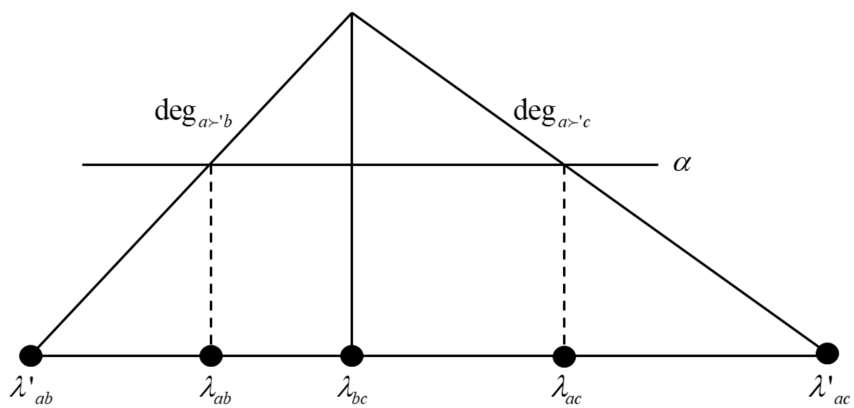

Figure 10: Why the existence of a uniform preference increase implies the line property

As, by the chain rule, $\frac{\Delta \deg_{a\succ' b}}{\Delta \deg_{a\succ' d}} = \frac{\Delta \deg_{a\succ' b}}{\Delta \deg_{a\succ' c}} \cdot \frac{\Delta \deg_{a\succ' c}}{\Delta \deg_{a\succ' d}}$ it follows that

$$\frac{\lambda_{ad} - \lambda_{bd}}{\lambda_{ab} - \lambda_{bd}} = \frac{\lambda_{ac} - \lambda_{bc}}{\lambda_{ab} - \lambda_{bc}} \cdot \frac{\lambda_{ad} - \lambda_{cd}}{\lambda_{ac} - \lambda_{cd}}$$

which establishes the line property.

We finally explain how strong transitivity and the line property enable the construction of a uniform preference increase. Consider the conditional preference relation $\succcurlyeq$ in Figure 11. Hence, there are three states x, y and z, and four choices a, b, c and d. For convenience, we have only indicated the indifference sets. It may be verified that $\succcurlyeq$ satisfies strong transitivity and the line property. To see the latter, consider the line L in Figure 11 which intersects each set $\langle P_{e\sim f}\rangle$ at a single point v_{ef}. It may be verified that this collection of six points satisfies the formula of the line property.

We construct a conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for a relative to $\succcurlyeq$, as follows. By the line property, we can find points v'_{ab}, v'_{ac} and v'_{ad} on the line L such that

$$\frac{\lambda'_{ab} - \lambda_{ab}}{\lambda_{bc} - \lambda_{ab}} = \frac{\lambda'_{ac} - \lambda_{ac}}{\lambda_{bc} - \lambda_{ac}}, \; \frac{\lambda'_{ac} - \lambda_{ac}}{\lambda_{cd} - \lambda_{ac}} = \frac{\lambda'_{ad} - \lambda_{ad}}{\lambda_{cd} - \lambda_{ad}} \text{ and } \frac{\lambda'_{ab} - \lambda_{ab}}{\lambda_{bd} - \lambda_{ab}} = \frac{\lambda'_{ad} - \lambda_{ad}}{\lambda_{bd} - \lambda_{ad}}.$$

See the points v'_{ab}, v'_{ac} and v'_{ad} in Figure 12. Let $H_{a\sim' b}$ be the unique hyperplane that is parallel to $\langle P_{a\sim b}\rangle$ and passes through v'_{ab}. Similarly, we construct the hyperplanes $H_{a\sim' c}$ and $H_{a\sim' d}$. See Figure 12 for an illustration. Set $H_{a\sim b} := \langle P_{a\sim b}\rangle$ and similarly for $H_{a\sim c}$ and $H_{a\sim d}$. Finally, let $\succcurlyeq'$ be the conditional

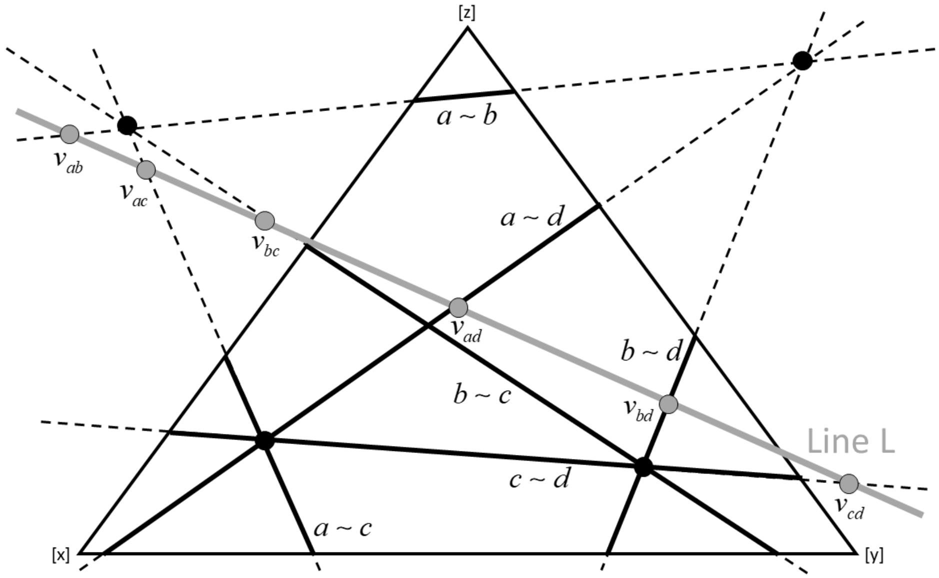

Figure 11: A conditional preference relation with four choices that satisfies strong transitivity and the line property

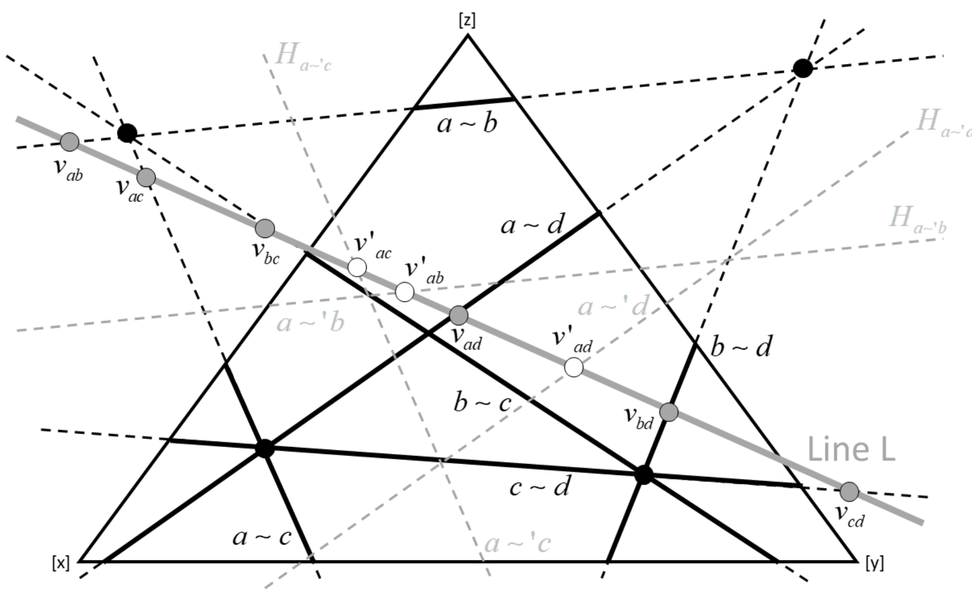

Figure 12: How to construct a uniform preference increase under strong transitivity and the line property

preference relation that coincides with $\succcurlyeq$ on $\{b, c, d\}$, and that is otherwise induced by the hyperplanes $H_{a\sim' b}, H_{a\sim b}, H_{a\sim' c}, H_{a\sim c}, H_{a\sim' d}$ and $H_{a\sim d}$ through condition (a) in Definition 4.1. In can be seen from Figure 12 that $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$. In the appendix we show that this construction works in general. That is, if there are no weakly dominated or equivalent choices, and strong transitivity and the line property are satisfied, we can always construct a uniform preference increase in this way.

7 Discussion

(a) Related literature. The paper that is closest to ours is Gilboa and Schmeidler (2003). It also provides an axiomatic characterization of conditional preference relations that admit an expected utility representation, but does so for a restricted class of utility functions – *diversified* utility functions. By this we mean utility matrices 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 choices 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 choices were examples of non-diversified conditional preference relations, having a non-diversified utility representation. By definition, diversity does not allow for weak dominance between choices. It may also be verified that the diversity condition exludes cases with two states and more than two choices, and cases with three states and more than three choices. Indeed, if we have two states and at least three choices, then there are 6 possible strict orderings on three choices, 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 choices, then there are 24 possible strict orderings on four choices, but at most 16 of these will be possible in a regular conditional preference relation. However, Gilboa and Schmeidler (2003) allow for infinitely many, even uncountably many, choices 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 in 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.

(b) Logical (in)dependence of the axioms. From conditions (a) and (b) in Definition 4.1 it follows that the axiom *existence of a uniform preference increase* logically implies each of the regularity axioms.

The other direction is not true: In Figure 2 we have seen a conditional preference relation that satisfies each of the regularity axioms, but for which there is no uniform preference increase. It may be shown, however, that the five regularity axioms are logically independent amongst each other.

