

ANDRÉS PEREA

A MODEL OF MINIMAL PROBABILISTIC BELIEF REVISION

ABSTRACT. In the literature there are at least two models for probabilistic belief revision: Bayesian updating and imaging [Lewis, D. K. (1973), *Counterfactuals*, Blackwell, Oxford; Gärdenfors, P. (1988), *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press, Cambridge, MA]. In this paper we focus on imaging rules that can be described by the following procedure: (1) Identify every state with some real valued vector of characteristics, and accordingly identify every probabilistic belief with an expected vector of characteristics; (2) For every initial belief and every piece of information, choose the revised belief which is compatible with this information and for which the expected vector of characteristics has minimal Euclidean distance to the expected vector of characteristics of the initial belief. This class of rules thus satisfies an intuitive notion of minimal belief revision. The main result in this paper is to provide an axiomatic characterization of this class of imaging rules.

KEY WORDS: belief revision, probabilistic beliefs, imaging

JEL CLASSIFICATION: C73, D81, D83

1. INTRODUCTION

Consider the following thought experiment: Suppose you may glimpse for a few seconds at a transparent urn which is filled with black balls and white balls. Afterwards, you must form a probabilistic belief about the number of black balls and white balls. More precisely, you are given three possible options:

- 30 black balls and 30 white balls (state a);
- 30 black balls and 20 white balls (state b);
- 20 black balls and 20 white balls (state c);

One of these options is true, you are said, but you do not know which one since the observation only gave you a vague impression of the number of balls of each color. Your task is to assign probabilities to each of these three options. You recall to have seen approximately the same number of black and white balls, but time was too short to count the balls you have seen. Based on these considerations, you deem the states a and c equally likely, but at the same time you deem the event “ a or c ” twice as likely as the state b . Hence, you assign probability $1/3$ to each of the states.

Now, you are told that state c is false. Your new task is to assign probabilities to the two remaining states a and b . How would you do this? *Bayesian updating* suggests that we should not change our belief about the relative likelihood of the two remaining states, and hence our revised belief by Bayesian updating would assign probability $1/2$ to states a and b . In other words, if we hear that c is false and use Bayesian updating, then we lose much of our confidence in the event that the number of black and white balls is equal. This seems counterintuitive, however, since you were rather confident that you have seen an equal number of black and white balls, but much less confident about the precise number of balls. Therefore, the information that c is false should not be a reason to drastically change your belief about the proportion of black and white balls. As an extreme case, one could decide to maintain the belief that the event “ a or c ” is twice as likely as state b , which would yield a revised belief in which you assign probability $2/3$ to a and probability $1/3$ to b .

In fact, the last way of revising your belief is an example of *imaging*, as proposed by Lewis (1973, 1976). The idea behind imaging is that, upon observing that some state x is impossible, you transfer the probability initially assigned to x completely towards the remaining state you deem most similar to x . In the example above, you deem state c much more similar to a than to b , since c and a share the same proportion of black and white balls, and the observation has given you a pretty good idea of this proportion. Therefore, if we hear that c is false and use imaging, then we transfer the probability

1/3 initially assigned to c completely towards the most similar state, a , and hence our revised belief would assign probability 2/3 to a and probability 1/3 to b .

The main difference between Bayesian updating and imaging in this context is thus that imaging makes explicit use of the perceived similarity between states when revising the belief, whereas Bayesian updating ignores this perceived similarity. Although the theory of imaging makes intuitive sense in contexts where similarity between states plays a prominent role, as the example above hopefully illustrates, it has received little attention, especially among decision theorists, game theorists and economists.

In Chapter 5 of his book, Gärdenfors (1988) extends Lewis' idea by allowing to transfer a part λ of the probability 1/3, initially assigned to c , towards state a and the remaining part $1 - \lambda$ to state b . Important is that these fractions λ and $1 - \lambda$ should be independent of the initial belief. That is, if the initial belief would not be (1/3, 1/3, 1/3) but (1/8, 3/8, 1/2), then the person should transfer the fraction λ of the probability 1/2 to a and the fraction $1 - \lambda$ of 1/2 to b , where λ is the *same* fraction as above. Gärdenfors' model is also known as *general imaging*.

A possible interpretation of λ is that it represents the similarity between c and a , as compared to the similarity between c and b . For instance, if λ is very close to 1, this indicates that the person deems c much more similar to a than to b , and for that reason shifts almost all the weight initially assigned to c towards a . If λ is close to 1/2, this indicates that the person deems c almost as similar to a as to b , and therefore redistributes the weight initially assigned to c almost equally among a and b . Lewis' model of imaging is thus a special case of Gärdenfors' model, by choosing λ equal to 0 or 1. Figure 1 graphically illustrates the differences between Bayesian updating, Lewis imaging and general imaging. The figure represents a situation with three states, a , b and c , where the person deems c more similar to a than to b . For three different initial beliefs, the picture indicates how the person would revise his belief in each of the three models, upon observing




Figure 1. Bayesian updating versus imaging.




Figure 2. Not every general imaging rule is plausible.

the event $\{a, b\}$. By the event $\{a, b\}$ we mean that the person observes that the true state must be in $\{a, b\}$.

In this paper we take the interpretation of λ as described above very seriously. Given this interpretation, however, not every general imaging rule seems reasonable. Suppose that the person initially holds the belief $(1/3, 1/3, 1/3)$, and would revise this belief to $(5/8, 3/8, 0)$ when observing the event $\{a, b\}$, would revise the belief to $(3/8, 0, 5/8)$ when observing $\{a, c\}$, and would revise the belief to $(0, 5/8, 3/8)$ when observing $\{b, c\}$. See Figure 2 for an illustration. These belief revisions indicate that, in the person's mind,

1. the distance between a and c is smaller than the distance between b and c ,
2. the distance between b and c is smaller than the distance between a and b ,
3. the distance between a and b is smaller than the distance between a and c ,

assuming that the person uses a *symmetric* distance measure to visualize his notion of similarity between states.¹ This is counterintuitive since 1 and 2 seem to suggest that the person should deem the distance between a and c smaller than the distance between a and b .

In order to overcome this problem, we assume in this paper that the similarity relation between states, used to govern the person's belief revisions, should be given by Euclidean distances between points in some Euclidean space. We thus take the notion of "perceived distance" between states very literally. More precisely, we assume that the person identifies each state with some real valued vector of characteristics, describing its main properties, and measures the distance between two states by taking the Euclidean distance between the corresponding vectors of characteristics. Every belief, being a probability distribution over states, thus corresponds to an *expected* vector of characteristics, and the distance between two beliefs can be measured by taking the Euclidean distance between the corresponding expected vectors of characteristics. The general imaging rule induced by it assigns to every initial belief and every observable event (that is, a subset of states) the belief that (1) assigns positive probability only to states in the event, and (2) for which the expected vector of characteristics has minimal Euclidean distance to the expected vector of characteristics of the initial belief. Important is that these vectors of characteristics should be independent of the initial belief and independent of the observed event. That is, the perceived similarity between states and beliefs should be an inherent characteristic of the person, and should therefore not depend on his initial belief, nor on the information he receives.




Figure 3. A non-plausible and a plausible general imaging rule.

Yet, not every imaging rule of this kind is plausible. Suppose the person identifies the state a with the vector of characteristics $\varphi([a]) = (0, 0)$, identifies b with $\varphi([b]) = (0, 1)$, and c with $\varphi([c]) = (2, 2)$. See the first triangle of Figure 3 for an illustration. If the person holds the initial belief $(1/4, 1/4, 1/2)$ and observes $\{a, b\}$, the revised belief would be the opinionated² belief $[b]$, assigning probability 1 to b . This is unreasonable, since there should be no reason to take away the probability $1/4$ from state a as long as the information does not rule out a . In general, if the initial belief assigns probability α to state x , and the observed event does not rule out x , the revised probability attached to x should be at least α . This property will be guaranteed if we require that the triangle in Figure 3, representing the possible expected vectors of characteristics, only has acute angles. In this case, the revised belief upon observing $\{a, b\}$ will be obtained by first taking the expected vector of characteristics of the initial belief, and then taking the orthogonal projection of this vector on the line segment through $\varphi([a])$ and $\varphi([b])$. See the second triangle of Figure 3. The same holds for the revised belief upon observing the events $\{a, c\}$ and $\{b, c\}$. In general, the requirement is that the mapping φ , assigning an expected vector of characteristics $\varphi(\beta)$ to every belief β , is “closed under orthogonal projections”. By the latter we mean that for every state $[x]$ and every subset of states $E = \{x_1, \dots, x_k\}$, the orthogonal projection of $\varphi([x])$ on the affine space spanned by $\{\varphi([x_1]), \dots, \varphi([x_k])\}$ should be in the