(c) How unique is the representation? In the numerical example from Section 5.1 we have seen that there will always be a whole range of different utility functions that represent a given conditional preference relation, provided it satisfies our axioms of course. In the example from that section there were four degrees of freedom. In general, the number of degrees of freedom may range anywhere between $|X|$ and $|C| \cdot |X|$. It is clear that for a given "baseline choice" a we can always choose the baseline utilities $u(a, x)$ in an arbitrary fashion, yielding $|X|$ degrees of freedom to start with. If, on the one extreme, the DM is always indifferent between all of his choices, the other utilities will be fixed, and hence the total number of degrees of freedom will be $|X|$. If, on the other extreme, the DM always holds the same strict preference relation between his choices, say $c_1 \succ c_2 \succ ... \succ c_n$, then the utilities for the other choices can be chosen freely at each of the states, as long as these utilities respect this strict preference relation. This would result in $|C| \cdot |X|$ degrees of freedom. For the special case where there are no weakly dominated or equivalent choices, and where the indifference sets $P_{a \sim b}$ are all pairwise different, it can be shown that there are $|X| + 1$ degrees of freedom. The single extra degree of freedom comes from the freedom to choose any $\alpha > 0$ in the utility design procedure from Section 5.1.

Savage (1954), on the other hand, has shown that for every preference relation over acts that satisfies the Savage axioms, the utility representation is unique up to a positive affine transformation, leaving much less freedom than is typically the case in our framework. The reason is that a DM in Savage's framework holds preferences over *all possible* acts, providing us with "more data" that restrict the possible utilities compared to a DM in our framework. Gilboa and Schmeidler (2003) have a result similar to Savage's, showing that their utility representation is unique up to the choice of the baseline utilities, and up to multiplying all utilities by the same positive number. This results in $|X| + 1$ degrees of freedom – less than the number of degrees of freedom in many of our cases. This is mainly due to their diversity axiom, which requires that all strict rankings between four choices or less should be obtained for at least one belief, generating a "large amount of data" that impose restrictions on the utilities.

(d) Uniform preference increase and counterfactual reasoning. The key axiom that has led to our expected utility theorem is the "existence of a uniform preference increase". This axioms asks whether the conditional preference relation $\succcurlyeq$ at hand allows for an alternative conditional preference relation $\succcurlyeq'$ in which the DM uniformly increases the preference for one of the choices. In a sense, it requires the DM to perform a thought experiment in which he counterfactually would increase his preference for a given choice in a uniform fashion, and see whether this would lead to inconsistencies. In Section 6 we have shown, however, that in the case of no weakly dominated choices this counterfactual reasoning can be replaced by checking some conditions that refer exclusively to the conditional preference relation $\succcurlyeq$ at hand. In a sense, counterfactual reasoning is also required in the settings of Savage and Anscombe and Aumann. There, the DM is required to rank *all* possible acts, also those that are not directly related to the decision problem at hand. For those acts, the DM must thus imagine he would be facing a decision problem that involves such as act.

(e) Utility differences reflecting degrees of preference. The utilities that represent a given conditional preference relation have a very clear and intuitive interpretation in our framework: The expected utility difference between a and b at a belief p reflects "by how much" the DM prefers a to b at that particular belief. Consequently, the utility difference between a and b at a state x reflects "by how much" he prefers a to b at state x. Consider, for instance, the table of utilities in Section 5.1 that represents the conditional preference relation $\succcurlyeq$ of Figure 7. At state x, the utility difference between a and b is $0 - 3 = -3$, whereas at state z this utility difference is $0 - (-2) = 2$. Interpreting these utility differences as "degrees of preference", we may say that the degree by which the DM prefers b to a at x is higher than the degree by which he prefers a to b at z. This is also reflected by the fact that at the belief $(\frac{1}{2}, 0, \frac{1}{2})$, which is halfway between $[x]$ and $[z]$, the DM still prefers b to a. We can say even more: The DM will only become indifferent between a and b at the belief $(\frac{2}{5}, 0, \frac{3}{5})$, where the probability assigned to z is 1.5 times the probability assigned to x. As such, we may say that the degree of preference between a and b at x is 1.5 times the degree of preference between a and b at z. This corresponds precisely to the relative utility differences between a and b at x and z, which are -3 and 2, respectively.

(f) Belief as a primitive notion. An important difference with Savage's (1954) framework is that we view the DM's belief as a primitive notion, from which we can derive his preference relation over choices. This is precisely how a conditional preference relation is defined: It takes the belief as an input, and delivers the preferences over choices as an output. One of the beautiful features of Savage's 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 is a whole 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 choices based on this belief, which finally allows him to make a choice based on this preference relation.

(g) Belief revision. A conditional preference relation does not only specify the DM's preferences over choices 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 choices C_i would be the set of actions in the game, and the states would be the set $X_i = \times_{j \neq i} C_j$ of opponents' choice combinations. Fix a conditional preference relation $\succcurlyeq^i$ 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 $\sigma_i \in \Delta(C_i)$ for every player i, such that $\sigma_i(c_i) > 0$ only if c_i is optimal for the induced preference relation $\succcurlyeq^i_{\sigma_{-i}}$. Here, σ_{-i} denotes the product of the probability distributions σ_j for $j \neq i$, which is a probability distribution over C_{-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' choices, 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 $C_i^0 := C_i$ for all players i, and

$$C_i^k := \{c_i \in C_i^{k-1} \mid c_i \text{ optimal for } \succcurlyeq_{p_i}^i \text{ for some } p_i \in \Delta(C_{-i}^{k-1})\}$$

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) Possible extensions. The analysis in this paper can be extended in various directions. First, one could replace "standard" probability distributions on X by lexicographic probability systems (Blume, Brandenburger and Dekel (1991)) or non-standard beliefs (Robinson (1973), Hammond (1994), Halpern (2010)) and find axioms which guarantee a representation by lexicographic or non-standard expected utility, respectively. Alternatively, one could attempt to extend the analysis to the case of infinitely many choices and/or states. Some of the arguments in the proofs heavily rely on the set of choices, and above all the set of states, being finite. It would be interesting to see how much, and what exactly, would have to be modified to generalize the results to the infinite case.

8 Appendix

8.1 Proof of Section 2

Proof of Lemma 2.1. (a) As the hyperplanes V and W are different and parallel, there are a vector $n \neq \underline{0}$ and two different numbers α, β such that

$$V = \{v \in \mathbf{R}^X \mid v \cdot n = \alpha\} \text{ and } W = \{v \in \mathbf{R}^X \mid v \cdot n = \beta\}.$$

Take some $v \in \mathbf{R}^X$. We first show that there is some number λ with $v \in (1-\lambda)V + \lambda W$. If $v \in V$ then we can choose $\lambda = 0$. Assume next that $v \notin V$. Choose the number λ such that $(1-\lambda)\alpha + \lambda\beta = v \cdot n$. As $v \cdot n \neq \alpha$ we conclude that $\lambda \neq 0$. Take some $v' \in V$ and consider the vector $v'' := \frac{1}{\lambda}v - \frac{1-\lambda}{\lambda}v'$. By the choice of λ it then follows that $v'' \cdot n = \beta$ and hence $v'' \in W$. As $v = (1-\lambda)v' + \lambda v''$ it follows that $v \in (1-\lambda)V + \lambda W$. We next show that the number λ is unique. Suppose that $v \in (1-\lambda)V + \lambda W$ and $v \in (1-\mu)V + \mu W$. Then, $v \cdot n = (1-\lambda)\alpha + \lambda\beta = (1-\mu)\alpha + \mu\beta$. As $\alpha \neq \beta$ it follows that $\lambda = \mu$.

(b) As V is a hyperplane, there are $|X|$ affinely independent vectors $v_1, ..., v_{|X|}$ in V. Take some arbitrary vector $v_{|X|+1} \in W$. We show that $v_1, ..., v_{|X|}, v_{|X|+1}$ are affinely independent. As every affine combination of $v_1, ..., v_{|X|}$ is in V, we know that $v_{|X|+1}$ is not an affine combination of $v_1, ..., v_{|X|}$. Suppose now that some vector in $\{v_1, ..., v_{|X|}\}$, say v_1, is an affine combination of the other vectors in $\{v_1, ..., v_{|X|}, v_{|X|+1}\}$. Then,

$$v_1 = \sum_{k=2}^{|X|+1} \alpha_k v_k \text{ for some } a_2, ..., \alpha_{|X|+1} \text{ with } \sum_{k=2}^{|X|+1} \alpha_k = 1. \tag{8.1}$$

As $v_1, ..., v_{|X|}$ are affinely independent it must be that $\alpha_{|X|+1} \neq 0$. But then, it follows from (8.1) that

$$v_{|X|+1} = \frac{1}{\alpha_{|X|+1}}(v_1 - \sum\nolimits_{k=2}^{|X|} \alpha_k v_k) \text{ with } \frac{1}{\alpha_{|X|+1}}(1 - \sum\nolimits_{k=2}^{|X|} \alpha_k) = 1.$$

In other words, $v_{|X|+1}$ is an affine combination of $v_1, ..., v_{|X|}$, which is impossible as we have seen. Thus, the vectors $v_1, ..., v_{|X|}, v_{|X|+1}$ are affinely independent.

(c) Let $X = \{x_1, ..., x_n\}$. Consider the system of linear equations, with $n+1$ variables $\beta_0, ..., \beta_n$ and $n+1$ equations, given by

$$1 \cdot \beta_0 + \sum_{k=1}^{n} v_m(x_k)\beta_k = \alpha_m \text{ for all } m \in \{1, ..., n+1\}. \tag{8.2}$$

Hence, the m-th equation has the row of coefficients $r_m := (1, v_m(x_1), ..., v_m(x_n))$. We show that the vectors $r_1, ..., r_{n+1}$ are linearly independent. Suppose not. Then, there is some vector, say r_1, that is a linear combination of the other vectors $r_2, ..., r_{n+1}$. Hence, there are numbers $\gamma_2, ..., \gamma_{n+1}$ with $r_1 = \sum_{m=2}^{n+1} \gamma_m r_m$. By definition of the vectors r_m we must then have that

$$1 = \sum\nolimits_{m=2}^{n+1} \gamma_m \text{ and } (v_1(x_1), ..., v_1(x_n)) = \sum\nolimits_{m=2}^{n+1} \gamma_m (v_m(x_1), ..., v_m(x_n)).$$

This implies, however, that v_1 is an affine combination of $v_2, ..., v_{n+1}$, which is a contradiction to the assumption that $v_1, ..., v_{n+1}$ are affinely independent. We thus conclude that the vectors of coefficients $r_1, ..., r_{n+1}$ are linearly independent. This guarantees that the system of linear equations in (8.2) has a unique solution $\beta_0, \beta_1, ..., \beta_n$. Consider the affine mapping $f : \mathbf{R}^X \to \mathbf{R}$ given by

$$f(v) := \beta_0 + \sum_{k=1}^{n} v(x_k)\beta_k.$$

Then, by (8.2), $f(v_m) = \alpha_m$ for all $m \in \{1, ..., n+1\}$. Moreover, f is also the only affine mapping with these property. To see this, note that every affine mapping $g : \mathbf{R}^X \to \mathbf{R}$ takes the form $g(v) = \delta_0 + \sum_{k=1}^{n} v(x_k)\delta_k$ for some numbers $\delta_0, ..., \delta_{n+1}$. If $g(v_m) = \alpha_m$ for all $m \in \{1, ..., n+1\}$, then the numbers $\delta_0, ..., \delta_{n+1}$ must solve the system of linear equations in (8.2). As this system has the unique solution $\beta_0, ..., \beta_{n+1}$, it follows that $g = f$. ■

8.2 Proof of Section 3

For the proof of Lemma 3.1 we need the following three properties.

Lemma 8.1 (Implications of regularity axioms) *Suppose that the conditional preference relation $\succcurlyeq$ is regular. Then, for every pair of choices a, b the following properties hold:*

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

(b) If $\succcurlyeq$ has preference reversals on $\{a,b\}$, then there are $|X|-1$ linearly independent beliefs in $P_{a\sim b}$;

(c) If a weakly dominates b under $\succcurlyeq$ then

$$P_{a\sim b} = \{p \in \Delta(X) \mid \sum_{x \in X_{a\sim b}} p(x) = 1\}.$$

Proof of Lemma 8.1. (a) Clearly, $P_{a\sim b} \subseteq \langle P_{a\sim b}\rangle \cap \Delta(X)$. It remains to show that $\langle P_{a\sim b}\rangle \cap \Delta(X) \subseteq P_{a\sim b}$. We prove, by induction on k, that every $p \in \langle P_{a\sim b}\rangle \cap \Delta(X)$ 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 $p \in \langle P_{a\sim b}\rangle \cap \Delta(X)$ 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 $p \in \Delta(X)$ and $p_1, ..., p_k \in \Delta(X)$, we have that $\sum_{x\in X} p(x) = \sum_{x \in X} p_m(x) = 1$ 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 $w \in \Delta(X)$. By construction, $\sum_{x\in X} w(x) = 1$, and it thus remains to show that $w(x) \geq 0$ for all x. We distinguish two cases: If $\lambda_1 > 0$, then $\lambda_m > 0$ for all m. As $p_2, ..., p_k \in \Delta(X)$, it follows that $w(x) \geq 0$ for all x. Suppose now that $\lambda_1 < 0$. Then, $w = \frac{1}{1-\lambda_1}(p - \lambda_1 p_1)$. As $p, p_1 \in \Delta(X)$ and $\lambda_1 < 0$, it follows that $w(x) \geq 0$ for all x. We thus conclude that $w \in \Delta(X)$. Hence, $w \in \langle P_{a\sim b}\rangle \cap \Delta(X)$ 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 $\Delta(X)$ with $p_1, w \in P_{a\sim b}$.

We will now show that $p \in P_{a\sim b}$. If $\lambda_1 \in [0,1]$, then it follows by preservation of indifference that $p = \lambda_1 p_1 + (1-\lambda_1)w \in P_{a\sim b}$. Suppose now that either $\lambda_1 < 0$ or $\lambda_1 < 1$. Assume, without loss of generality, that $\lambda_1 < 0$. Then $w = \frac{1}{1-\lambda_1}((\lambda_1 p_1 + (1-\lambda_1)w) - \lambda_1 p_1)$, where $\frac{1}{1-\lambda_1}, -\frac{\lambda_1}{1-\lambda_1} \in (0,1)$. Suppose, contrary to what we want to show, that $\lambda_1 p_1 + (1-\lambda_1)w \notin P_{a\sim b}$. 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 $p = \lambda_1 p_1 + (1-\lambda_1)w \in P_{a\sim b}$.

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 $\langle P_{a\sim b}\rangle \cap \Delta(X) \subseteq P_{a\sim b}$.

(b) As $\succcurlyeq$ has preference reversals on $\{a,b\}$, the sets $X_{a\succ b}$ and $X_{b\succ a}$ must both be non-empty. Indeed, suppose that $X_{a\succ b}$ would be empty. Then, $[x] \in P_{b\succcurlyeq a}$ for all $x \in X$. But then it would follow by preservation of indifference and preservation of strict preference that $p \in P_{b\succcurlyeq a}$ for all beliefs $p \in \Delta(X)$, which would be a contradiction to our assumption that $\succcurlyeq$ has preference reversals on $\{a,b\}$. Hence, we see that $X_{a\succ b}$ cannot be empty. By a similar argument it can be shown that $X_{b\succ a}$ cannot be empty either.

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

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

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

$$p_{yx} := (1 - \lambda_{yx})[y] + \lambda_{yx}[x] \in P_{a \sim b}, \tag{8.4}$$

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

$$p_{zx} := (1 - \lambda_{zx})[z] + \lambda_{zx}[x] \in P_{a \sim b}. \tag{8.5}$$

Consider the set

$$B := \{[x] \mid x \in X_{a \sim b}\} \cup \{p_{yz}\} \cup \{p_{yx} \mid x \in X_{b \succ a} \backslash \{z\}\} \cup \{p_{zx} \mid x \in X_{a \succ b} \backslash \{y\}\},$$

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

Take some numbers α_x for $x \in X_{a \sim b}$, some number α_{yz}, some numbers α_{yx} for $x \in X_{b \succ a} \backslash \{z\}$ and some numbers α_{zx} for $x \in X_{a \succ b} \backslash \{y\}$ such that

$$\sum_{x \in X_{a \sim b}} \alpha_x[x] + \alpha_{yz} p_{yz} + \sum_{x \in X_{b \succ a} \backslash \{z\}} \alpha_{yx} p_{yx} + \sum_{x \in X_{a \succ b} \backslash \{z\}} \alpha_{zx} p_{zx} = \underline{0}.$$

By (8.3), (8.4) and (8.5), this sum is equal to

$$\begin{aligned}
\sum_{x \in X_{a \sim b}} \alpha_x[x] + \alpha_{yz}((1 - \lambda_{yz})[y] + \lambda_{yz}[z]) + \sum_{x \in X_{b \succ a} \backslash \{z\}} \alpha_{yx}((1 - \lambda_{yx})[y] + \lambda_{yx}[x]) \\
+ \sum_{x \in X_{a \succ b} \backslash \{y\}} \alpha_{zx}((1 - \lambda_{zx})[z] + \lambda_{zx}[x]) \\
= \sum_{x \in X_{a \sim b}} \alpha_x[x] + \left[\alpha_{yz}((1 - \lambda_{yz}) + \sum_{x \in X_{b \succ a} \backslash \{z\}} \alpha_{yx}(1 - \lambda_{yx}) \right] [y] \\
+ \left[\alpha_{yz}\lambda_{yz} + \sum_{x \in X_{a \succ b} \backslash \{y\}} \alpha_{zx}(1 - \lambda_{zx}) \right] [z] + \sum_{x \in X_{b \succ a} \backslash \{z\}} \alpha_{yx}\lambda_{yx}[x] + \sum_{x \in X_{a \succ b} \backslash \{y\}} \alpha_{zx}\lambda_{zx}[x] = \underline{0}.
\end{aligned}$$

As all vectors in $\{[x] \mid x \in X\}$ are linearly independent, it follows that $\alpha_x = 0$ for all $x \in X_{a \sim b}$, that $\alpha_{yx}\lambda_{yx} = 0$ for all $x \in X_{b \succ a} \backslash \{z\}$, and that $\alpha_{zx}\lambda_{zx} = 0$ for all $x \in X_{a \succ b} \backslash \{y\}$. Since $\lambda_{yx} \in (0,1)$ for all $x \in X_{b \succ a} \backslash \{z\}$ and $\lambda_{zx} \in (0,1)$ for all $x \in X_{a \succ b} \backslash \{y\}$, it follows that $\alpha_{yx} = 0$ for all $x \in X_{b \succ a} \backslash \{z\}$ and $\alpha_{zx} = 0$ for all $x \in X_{a \succ b} \backslash \{y\}$. 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 $|X| - 1$ linearly independent beliefs.

Consider the set

$$B := \{[x] \mid x \in X_{a\sim b}\} \cup \{p_{yz}\} \cup \{p_{yx} \mid x \in X_{b\succ a}\backslash\{z\}\} \cup \{p_{zx} \mid x \in X_{a\succ b}\backslash\{y\}\},$$

which contains $|X|-1$ vectors in $P_{a\sim b}$. It may be verified that all vectors in B are linearly independent, and hence $P_{a\sim b}$ contains $|X|-1$ linearly independent beliefs.