convex hull of $\{\varphi([x_1]), \dots, \varphi([x_k])\}$. This property guarantees that, for every initial belief β and observed event $E = \{x_1, \dots, x_k\}$, the revised belief will always be obtained by first taking $\varphi(\beta)$, and then taking the orthogonal projection of $\varphi(\beta)$ on the convex hull of $\{\varphi([x_1]), \dots, \varphi([x_k])\}$.³

The road taken so far has thus led us to the following class of general imaging rules:

1. Construct a linear one-to-one mapping φ , which uniquely identifies every belief β with some expected vector of characteristics $\varphi(\beta)$.
2. Make sure that φ is closed under orthogonal projections.
3. For every initial belief β and observed event E , choose the revised belief β' that attaches only positive probability to states in E , and for which $\varphi(\beta')$ has minimal Euclidean distance to $\varphi(\beta)$.

For the remainder of this paper, we shall refer to this class as *Euclidean distance based imaging rules*. Such rules thus take the idea of minimal belief revision very literally, as they always look for the expected vector of characteristics that is as close as possible, in the sense of Euclidean distance, to the expected vector of characteristics of the initial belief. Schulte (2002) provides an excellent discussion of the idea of minimal belief revision, and gives an overview of the various ways in which this idea has been implemented in belief revision theory.

The main result in this paper, Theorem 3.3, is to show that the class of Euclidean distance based imaging rules can be characterized by three axioms: *linearity*, *transitivity* and *information-order independence*.

Linearity states that for a fixed observable event $\{a, b\}$ with two states, the revised belief should depend linearly upon the initial belief. In the example of Figure 1, this means that for the fixed event $\{a, b\}$, the arrows connecting the initial beliefs with their corresponding revised beliefs should all be parallel. Intuitively, linearity can be defended as follows. Consider three different states a, b and c and a person who initially assigns probability α to c . If this person finds out later that

the true state must be in $\{a, b\}$, he should shift some fraction $\lambda\alpha$ towards a , and the remaining fraction $(1 - \lambda)\alpha$ towards b . Linearity states that these fractions λ and $1 - \lambda$ should be independent of the initial belief, since λ reflects the perceived similarity between c and a , as compared to the perceived similarity between c and b . Among the three axioms, linearity is the one that rules out Bayesian updating as a possible candidate. It is easily seen that Bayesian updating violates linearity, but satisfies the other two axioms to be defined below.

The axiom of linearity is almost identical to Gärdenfors' (1988) *homomorphism* condition. There are two subtle differences. First, the homomorphism condition requires that the revised belief should also depend linearly upon the initial belief if the observed event contains more than two states. The reason why we do not choose this stronger version is that the intuition behind linearity, as given in the previous paragraph, would not work as nicely if we consider an event with more than two states. For instance, if there are four states $\{a, b, c, d\}$ and the initial belief $[d]$ is mapped to $(1/8, 3/8, 1/2, 0)$ upon observing the event $\{a, b, c\}$, the numbers $1/8, 3/8$ and $1/2$ cannot be interpreted that easily. For instance, these numbers do not automatically imply that d is deemed more similar to b than to a , since it may well be that the initial belief is mapped to $(5/8, 3/8, 0, 0)$ upon observing the event $\{a, b\}$. The second difference is that our linearity axiom assumes that if the initial belief assigns probability 1 to a , then the revised belief upon observing $\{a, b\}$ should also assign probability 1 to a . Formally, the homomorphism condition does not require this. In a non-probabilistic framework, Katsuno and Mendelzon (1992) formulate a condition called *disjunction rule* which is similar in spirit to linearity and the homomorphism condition.