(c) Let $A = \{p \in \Delta(X) \mid \sum_{x\in X_{a\sim b}} p(x) = 1\}$. 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, $p(x) > 0$ for some $x \in X_{a\succ b}$, and hence $p = p(x)\cdot[x]+(1-p(x))\cdot q$ for some $q \in \Delta(X)$. As $[x] \in P_{a\succ b}$, $q \in P_{a\succcurlyeq b}$ and $p(x) > 0$, it follows by preservation of strict preference that $p \in P_{a\succ b}$, which is a contradiction to the assumption that $p \in P_{a\sim b}$. We thus conclude that $p \in A$. To show that $A \subseteq P_{a\sim b}$, take some $p \in A$. Then, it follows by preservation of indifference that $p \in P_{a\sim b}$. We thus see that $P_{a\sim b} = A$. ■

Proof of Lemma 3.1. (a) Assume first that a and b are equivalent under $\succcurlyeq$. Then we can choose the hyperplane $H_{a\sim b} = \{v \in \mathbf{R}^X \mid \sum_{x\in X} v(x) = 1\}$, which guarantees that $P_{a\sim b} = \Delta(X) = H_{a\sim b} \cap \Delta(X)$.

If $\succcurlyeq$ has preference reversals on $\{a, b\}$ then we know from Lemma 8.1 (b) that $\dim(\langle P_{a\sim b}\rangle) \geq |X|-1$. Note that $\dim(\langle P_{a\sim b}\rangle) \neq |X|$ since otherwise $\langle P_{a\sim b}\rangle = \mathbf{R}^X$, which would imply by Lemma 8.1 (a) that $P_{a\sim b} = \langle P_{a\sim b}\rangle \cap \Delta(X) = \Delta(X)$. That would be a contradiction to the assumption that $\succcurlyeq$ has preference reversals on $\{a, b\}$. Hence, $\dim(\langle P_{a\sim b}\rangle) = |X|-1$, which implies that $\langle P_{a\sim b}\rangle$ is a hyperplane. By choosing $H_{a\sim b} = \langle P_{a\sim b}\rangle$ we know by Lemma 8.1 (a) that $P_{a\sim b} = H_{a\sim b} \cap \Delta(X)$.

Suppose next that a weakly dominates b under $\succcurlyeq$. Then, we know by Lemma 8.1 (c) that

$$P_{a\sim b} = \{p \in \Delta(X) \mid \sum_{x\in X_{a\sim b}} p(x) = 1\}. \tag{8.6}$$

Let $n \in \mathbf{R}^X$ be the vector with $n(x) = 0$ for all $x \in X_{a\sim b}$ and $n(x) = 1$ for all $x \in X\backslash X_{a\sim b}$. As a is not equivalent to b under $\succcurlyeq$ we know that $X_{a\sim b} \neq X$, and hence $n \neq \underline{0}$. Define $H_{a\sim b} := \{v \in \mathbf{R}^X \mid v \cdot n = 0\}$, which is a hyperplane. Then, it follows from (8.6) that $P_{a\sim b} = H_{a\sim b} \cap \Delta(X)$.

(b) This part follows from our arguments above. ■

8.3 Proof of Section 6

For proving Theorem 6.1 we need the following result.

Lemma 8.2 (Preference and parallel hyperplanes) *Let $\succcurlyeq$ be a regular conditional preference relation. Consider two choices a, b and two different, parallel hyperplanes H and H' such that $P_{a\sim b} = H \cap \Delta(X)$.*

(a) If there is some $q \in P_{a\succ b}$ and $\lambda < 1$ with $q \in (1-\lambda)H' + \lambda H$ then

$$P_{a\succcurlyeq b} = \{p \in \Delta(X) \mid \text{there is some } \lambda \leq 1 \text{ with } p \in (1-\lambda)H' + \lambda H\}.$$

(b) If there is some $q \in P_{b\succ a}$ and $\lambda < 1$ with $q \in (1-\lambda)H' + \lambda H$ then

$$P_{a\succcurlyeq b} = \{p \in \Delta(X) \mid \text{there is some } \lambda \geq 1 \text{ with } p \in (1-\lambda)H' + \lambda H\}.$$

Proof. (a) Suppose there is some $q \in P_{a \succ b}$ and $\mu < 1$ with $q \in (1-\mu)H' + \mu H$. Let A be the set on the righthand side of the equality in (a). To show that $P_{a \succcurlyeq b} \subseteq A$ take some $p \in P_{a \succcurlyeq b}$, and let λ be such that $p \in (1-\lambda)H' + \lambda H$. Suppose, contrary to what we want to show, that $\lambda > 1$. Then, there is some $\alpha \in (0,1)$ such that $(1-\alpha)q + \alpha p \in H \cap \Delta(X) = P_{a \sim b}$. However, as $q \in P_{a \succ b}$ and $p \in P_{a \succcurlyeq b}$, it follows by preservation of strict preference that $(1-\alpha)q + \alpha p \in P_{a \succ b}$ for every $\alpha \in (0,1)$, which is a contradiction. Hence, $\lambda \leq 1$, which implies that $p \in A$.

To show that $A \subseteq P_{a \succcurlyeq b}$, take some $p \in A$. Hence, $p \in (1-\lambda)H' + \lambda H$ for some $\lambda \leq 1$. If $\lambda = 1$ then $p \in H \cap \Delta(X) = P_{a \sim b}$. Suppose now that $\lambda < 1$. As $\mu < 1$ also, we know that p must be on the same side of the hyperplane H as q. Since $P_{a \sim b} = H \cap \Delta(X)$ and $q \in P_{a \succ b}$, it follows by continuity and preservation of strict preference that $p \in P_{a \succ b}$ also. We thus have that $p \in P_{a \succcurlyeq b}$ in both cases. This completes the proof for (a).

Property (b) can be shown in a similar fashion. ■

Proof of Theorem 6.1. (a) Suppose that $\succcurlyeq$ is regular and satisfies existence of a uniform preference increase. Then, by Theorem 5.1, there is a utility function u that represents $\succcurlyeq$. To show strong transitivity, consider three choices a, b and c. As $\succcurlyeq$ has preference reversals on every pair of choices, it follows from Lemma 3.1 (b) that

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

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 $\langle P_{e \sim f} \rangle \neq \langle P_{e \sim g} \rangle$. To see this, select a vector w such that $w \notin \langle P_{e \sim f} \rangle$ for every $e, f \in C$. 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 $\langle P_{e \sim f} \rangle \neq \langle P_{e \sim g} \rangle$.

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

$$(\lambda_{ab} - \lambda_{bd})(\lambda_{ac} - \lambda_{bc})(\lambda_{ad} - \lambda_{cd}) = (\lambda_{ab} - \lambda_{bc})(\lambda_{ac} - \lambda_{cd})(\lambda_{ad} - \lambda_{bd}). \tag{8.7}$$

Assume first that $\lambda_{ef} = \lambda_{eg}$ for some $e, f, g \in \{a, b, c, d\}$. Then, by transitivity of $\succcurlyeq$, we have that $\lambda_{ef} = \lambda_{eg} = \lambda_{fg}$, and (8.7) 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

$$\delta_{ab}(\lambda) := u(a, v + \lambda w) - u(b, v + \lambda w), \ \delta_{ac}(\lambda) := u(a, v + \lambda w) - u(c, v + \lambda w) \text{ and} \tag{8.8}$$
$$\delta_{ad}(\lambda) := u(a, v + \lambda w) - u(d, v + \lambda w) \text{ for all } \lambda \in \mathbf{R}.$$

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

$$\delta_{ab}(\lambda) - \delta_{ab}(\mu) = D_{ab} \cdot (\lambda - \mu) \tag{8.9}$$
$$\delta_{ac}(\lambda) - \delta_{ac}(\mu) = D_{ac} \cdot (\lambda - \mu) \text{ and} \tag{8.10}$$
$$\delta_{ad}(\lambda) - \delta_{ad}(\mu) = D_{ad} \cdot (\lambda - \mu) \tag{8.11}$$

for all $\lambda, \mu \in \mathbf{R}$.

We will now show that

$$\frac{D_{ab}}{D_{ac}} = \frac{\lambda_{ac} - \lambda_{bc}}{\lambda_{ab} - \lambda_{bc}}. \tag{8.12}$$

Note that $\frac{D_{ab}}{D_{ac}}$ plays the same role as $\frac{\Delta \deg_{a \succ' c}}{\Delta \deg_{a \succ' d}}$ in the proof sketch of Section 6.2. By taking $\lambda = \lambda_{ab}$ and $\mu = \lambda_{bc}$, we obtain from (8.9) that

$$D_{ab} = \frac{\delta_{ab}(\lambda_{ab}) - \delta_{ab}(\lambda_{bc})}{\lambda_{ab} - \lambda_{bc}} = -\frac{\delta_{ab}(\lambda_{bc})}{\lambda_{ab} - \lambda_{bc}} \tag{8.13}$$

since $\delta_{ab}(\lambda_{ab}) = 0$. Similarly, by taking $\lambda = \lambda_{ac}$ and $\mu = \lambda_{bc}$, we obtain from (8.10) that

$$D_{ac} = \frac{\delta_{ac}(\lambda_{ac}) - \delta_{ac}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} = -\frac{\delta_{ac}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} = -\frac{\delta_{ab}(\lambda_{bc})}{\lambda_{ac} - \lambda_{bc}} \tag{8.14}$$

since $\delta_{ac}(\lambda_{ac}) = 0$ and $\delta_{ac}(\lambda_{bc}) = \delta_{ab}(\lambda_{bc})$. By combining (8.13) and (8.14) we obtain (8.12).

In a similar fashion it can be shown that

$$\frac{D_{ac}}{D_{ad}} = \frac{\lambda_{ad} - \lambda_{cd}}{\lambda_{ac} - \lambda_{cd}} \text{ and } \frac{D_{ab}}{D_{ad}} = \frac{\lambda_{ad} - \lambda_{bd}}{\lambda_{ab} - \lambda_{bd}}. \tag{8.15}$$

As

$$\frac{D_{ab}}{D_{ad}} = \frac{D_{ab}}{D_{ac}} \cdot \frac{D_{ac}}{D_{ad}}$$

it follows from (8.12) and (8.15) that

$$\frac{\lambda_{ad} - \lambda_{bd}}{\lambda_{ab} - \lambda_{bd}} = \frac{\lambda_{ac} - \lambda_{bc}}{\lambda_{ab} - \lambda_{bc}} \cdot \frac{\lambda_{ad} - \lambda_{cd}}{\lambda_{ac} - \lambda_{cd}},$$

which yields (8.7). Hence, $\succcurlyeq$ satisfies strong transitivity and the line property, which was to show.

(b) Assume now that $\succcurlyeq$ is regular, and satisfies strong transitivity and the line property. Take an arbitrary choice a. We construct a conditional preference relation $\succcurlyeq'$ that uniformly increases the preference for a relative to $\succcurlyeq$, distinguishing two cases: (1) there are some $b, c \neq a$ with $P_{a \sim b} \neq P_{a \sim c}$, and (2) $P_{a \sim b} = P_{a \sim c}$ for all $b, c \neq a$.

Case 1. Suppose there are some $e, f \neq a$ with $P_{a\sim e} \neq P_{a\sim f}$. Since $\succcurlyeq$ has preference reversals on all pairs of choices we know by Lemma 3.1 (b) that every set $\langle P_{b\sim c}\rangle$ is a hyperplane. By the same argument as in (a) we can find a line $L = \{v^* + \lambda w^* \mid \lambda \in \mathbf{R}\}$ that intersects each of the sets $\langle P_{b\sim c}\rangle$ exactly once, say at $v_{bc} = v^* + \lambda_{bc} w^*$, and such that $\lambda_{bc} \neq \lambda_{bd}$ whenever $\langle P_{b\sim c}\rangle \neq \langle P_{b\sim d}\rangle$.

For the purpose of this proof we define, for every $b, c \neq a$ with $\lambda_{ac} \neq \lambda_{bc}$, the number

$$D_{abc} := \frac{\lambda_{ab} - \lambda_{bc}}{\lambda_{ac} - \lambda_{bc}}. \tag{8.16}$$

As $\lambda_{ac} \neq \lambda_{bc}$, it follows by transitivity that $\lambda_{ab}, \lambda_{ac}$ and λ_{bc} are pairwise different, and hence $D_{abc} \neq 0$. By the line property we have that

$$D_{abd} = D_{abc} \cdot D_{acd} \tag{8.17}$$

for all $b, c, d \neq a$ with $\lambda_{ad} \neq \lambda_{bd}, \lambda_{ac} \neq \lambda_{bc}$ and $\lambda_{ad} \neq \lambda_{cd}$.

We will now define a hyperplane $H_{a\sim' b}$ for every choice $b \neq a$, as follows. Recall that choices $e, f \neq a$ are such that $P_{a\sim e} \neq P_{a\sim f}$. As $\succcurlyeq$ has preference reversals on $\{a, e\}$, there is some belief $p^* \in P_{e\succ a}$. Since, by Lemma 3.1 (a), $P_{a\sim e} = \langle P_{a\sim e}\rangle \cap \Delta(X)$, we know that $p^* \notin \langle P_{a\sim e}\rangle$. Let $H_{a\sim' e}$ be the unique hyperplane that is parallel to $\langle P_{a\sim e}\rangle$ and that passes through p^*. Then, the line L intersects $H_{a\sim' e}$ at a single point, say $v'_{ae} = v^* + \lambda'_{ae} w^*$. Note that $\lambda'_{ae} \neq \lambda_{ae}$, since $H_{a\sim' e}$ is parallel to, but different from, $\langle P_{a\sim e}\rangle$.

As $P_{a\sim e} \neq P_{a\sim f}$ and, by Lemma 3.1 (a), we have that $P_{a\sim e} = \langle P_{a\sim e}\rangle \cap \Delta(X)$ and $P_{a\sim f} = \langle P_{a\sim f}\rangle \cap \Delta(X)$, it follows that $\langle P_{a\sim e}\rangle \neq \langle P_{a\sim f}\rangle$. By the way we have chosen the line L we then conclude that $\lambda_{ae} \neq \lambda_{af}$. Thus, $v_{ae} \in \langle P_{a\sim e}\rangle \setminus \langle P_{a\sim f}\rangle$. Since, by strong transitivity, $\langle P_{a\sim e}\rangle \cap \langle P_{e\sim f}\rangle \subseteq \langle P_{a\sim f}\rangle$, we conclude that $v_{ae} \notin \langle P_{e\sim f}\rangle$ and hence $\lambda_{ae} \neq \lambda_{ef}$. Similarly, $\lambda_{af} \neq \lambda_{ef}$. That is, $\lambda_{ae}, \lambda_{af}$ and λ_{ef} are pairwise different. Let λ'_{af} be the unique number such that

$$\lambda'_{ae} - \lambda_{ae} = D_{aef} \cdot (\lambda'_{af} - \lambda_{af}), \tag{8.18}$$

and let $H_{a\sim' f}$ be the unique hyperplane that is parallel to $\langle P_{a\sim f}\rangle$ and that passes through $v'_{af} = v^* + \lambda'_{af} w^*$. Note that D_{aef} is well-defined and not equal to 0 as $\lambda_{ae}, \lambda_{af}$ and λ_{ef} are pairwise different. This is exactly the way we defined λ'_{ae} and λ'_{af} in the proof sketch of Section 6.2.

Next, take a choice $b \neq a, e, f$. Then either $\lambda_{ab} \neq \lambda_{ae}$ or $\lambda_{ab} \neq \lambda_{af}$. Assume first that $\lambda_{ab} \neq \lambda_{ae}$. By strong transitivity we can then show, in a similar way as above, that $\lambda_{ab}, \lambda_{ae}$ and λ_{be} are pairwise different. Let λ'_{ab} be the unique number such that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot (\lambda'_{ae} - \lambda_{ae}) \tag{8.19}$$

and let $H_{a\sim' b}$ be the unique hyperplane that is parallel to $\langle P_{a\sim b}\rangle$ and that passes through $v'_{ab} = v^* + \lambda'_{ab} w^*$.

If $\lambda_{ab} = \lambda_{ae}$ then it must be $\lambda_{ab} \neq \lambda_{af}$. By strong transitivity, it then follows that $\lambda_{ab}, \lambda_{af}$ and λ_{bf} are pairwise different. Let λ'_{ab} be the unique number such that

$$\lambda'_{ab} - \lambda_{ab} = D_{abf} \cdot (\lambda'_{af} - \lambda_{af}), \tag{8.20}$$

and let $H_{a\sim' b}$ be the unique hyperplane that is parallel to $\langle P_{a\sim b}\rangle$ and that passes through $v'_{ab} = v^* + \lambda'_{ab} w^*$.

We now show that, for all $b, c \neq a$ with $\lambda_{ab} \neq \lambda_{ac}$,

$$\lambda'_{ab} - \lambda_{ab} = D_{abc} \cdot (\lambda'_{ac} - \lambda_{ac}). \tag{8.21}$$

In view of (8.18), (8.19) and (8.20) it only remains to show (8.21) for the case where $c = f$ and $\lambda_{ab} \neq \lambda_{ae}$, and for the case where $b, c \neq e, f$.

Consider first the case where $c = f$ and $\lambda_{ab} \neq \lambda_{ae}$. Then we have, by (8.18) and (8.19), that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot (\lambda'_{ae} - \lambda_{ae}) \text{ and } \lambda'_{ae} - \lambda_{ae} = D_{aef} \cdot (\lambda'_{af} - \lambda_{af}),$$

which implies that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot D_{aef} \cdot (\lambda'_{af} - \lambda_{af}).$$

As, by the line property, $D_{abe} \cdot D_{aef} = D_{abf}$, we obtain that $\lambda'_{ab} - \lambda_{ab} = D_{abf} \cdot (\lambda'_{af} - \lambda_{af})$, which was to show.

Suppose next that $b, c \neq e, f$. If $\lambda_{ab} \neq \lambda_{ae}$ and $\lambda_{ac} \neq \lambda_{ae}$, then it follows from (8.19) that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot (\lambda'_{ae} - \lambda_{ae}) \text{ and } \lambda'_{ac} - \lambda_{ac} = D_{ace} \cdot (\lambda'_{ae} - \lambda_{ae})$$

and hence

$$\lambda'_{ab} - \lambda_{ab} = \frac{D_{abe}}{D_{ace}} \cdot (\lambda'_{ac} - \lambda_{ac}).$$

As, by definition, $D_{ace} = \frac{1}{D_{aec}}$, it follows that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot D_{aec} \cdot (\lambda'_{ac} - \lambda_{ac}) = D_{abc} \cdot (\lambda'_{ac} - \lambda_{ac})$$

since, by the line property, $D_{abe} \cdot D_{aec} = D_{abc}$.