While linearity basically makes sure that the belief revision rule is a general imaging rule (at least if we restrict to observable events with two states), the other two axioms guarantee that the perceived similarity between beliefs can be represented by expected vectors of characteristics. The second axiom, transitivity, states that the "equally similar to"

relation between beliefs should be transitive. That is, whenever a person deems the belief β equally similar to the opinionated beliefs $[a]$ and $[b]$, and deems β equally similar to $[b]$ and $[c]$, then he should deem β equally similar to $[a]$ and $[c]$. In terms of belief revision, this means that if the revised belief upon observing $\{a, b\}$ would assign equal probabilities to a and b , and the revised belief upon observing $\{b, c\}$ would assign equal probabilities to b and c , then the revised belief upon observing $\{a, c\}$ should assign equal probabilities to a and c . This condition rules out intuitive problems as in Figure 2.

The last axiom, information-order independence, states that the revised belief should not depend on the order in which information is received. For instance, revising the belief upon observing the event $\{x_1, x_2\}$ at once should give the same result as first observing the event $\{x_1, x_2, x_3\}$, and then observing the event $\{x_1, x_2\}$. This axiom thus guarantees that iterated belief revision does not lead to problems. Majumdar (2004) calls this condition path independence.

It is easy to show that every Euclidean distance based imaging rule satisfies these three axioms. Showing the opposite direction, namely that every belief revision function satisfying these axioms is an Euclidean distance based imaging rule, turned out to be a very difficult task. The complete proof involves many steps. In order to help the reader, we provide in Section 4 an overview of the proof of our main theorem. In that section, we present the lemmas that are used to show the main result, and present for each lemma an intuitive, often geometrical, argument. The full algebraic proofs are presented in the Appendix. The reader is advised to first read the intuitive arguments before going through the formal proof.

The outline of the paper is as follows. In Section 2, we introduce the concept of a belief revision function and the axioms of linearity, transitivity and information-order independence. In Section 3, we present the main theorem. Section 4 provides a non-technical overview of the proof. We conclude in Section 5 with a discussion.

2. BELIEF REVISION FUNCTIONS

In this section we introduce the notion of a belief revision function, and propose the axioms of *linearity*, *transitivity* and *information-order independence*. Let $X = \{x_1, x_2, \dots, x_n\}$ be some finite set of states. A *probabilistic belief* (or simply *belief*) on X is a probability distribution on X . We denote by $B(X)$ the set of all beliefs on X . For two beliefs β_1, β_2 in $B(X)$ and a number $\lambda \in [0, 1]$, let $\lambda\beta_1 + (1 - \lambda)\beta_2$ be the belief that assigns to every state $x \in X$ the probability $\lambda\beta_1(x) + (1 - \lambda)\beta_2(x)$. For a given state x , let $[x]$ be the opinionated belief that assigns probability one to x . A subset $E \subseteq X$ of states is called an *event*. By $B(X|E)$ we denote the set of beliefs on X that assign positive probability only to states in E .

DEFINITION 2.1. (Belief revision function). *A belief revision function on X is a function br that assigns to every belief $\beta \in B(X)$ and every event $E \subseteq X$ some belief $br(\beta|E) \in B(X|E)$.*

Here, β represents the *initial* belief a person holds about the state in X , whereas $br(\beta|E)$ represents the *revised* belief after receiving the information that the state is in E .