If $\lambda_{ab} \neq \lambda_{ae}$ and $\lambda_{ac} = \lambda_{ae}$, then it follows from (8.19) and (8.20) that

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot (\lambda'_{ae} - \lambda_{ae}) \text{ and } \lambda'_{ac} - \lambda_{ac} = D_{acf} \cdot (\lambda'_{af} - \lambda_{af}).$$

Combined with (8.18) we get

$$\lambda'_{ab} - \lambda_{ab} = D_{abe} \cdot D_{aef} \cdot (\lambda'_{af} - \lambda_{af}) \text{ and } \lambda'_{ac} - \lambda_{ac} = D_{acf} \cdot (\lambda'_{af} - \lambda_{af}),$$

and hence

$$\lambda'_{ab} - \lambda_{ab} = \frac{D_{abe} \cdot D_{aef}}{D_{acf}} (\lambda'_{ac} - \lambda_{ac}).$$

As, by the line property, $D_{abe} \cdot D_{aef} = D_{abf}$, and $\frac{D_{abf}}{D_{acf}} = D_{abf} \cdot D_{afc} = D_{abc}$, it follows that $\lambda'_{ab} - \lambda_{ab} = D_{abc} \cdot (\lambda'_{ac} - \lambda_{ac})$.

The case where $\lambda_{ab} = \lambda_{ae}$ and $\lambda_{ac} \neq \lambda_{ae}$, and the case where $\lambda_{ab} = \lambda_{ae}$ and $\lambda_{ac} = \lambda_{ae}$ can be shown in a similar fashion as above. We have thus established (8.21) for every $b, c \neq a$ with $\lambda_{ab} \neq \lambda_{ac}$.

In this way we have constructed, for every $b \neq a$, a hyperplane $H_{a\sim' b}$. For every choice $b \neq a$ set $H_{a\sim b} := \langle P_{a\sim b}\rangle$, which is a hyperplane as well by Lemma 3.1 (b). Let $\succcurlyeq'$ be the conditional preference relation that coincides with $\succcurlyeq$ on every pair of choices $\{b, c\}$ with $b, c \neq a$, and such that for every choice $b \neq a$

$$P_{a\succ' b} = \{p \in \Delta(X) \mid \text{there is } \lambda > 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}, \tag{8.22}$$

$$P_{a\sim' b} = H_{a\sim' b} \cap \Delta(X), \text{ and} \tag{8.23}$$

$$P_{b\succ' a} = \{p \in \Delta(X) \mid \text{there is } \lambda < 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}. \tag{8.24}$$

By (8.18), (8.19), (8.20), the fact that $D_{abc} \neq 0$ for all $b, c \neq a$ with $\lambda_{ab} \neq \lambda_{ac}$, and the fact that $\lambda_{ae} \neq \lambda'_{ae}$, it follows that the hyperplane $H_{a\sim b}$ is always different from $H_{a\sim' b}$ for every $b \neq a$.

We will show $\succcurlyeq'$ is regular, and that these choices of $\succcurlyeq'$, $H_{a\sim b}$ and $H_{a\sim' b}$ satisfy properties (a) and (b) in Definition 4.1. We start by showing that

$$\langle P_{b\sim c}\rangle = \langle (\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle) \cup \{v_{bc}\}\rangle \tag{8.25}$$

for all $b, c \neq a$.

To see this, assume first that $\lambda_{ab} = \lambda_{ac}$. By the way we have chosen the line L it then follows that $\langle P_{a\sim b}\rangle = \langle P_{a\sim c}\rangle$, and hence, by transitivity, $\langle P_{a\sim b}\rangle = \langle P_{a\sim c}\rangle = \langle P_{b\sim c}\rangle$. Thus, $v_{bc} \in \langle P_{a\sim b}\rangle = \langle P_{a\sim c}\rangle$, and the (8.25) holds.

Assume next that $\lambda_{ab} \neq \lambda_{ac}$. Then, by transitivity, $\lambda_{ab}, \lambda_{ac}$ and λ_{bc} are pairwise different, which means that $v_{bc} \notin \langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$. Moreover, as $\lambda_{ab} \neq \lambda_{ac}$, we must have that $\langle P_{a\sim b}\rangle \neq \langle P_{a\sim c}\rangle$. Since, by Lemma 3.1 (b), $\dim(\langle P_{a\sim b}\rangle) = \dim(\langle P_{a\sim c}\rangle) = |X| - 1$, it follows that $\dim(\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle) = |X| - 2$. As $v_{bc} \notin \langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle$ we conclude that $\dim(\langle (\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle) \cup \{v_{bc}\}\rangle) = |X| - 1$. By strong transitivity we know that $\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle \subseteq \langle P_{b\sim c}\rangle$. Since $v_{bc} \in \langle P_{b\sim c}\rangle$ by construction, it follows that $\langle (\langle P_{a\sim b}\rangle \cap \langle P_{a\sim c}\rangle) \cup \{v_{bc}\}\rangle \subseteq \langle P_{b\sim c}\rangle$. As both linear subspaces have dimension $|X| - 1$, they must be equal, and hence (8.25) applies.

We next show that

$$\begin{aligned} \langle P_{b\sim c}\rangle = \{v \in \mathbf{R}^X \mid &\text{ there is some } \lambda \in \mathbf{R} \text{ with} \\ v \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} &\text{ and } v \in (1-\lambda)H_{a\sim' c} + \lambda H_{a\sim c}\} \end{aligned} \tag{8.26}$$

for all $b, c \neq a$.

We first show this property for all $b, c \neq a$ with $\lambda_{ab} = \lambda_{ac}$. By the choice of the line L, it must then be that $\langle P_{a\sim b}\rangle = \langle P_{a\sim c}\rangle$, and hence, by transitivity, $\langle P_{a\sim b}\rangle = \langle P_{a\sim c}\rangle = \langle P_{b\sim c}\rangle$. Let A be the righthand side of (8.26). Then,

$$A = \{v \in \mathbf{R}^X \mid v \in 0H_{a\sim' b} + 1H_{a\sim b} \text{ and } v \in 0H_{a\sim' c} + 1H_{a\sim c}\} = H_{a\sim b} \cap H_{a\sim c} = \langle P_{b\sim c}\rangle ,$$

which establishes (8.26).

We next show (8.26) for $b, c \neq a$ with $\lambda_{ab} \neq \lambda_{ac}$. Let $\lambda := \frac{\lambda_{ab}-\lambda_{bc}}{\lambda_{ab}-\lambda'_{ab}}$. Then, it may be verified that $v_{bc} = (1-\lambda)v_{ab} + \lambda v'_{ab}$. On the basis of (8.21) we may conclude that $\lambda = \frac{\lambda_{ac}-\lambda_{bc}}{\lambda_{ac}-\lambda'_{ac}}$ also, and hence it follows similarly that $v_{bc} = (1-\lambda)v_{ac} + \lambda v'_{ac}$. As such, there is some number $\lambda \neq 0$ with

$$v_{bc} \in \lambda H_{a\sim' b} + (1-\lambda)H_{a\sim b} \text{ and } v_{bc} \in \lambda H_{a\sim' c} + (1-\lambda)H_{a\sim c}. \tag{8.27}$$

Let A be the set on the righthand side of (8.26). To show that $\langle P_{b\sim c}\rangle \subseteq A$, take some $v \in \langle P_{b\sim c}\rangle$. Then, by (8.25), there is some vector $w \in H_{a\sim b} \cap H_{a\sim c}$ and some number α such that $v = w + \alpha v_{bc}$. Together with (8.27) it then follows that

$$v \in \alpha\lambda H_{a\sim' b} + (1-\alpha\lambda)H_{a\sim b} \text{ and } v \in \alpha\lambda H_{a\sim' c} + (1-\alpha\lambda)H_{a\sim c}$$

and hence $v \in A$.

To show that $A \subseteq \langle P_{b\sim c}\rangle$ take some $v \in A$. Hence, there is some number μ with

$$v \in (1-\mu)H_{a\sim' b} + \mu H_{a\sim b} \text{ and } v \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}. \tag{8.28}$$

We show that

$$w := \lambda v - (1-\mu)v_{bc} \in H_{a\sim b} \cap H_{a\sim c}. \tag{8.29}$$

By (8.27) and (8.28), there are $h_{ab}, k_{ab} \in H_{a\sim b}$ and $h'_{ab}, k'_{ab} \in H_{a\sim' b}$ such that $v = (1-\mu)h'_{ab} + \mu h_{ab}$ and $v_{bc} = \lambda k'_{ab} + (1-\lambda)k_{ab}$. Hence,

$$\begin{aligned} w &= \lambda((1-\mu)h'_{ab} + \mu h_{ab}) - (1-\mu)(\lambda k'_{ab} + (1-\lambda)k_{ab}) \\ &= \lambda\mu h_{ab} - (1-\lambda)(1-\mu)k_{ab} + \lambda(1-\mu)(h'_{ab} - k'_{ab}). \end{aligned}$$

As the hyperplanes $H_{a\sim' b}$ and $H_{a\sim b}$ are parallel and $H_{a\sim b} = \langle P_{a\sim b}\rangle$ is a linear subspace, it follows that $h'_{ab} - k'_{ab} \in H_{a\sim b}$ and hence $w \in H_{a\sim b}$. In a similar way it can be shown that $w \in H_{a\sim c}$, and hence we have established (8.29).

As $\lambda \neq 0$ it follows from (8.29) that $v = \frac{1}{\lambda}w + \frac{1-\mu}{\lambda}v_{bc} \in \langle P_{b\sim c}\rangle$ by (8.25) and (8.29). We have thus established (8.26) for all $b, c \neq a$.

We now show property (a) in Definition 4.1. Take some $b \neq a$. Note that, by (8.22) and (8.23),

$$P_{a\succcurlyeq' b} = \{p \in \Delta(X) \mid \text{there is } \lambda \geq 0 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}.$$

Hence, it remains to show that

$$P_{a\succcurlyeq b} = \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}. \tag{8.30}$$

We first show (8.30) for $b = e$. Recall that we have chosen the belief p^* such that $p^* \in P_{e\succ a}$ and $p^* \in H_{a\sim' e}$. As $H_{a\sim e} \cap \Delta(X) = P_{a\sim e}$, it follows by Lemma 8.2 (b) that (8.30) obtains for $b = e$.

We now show (8.30) for $b = f$. Recall that $P_{a\sim e} \neq P_{a\sim f}$ and hence $\langle P_{a\sim e}\rangle \neq \langle P_{a\sim f}\rangle$. Then, by strong transitivity, $\langle P_{a\sim e}\rangle \neq \langle P_{e\sim f}\rangle$ and $\langle P_{a\sim f}\rangle \neq \langle P_{e\sim f}\rangle$. Hence, there is some $p_{ef} \in P_{e\sim f}$ with $p_{ef} \notin P_{a\sim e} \cup P_{a\sim f}$. Then, by (8.26), there is some λ such that

$$p_{ef} \in (1-\lambda)H_{a\sim' e} + \lambda H_{a\sim e} \text{ and } p_{ef} \in (1-\lambda)H_{a\sim' f} + \lambda H_{a\sim f}$$

where $\lambda \neq 1$.

If $\lambda > 1$ then it follows from (8.30) for $b = e$ that $p_{ef} \in P_{a\succ e}$. As $p_{ef} \in P_{e\sim f}$ we conclude that $p_{ef} \in P_{a\succ f}$. Since $p_{ef} \in (1-\lambda)H_{a\sim' f} + \lambda H_{a\sim f}$ and $\lambda > 1$ we conclude from Lemma 8.2 that $(1-\mu)H_{a\sim' f} + \mu H_{a\sim f} \cap \Delta(X) \subseteq P_{f\succ a}$ for all $\mu < 1$. Since $\succcurlyeq$ has preference reversals on $\{a, f\}$, there is some $\mu < 1$ such that $(1-\mu)H_{a\sim' f} + \mu H_{a\sim f} \cap \Delta(X)$ is non-empty. Take some $q \in (1-\mu)H_{a\sim' f} + \mu H_{a\sim f}$. Then, $q \in P_{f\succ a}$. Hence, by Lemma 8.2 (b), (8.30) holds for $b = f$.

If $\lambda < 1$ then it follows from (8.30) for $b = e$ that $p_{ef} \in P_{e\succ a}$. As $p_{ef} \in P_{e\sim f}$ we conclude that $p_{ef} \in P_{f\succ a}$. Since $p_{ef} \in (1-\lambda)H_{a\sim' f} + \lambda H_{a\sim f}$ and $\lambda < 1$ we conclude from Lemma 8.2 that (8.30) holds for $b = f$.

We finally show (8.30) for all $b \neq e, f$. Since $P_{a\sim e} \neq P_{a\sim f}$ it follows that either $P_{a\sim b} \neq P_{a\sim e}$ or $P_{a\sim b} \neq P_{a\sim f}$. Assume that $P_{a\sim b} \neq P_{a\sim e}$, which implies by strong transitivity that $\langle P_{a\sim b}\rangle \neq \langle P_{b\sim e}\rangle$ and $\langle P_{a\sim e}\rangle \neq \langle P_{b\sim e}\rangle$. Hence, we can choose some $p_{be} \in P_{b\sim e}$ with $p_{be} \notin P_{a\sim b} \cup P_{a\sim e}$. Then, by (8.26), there is some λ such that

$$p_{be} \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p_{be} \in (1-\lambda)H_{a\sim' e} + \lambda H_{a\sim e}$$

where $\lambda \neq 1$. But then, we can show in the same way as above that (8.30) holds. If $P_{a\sim b} \neq P_{a\sim f}$ the proof will be similar since we know that (8.30) holds for $b = f$. We have thus established (8.30) for all $b \neq a$, and hence we have shown property (a) in Definition 4.1.

We continue by proving property (b) in Definition 4.1. Hence, we must show for all $b, c \neq a$ that

$$\begin{aligned} P_{b\succcurlyeq c} = \{p \in \Delta(X) \mid & \text{ there are } \lambda \leq \mu \text{ with} \\ & p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}. \end{aligned} \tag{8.31}$$

Note that by (8.26) and Lemma 3.1 (b) we know that

$$\begin{aligned} P_{b\sim c} = \{p \in \Delta(X) \mid & \text{ there is } \lambda \text{ with} \\ & p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\lambda)H_{a\sim' c} + \lambda H_{a\sim c}\}. \end{aligned} \tag{8.32}$$

We distinguish two cases: (i) $P_{a\sim b} \neq P_{a\sim c}$, and (ii) $P_{a\sim b} = P_{a\sim c}$.

(i) Assume first that $P_{a\sim b} \neq P_{a\sim c}$. As $\succcurlyeq$ has preference reversals on $\{a, b\}$ and $\{a, c\}$, there is some $p \in P_{a\sim b} \backslash P_{a\sim c}$. Suppose that $p \in P_{a\succ c}$. By transitivity, $p \in P_{b\succ c}$. Moreover, since $p \in P_{a\sim b} \cap P_{a\succ c}$ we know by (8.30) that $p \in H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$ for some $\mu > 1$. Hence, we have found some $p \in P_{b\succ c}$ with $1 = \lambda < \mu$ such that $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$. As $\succcurlyeq$ is regular, it follows from (8.32) that

$$\begin{aligned} P_{b\succ c} = \{p \in \Delta(X) \mid & \text{ there are } \lambda < \mu \text{ with} \\ & p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}. \end{aligned}$$

Together with (8.32) we thus obtain (8.31). If $p \in P_{c \succ a}$, we can use a similar argument to show (8.31).

(ii) Assume next that $P_{a \sim b} = P_{a \sim c}$. Recall that $P_{a \sim e} \neq P_{a \sim f}$. Hence, $P_{a \sim b} \neq P_{a \sim e}$ or $P_{a \sim b} \neq P_{a \sim f}$. Assume, without loss of generality, that $P_{a \sim b} \neq P_{a \sim e}$.

As $P_{a \sim b} = P_{a \sim c}$, it follows by transitivity of $\succcurlyeq$ that $P_{a \sim b} = P_{a \sim c} = P_{b \sim c}$. We also have that $P_{b \sim e}$ and $P_{c \sim e}$ are different. To see this, assume on the contrary that $P_{b \sim e} = P_{c \sim e}$. Then, by transitivity of $\succcurlyeq$ we would have that $P_{b \sim e} = P_{c \sim e} = P_{b \sim c}$. As $P_{b \sim c} = P_{a \sim b}$, it would follow that $P_{b \sim e} = P_{a \sim b}$, which would imply that $P_{a \sim b} = P_{a \sim e}$, which is a contradiction. Hence, we conclude that $P_{b \sim e}$ and $P_{c \sim e}$ are different.

Take some $p \in P_{b \sim e} \backslash P_{c \sim e}$. Then, by (8.32), there is some λ with

$$p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b} \text{ and } p \in (1-\lambda) H_{a \sim' e} + \lambda H_{a \sim e}. \tag{8.33}$$

As $p \notin P_{c \sim e}$, we must have that either $p \in P_{c \succ e}$ or $p \in P_{e \succ c}$.

Assume first that $p \in P_{c \succ e}$. As $p \in P_{b \sim e}$ it follows by transitivity that $p \in P_{c \succ b}$. Moreover, as $p \in P_{c \succ e}$ and $P_{a \sim c} = P_{a \sim b} \neq P_{a \sim e}$, we know by case (i) above and the second equation in (8.33) that

$$p \in (1-\mu) H_{a \sim' c} + \mu H_{a \sim c} \text{ for some } \mu < \lambda. \tag{8.34}$$

By combining (8.33) and (8.34) we have thus found some $p \in P_{c \succ b}$ with

$$p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b} \text{ and } p \in (1-\mu) H_{a \sim' c} + \mu H_{a \sim c} \text{ for some } \lambda > \mu. \tag{8.35}$$

As $\succcurlyeq$ is regular, it follows from (8.35) and (8.32) that

$$\begin{gathered} P_{c \succ b} = \{p \in \Delta(X) \mid \text{there are } \lambda > \mu \text{ with} \\ p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b} \text{ and } p \in (1-\mu) H_{a \sim' c} + \mu H_{a \sim c}\}. \end{gathered}$$

Together with (8.32) we thus obtain (8.31). If $p \in P_{e \succ c}$, we can use a similar argument as above to derive (8.31).

We will finally show that $\succcurlyeq'$ is regular. By (8.22), (8.23) and (8.24) it immediately follows that $\succcurlyeq'$ is complete and satisfies continuity, preservation of indifference and preservation of strict preference. It remains to show that $\succcurlyeq'$ is transitive.

Take some $b, c, d \in C$ and some $p \in P_{b \succcurlyeq' c} \cap P_{c \succcurlyeq' d}$. We must show that $p \in P_{b \succcurlyeq' d}$. If $b, c, d \neq a$ this holds because $\succcurlyeq'$ coincides with $\succcurlyeq$ on $\{b, c, d\}$ and $\succcurlyeq$ is transitive.

Assume now that $b = a$. As $p \in P_{a \succcurlyeq' c}$ it follows from (8.22) and (8.23) that $p \in (1-\lambda) H_{a \sim' c} + \lambda H_{a \sim c}$ for some $\lambda \geq 0$. Since $p \in P_{c \succcurlyeq' d} = P_{c \succcurlyeq d}$ we know from (8.31) that $p \in (1-\mu) H_{a \sim' d} + \mu H_{a \sim d}$ for some $\mu \geq \lambda$. By combining these two facts we conclude that $p \in (1-\mu) H_{a \sim' d} + \mu H_{a \sim d}$ for some $\mu \geq 0$. From (8.22) and (8.23) it then follows that $p \in P_{a \succcurlyeq' d}$, as was to show.