In order to introduce the axiom of linearity, assume that the decision maker initially assigns probability one to some state c , but later observes that the true state is in $\{a, b\}$. Suppose that, after observing this event, he attaches probability $\alpha > 1/2$ to a and probability $1 - \alpha < 1/2$ to b , that is, $br([c]|\{a, b\}) = \alpha[a] + (1 - \alpha)[b]$. Then, the revised belief reveals that the decision maker deems c more similar to a than to b , since he shifts more weight towards a than towards b . Moreover, the precise numbers α and $1 - \alpha$ reveal *how much* more similar he deems c to a compared to b : the higher α , the higher the perceived similarity between c and a compared to the perceived similarity between c and b . Consider now a different situation in which he initially assigns probability $\beta(c) < 1$ to c , and later observes the event $\{a, b\}$. He then faces a similar task as before, since he must shift the total weight $\beta(c)$ initially assigned to c towards the non-excluded states

a and b . Since we may assume that the relative perceived similarities between c and a as compared to c and b are still given by the numbers α and $1 - \alpha$ above, he should shift the fraction α of $\beta(c)$ to a , and shift the fraction $1 - \alpha$ of $\beta(c)$ to b . That is, the fractions of weight $\beta(c)$ that are shifted towards a and b , respectively, are given by the revised belief $br([c]|{a, b})$, and are independent of the initial weight $\beta(c)$ assigned to the excluded state c . By applying this reasoning to every state x that is excluded by the event $\{a, b\}$, we obtain the axiom of linearity.

AXIOM 2.2. (Linearity). *For every initial belief $\beta \in B(X)$ and every two states $a, b \in X$, it should hold that*

$$br(\beta|{a, b}) = \beta(a)[a] + \beta(b)[b] + \sum_{x \in X \setminus \{a, b\}} \beta(x)br([x]|{a, b}).$$

The second axiom, transitivity, states that, whenever a belief β is perceived “equally similar” to the opinionated beliefs $[a]$ and $[b]$, and is perceived “equally similar” to $[b]$ and $[c]$, then the belief β should be perceived “equally similar” to $[a]$ and $[c]$. Here, this “equally similar to” relation may be deduced from the belief revision function. Suppose, namely, that the initial belief $\beta \in B(X)$ and the states $a, b \in X$ are such that $br(\beta|{a, b}) = \frac{1}{2}[a] + \frac{1}{2}[b]$. Then, the revision of the belief β upon observing that all states but a and b are excluded is exactly halfway between $[a]$ and $[b]$, and hence, intuitively, the initial belief β was deemed equally similar to the opinionated beliefs $[a]$ and $[b]$. The transitivity axiom simply imposes that this “equally similar to” relation be transitive.

AXIOM 2.3. (Transitivity). *For every initial belief $\beta \in B(X)$ and every three different states $a, b, c \in X$ for which $br(\beta|{a, b}) = \frac{1}{2}[a] + \frac{1}{2}[b]$ and $br(\beta|{b, c}) = \frac{1}{2}[b] + \frac{1}{2}[c]$, it should also hold that $br(\beta|{a, c}) = \frac{1}{2}[a] + \frac{1}{2}[c]$.*

Geometrically speaking, within the state space $X = \{a, b, c\}$ the transitivity axiom connects the direction of revision upon receiving information $\{a, c\}$ to the revision directions with respect to the events $\{a, b\}$ and $\{b, c\}$. Figure 4 provides an




Figure 4. Transitivity.

illustration of this fact. The last axiom, information-order independence, states that it should not matter in which particular form information is received. For instance, it should not make a difference whether the information that states a and b are no longer possible is received at once, or that first the information excluding a is received followed by the information excluding b .

AXIOM 2.4. (Information-order independence). *For every initial belief $\beta \in B(X)$ and every two events $E_1, E_2 \subseteq X$ with $E_2 \subseteq E_1$, it holds that*

$$br(\beta|E_2) = br(br(\beta|E_1)|E_2).$$

Figure 5 illustrates this axiom for the case of four states, with state space $X = \{a, b, c, d\}$, and events $E_1 = \{a, b, c\}$ and $E_2 = \{b, c\}$.

3. REPRESENTATION THEOREM

Let the set X contain n states. Let φ be a one-to-one function from the set $B(X)$ of beliefs to some Euclidean space \mathbb{R}^m , with $m \geq n - 1$.




Figure 5. Information-order independence.