Assume next that $c = a$. As $p \in P_{b \succcurlyeq' a}$ it follows from (8.23) and (8.24) that $p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b}$ for some $\lambda \leq 0$. Since $p \in P_{a \succcurlyeq' d}$ we know from (8.22) and (8.23) that $p \in (1-\mu) H_{a \sim' d} + \mu H_{a \sim d}$ for some $\mu \geq 0$. By combining these two facts we conclude that $p \in (1-\lambda) H_{a \sim' b} + \lambda H_{a \sim b}$ and $p \in (1-\mu) H_{a \sim' d} + \mu H_{a \sim d}$ for some $\lambda \leq \mu$. By (8.31) it then follows that $p \in P_{b \succcurlyeq' d}$, as was to show.

Assume finally that $d = a$. As $p \in P_{b \succcurlyeq' c} = P_{b \succcurlyeq c}$ it follows from (8.31) that $p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}$ and $p \in (1-\mu)H_{a \sim' c} + \mu H_{a \sim c}$ for some $\lambda \leq \mu$. Since $p \in P_{c \succcurlyeq' a}$ it follows from (8.23) and (8.24) that $\mu \leq 0$. By combining these two facts we conclude that $p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}$ for some $\lambda \leq 0$. From (8.23) and (8.24) it then follows that $p \in P_{b \succcurlyeq' a}$, as was to show.

We thus conclude that $\succcurlyeq'$ is transitive. Summarizing, we see that $\succcurlyeq'$ is regular, and that conditions (a) and (b) in Definition 4.1 are satisfied. As such, $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$. This completes Case 1.

Case 2. Assume now that $P_{a \sim b} = P_{a \sim c}$ for all $b, c \neq a$. By transitivity it then follows that $P_{a \sim b} = P_{a \sim c} = P_{b \sim c}$ for all $b, c \neq a$. By Lemma 3.1 (b) we know, for every $b \neq a$, that $P_{a \sim b} = \langle P_{a \sim b} \rangle \cap \Delta(X)$ and that $\langle P_{a \sim b} \rangle$ is a hyperplane. Set $H := \langle P_{a \sim b} \rangle$ for some $b \neq a$. Hence, $H = \langle P_{a \sim b} \rangle$ for all $b \neq a$. Let H^+ and H^- be sets of vectors such that H separates H^+ from H^- and $H^+ \cup H^- \cup H = \mathbf{R}^X$.

As $\succcurlyeq$ has preference reversals on all pairs of choices, there are choices $b_1, b_2, ..., b_K, c_1, c_2, ..., c_M \neq a$ (where $\{b_1, ..., b_K\}$ or $\{c_1, ..., c_M\}$ may be empty) such that

$$b_1 \succ_p b_2 \succ_p ... \succ_p b_K \succ_p a \succ_p c_1 \succ_p c_2 \succ_p ... \succ_p c_M \text{ for all } p \in H^+ \tag{8.36}$$

and

$$c_M \succ_p c_{M-1} \succ_p ... \succ_p c_1 \succ_p a \succ_p b_K \succ_p b_{K-1} \succ_p ... \succ_p b_1 \text{ for all } p \in H^-. \tag{8.37}$$

Take a vector $v \in H$ and a vector $w \neq \underline{0}$ such that the line $L := \{v + \lambda w \mid \lambda \in \mathbf{R}\}$ intersects the hyperplane H only at v, and such that for every $p \in H^+ \cap L$ there is some $\lambda > 0$ with $p = v + \lambda w$. Choose a number λ'_{ab} for every $b \neq a$ such that

$$\lambda'_{ab_K} > \lambda'_{ab_{K-1}} > ... > \lambda'_{ab_1} > 0 > \lambda'_{ac_M} > \lambda'_{ac_{M-1}} > ... > \lambda'_{ac_1}. \tag{8.38}$$

For every $b \neq a$ let $H_{a \sim' b}$ be the unique hyperplane that is parallel to H and passes through $v'_{ab} := v + \lambda'_{ab} w$. Hence, by construction,

$$H_{a \sim' b} \subseteq H^+ \text{ for all } b \in \{b_1, ..., b_K\} \text{ and } H_{a \sim' b} \subseteq H^- \text{ for all } b \in \{c_1, ..., c_M\}. \tag{8.39}$$

Moreover, we can choose the numbers λ'_{ab} close enough to 0 such that $H_{a \sim' b} \cap \Delta(X)$ is nonempty for all $b \neq a$.

Set $H_{a \sim b} := H$. Let $\succcurlyeq'$ be the conditional preference relation that coincides with $\succcurlyeq$ on all pairs $\{b, c\}$ with $b, c \neq a$, and where for every $b \neq a$

$$P_{a \succ' b} = \{p \in \Delta(X) \mid \text{there is } \lambda > 0 \text{ with } p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}\}, \tag{8.40}$$

$$P_{a \sim' b} = H_{a \sim' b} \cap \Delta(X), \text{ and} \tag{8.41}$$

$$P_{b \succ' a} = \{p \in \Delta(X) \mid \text{there is } \lambda < 0 \text{ with } p \in (1-\lambda)H_{a \sim' b} + \lambda H_{a \sim b}\}. \tag{8.42}$$

We will show that $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$. That is, we must prove that $\succcurlyeq'$ is regular, and that the choice of $\succcurlyeq'$, $H_{a \sim b}$ and $H_{a \sim' b}$ satisfies the conditions (a) and (b) in Definition 4.1.

We start by showing property (a) in Definition 4.1. In view of (8.40) and (8.41) it only remains to show that

$$P_{a \succcurlyeq b} = \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\}. \tag{8.43}$$

Take some $b \neq a$, and assume first that $b \in \{b_1, ..., b_K\}$. Hence, by (8.39), $H_{a\sim' b} \subseteq H^+$. Moreover, by (8.36) and (8.37), $P_{a\succcurlyeq b} = H^- \cup H$. Since

$$H^- \cup H = \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\},$$

(8.43) obtains.

Take next some $b \in \{c_1, ..., c_M\}$. Hence, by (8.39), $H_{a\sim' b} \subseteq H^-$. Moreover, by (8.36) and (8.37), $P_{a\succcurlyeq b} = H^+ \cup H$. Since

$$H^+ \cup H = \{p \in \Delta(X) \mid \text{there is } \lambda \geq 1 \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}\},$$

(8.43) obtains. We thus conclude that property (a) in Definition 4.1 holds.

We now prove property (b) in Definition 4.1. Hence, we must show for every $b, c \neq a$ that

$$\begin{aligned} P_{b\succcurlyeq c} = \{&p \in \Delta(X) \mid \text{there are } \lambda \leq \mu \text{ with} \\ &p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}. \end{aligned} \tag{8.44}$$

Recall that $H_{a\sim b} = H$ for all $b \neq a$. Moreover, by construction, $H_{a\sim' b} \neq H_{a\sim' c}$ for all $b \neq c$. Hence,

$$\begin{aligned} \{p \in \Delta(X) \mid \text{there is } \lambda \text{ with } p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\lambda)H_{a\sim' c} + \lambda H_{a\sim c}\} \\ = H \cap \Delta(X) = P_{b\sim c}. \end{aligned} \tag{8.45}$$

Now, take some $b, c \neq a$. Suppose first that $p \in P_{b\succ c}$ for some $p \in H^+$. Then, by (8.36) and (8.38) there are $\lambda < \mu$ such that $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$. Hence, we have found some $p \in P_{b\succ c}$ such that there are $\lambda < \mu$ with $p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b}$ and $p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}$. Together with (8.45) it follows that

$$\begin{aligned} P_{b\succ c} = \{&p \in \Delta(X) \mid \text{there are } \lambda < \mu \text{ with} \\ &p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}, \end{aligned}$$

which, together with (8.45), implies (8.44).

Suppose next that $p \in P_{c\succ b}$ for some $p \in H^+$. In the same way as above, it then follows that

$$\begin{aligned} P_{c\succ b} = \{&p \in \Delta(X) \mid \text{there are } \lambda > \mu \text{ with} \\ &p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}. \end{aligned}$$

Together with (8.45) we conclude that

$$\begin{aligned} P_{b\succ c} = \{&p \in \Delta(X) \mid \text{there are } \lambda < \mu \text{ with} \\ &p \in (1-\lambda)H_{a\sim' b} + \lambda H_{a\sim b} \text{ and } p \in (1-\mu)H_{a\sim' c} + \mu H_{a\sim c}\}, \end{aligned}$$

which, together with (8.45), implies (8.44). We thus have shown property (b) in Definition 4.1.

It remains to show that $\succcurlyeq'$ is regular. This can be shown in the same way as in Case 1, using (8.40), (8.41), (8.42) and (8.44). We thus see that $\succcurlyeq'$ uniformly increases the preference for a relative to $\succcurlyeq$. 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. and A. Brandenburger (1995), Epistemic conditions for Nash equilibrium, *Econometrica* **63,** 1161–1180.

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

[4] Blume, L.E., Brandenburger, A. and E. Dekel (1991), Lexicographic probabilities and choice under uncertainty, *Econometrica* **59,** 61–79.

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

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

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

[8] 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.

[9] Halpern, J.Y. (2010), Lexicographic probability, conditional probability, and nonstandard probability, *Games and Economic Behavior* **68,** 155–179.

[10] Hammond, P.J. (1994), Elementary non-archimedean representations of probability for decision theory and games, In: Humphreys, P. (Ed.), Patrick Suppes: Scientific Philosopher, Vol. 1. Kluwer, Dordrecht, pp. 25–49.

[11] 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.

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

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

[14] Robinson, A. (1973), Function theory on some nonarchimedean fields, *The American Mathematical Monthly* **80,** 87–109.

[15] Savage, L.J. (1954, 1972), *The Foundation of Statistics,* Wiley, New York.

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