DEFINITION 3.1. (Linear function). *The function $\varphi: B(X) \rightarrow \mathbb{R}^m$ is called linear if for every belief $\beta \in B(X)$ we have that*

$$\varphi(\beta) = \sum_{x \in X} \beta(x) \varphi([x]).$$

In this paper, we interpret a linear function φ as follows: For every state x , the vector $\varphi([x])$ is chosen as a *vector of characteristics* for x . If the function φ is linear, the vector $\varphi(\beta)$ for a given belief β may then be interpreted as the *expected vector of characteristics* induced by the belief β and the vectors $\{\varphi([x]) \mid x \in X\}$. For every two beliefs β and β' , let $\|\varphi(\beta) - \varphi(\beta')\|$ denote the Euclidean distance between the corresponding expected vectors of characteristics.

DEFINITION 3.2. (Closed under orthogonal projections). *A linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$ is called closed under orthogonal projections if for every $\beta_1 \in B(X)$ and every $E \subseteq X$ there is some $\beta_2 \in B(X|E)$ such that*

$$(\varphi(\beta_2) - \varphi(\beta_1)) \perp (\varphi(\beta_3) - \varphi(\beta_4))$$

for all $\beta_3, \beta_4 \in B(X|E)$.

Here, \perp means “orthogonal to”. Hence, $\varphi(\beta_2)$ is the orthogonal projection of $\varphi(\beta_1)$ on $\varphi(B(X|E))$. Note that this orthogonal projection is always unique. We are now ready to present the main theorem of this paper.

THEOREM 3.3. (Representation Theorem). *Let X be a finite set of n states, and br a belief revision function on X . Then, the following two statements are equivalent:*

- (1) *The belief revision function br is linear, transitive and information-order independent;*
- (2) *There is a Euclidean space \mathbb{R}^m ($m \geq n - 1$) and a linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$, closed under orthogonal projections, such that for every initial belief β and every event E*

$$br(\beta|E) = \arg \min_{\beta' \in B(X|E)} \|\varphi(\beta') - \varphi(\beta)\|.$$

For the remainder of this paper, whenever we say that φ represents the belief revision function br , we mean that br is given by the formula above. In the proof of this theorem we will show that, if the belief revision functions br satisfies the three axioms above, then it can be represented by some $\varphi: B(X) \rightarrow \mathbb{R}^{n-1}$. That is, theoretically one can always choose $m = n - 1$. However, in practical cases it is often natural to identify states x with vectors $\varphi([x])$ of dimension larger than $n - 1$. For that reason, we allow in the theorem above for the possibility that $m > n - 1$, although theoretically we do not need it.

4. PROOF OF THE REPRESENTATION THEOREM

4.1. (2) *Implies (1)*

The easy part is to show the implication from (2) to (1). Assume that br is represented by a linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$ that is closed under orthogonal projections. We show that br is linear, transitive and information-order independent.

Linearity. Choose a pair of states $a, b \in X$. We first show that $br(\beta_1|a, b) = \beta_1$ whenever $\beta_1 \in B(X|\{a, b\})$. Let $\beta_1 \in B(X|\{a, b\})$ and $\beta_2 = br(\beta_1|a, b)$. Then,

$$\|\varphi(\beta_2) - \varphi(\beta_1)\| \leq \|\varphi(\beta_3) - \varphi(\beta_1)\|$$

for all $\beta_3 \in B(X|\{a, b\})$. However, as $\beta_1 \in B(X|\{a, b\})$ and φ is one-to-one, this is only possible when $\beta_2 = \beta_1$, hence $br(\beta_1|a, b) = \beta_1$.

Now, choose $\beta_1, \beta_2 \in B(X)$ and $\alpha \in [0, 1]$. Let $\beta := \alpha\beta_1 + (1 - \alpha)\beta_2$, and let $\beta_1^{\text{rev}} := br(\beta_1|a, b)$, $\beta_2^{\text{rev}} := br(\beta_2|a, b)$, $\beta^{\text{rev}} := br(\beta|a, b)$. We show that $\beta^{\text{rev}} = \alpha\beta_1^{\text{rev}} + (1 - \alpha)\beta_2^{\text{rev}}$. Since φ represents br ,

$$\|\varphi(\beta_1^{\text{rev}}) - \varphi(\beta_1)\| \leq \|\varphi(\beta_3) - \varphi(\beta_1)\|$$

for all $\beta_3 \in B(X|\{a, b\})$. Since φ is closed under orthogonal projections, this is equivalent to

$$\varphi(\beta_1^{\text{rev}}) - \varphi(\beta_1) \perp \varphi([a]) - \varphi([b]). \tag{1}$$

Similarly,

$$\varphi(\beta_2^{\text{rev}}) - \varphi(\beta_2) \perp \varphi([a]) - \varphi([b]). \tag{2}$$

Define $\tilde{\beta}^{\text{rev}} := \alpha\beta_1^{\text{rev}} + (1 - \alpha)\beta_2^{\text{rev}}$. Then, by (1), (2), and linearity of φ ,

$$\varphi(\tilde{\beta}^{\text{rev}}) - \varphi(\beta) \perp \varphi([a]) - \varphi([b]),$$

which implies that $\beta^{\text{rev}} = \tilde{\beta}^{\text{rev}}$. Hence, $\beta^{\text{rev}} = \alpha\beta_1^{\text{rev}} + (1 - \alpha)\beta_2^{\text{rev}}$. Together with the insight that $br(\beta_1|a, b) = \beta_1$ whenever $\beta_1 \in B(X|\{a, b\})$, it follows that br is linear.

Transitivity. Suppose that $br(\beta|a, b) = \frac{1}{2}[a] + \frac{1}{2}[b]$ and $br(\beta|b, c) = \frac{1}{2}[b] + \frac{1}{2}[c]$. Since φ is closed under orthogonal projections, this is equivalent to stating that

$$\begin{aligned} \varphi\left(\frac{1}{2}[a] + \frac{1}{2}[b]\right) - \varphi(\beta) &\perp \varphi([a]) - \varphi([b]), \\ \varphi\left(\frac{1}{2}[b] + \frac{1}{2}[c]\right) - \varphi(\beta) &\perp \varphi([b]) - \varphi([c]). \end{aligned}$$

This, in turn, implies that

$$\begin{aligned} \|\varphi(\beta) - \varphi(a)\| &= \|\varphi(\beta) - \varphi(b)\| \quad \text{and} \\ \|\varphi(\beta) - \varphi(b)\| &= \|\varphi(\beta) - \varphi(c)\|. \end{aligned}$$

As such,

$$\|\varphi(\beta) - \varphi(a)\| = \|\varphi(\beta) - \varphi(c)\|$$

yielding

$$\varphi\left(\frac{1}{2}[a] + \frac{1}{2}[c]\right) - \varphi(\beta) \perp \varphi([a]) - \varphi([c])$$

and hence, $br(\beta|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$. We may thus conclude that br is transitive.

Information-order independence. Let E_1, E_2 be two events with $E_2 \subseteq E_1$, and let β be some initial belief. Then, we have that

$$\varphi(br(\beta|E_1)) - \varphi(\beta) \perp \varphi([a]) - \varphi([b])$$

for all $a, b \in E_1$, and

$$\varphi(br(br(\beta|E_1)|E_2)) - \varphi(br(\beta|E_1)) \perp \varphi([a]) - \varphi([b])$$

for all $a, b \in E_2$. By combining these two facts, we may conclude that

$$\begin{aligned} \varphi(br(br(\beta|E_1)|E_2)) - \varphi(\beta) &= (\varphi(br(\beta|E_1)) - \varphi(\beta)) \\ &\quad + (\varphi(br(br(\beta|E_1)|E_2)) \\ &\quad - \varphi(br(\beta|E_1))) \\ &\quad \perp \varphi([a]) - \varphi([b]) \end{aligned}$$

for all $a, b \in E_2$. As such, $br(\beta|E_2) = br(br(\beta|E_1)|E_2)$, and hence we may conclude that br is information-order independent.

4.2. (1) Implies (2)

The difficult part is to show that the axioms linearity (LIN), transitivity (TRA) and information-order independence (IOI) imply that br is represented by some linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$ that is closed under orthogonal projections. For a given such function φ , let br_φ be the (unique) belief revision function that is represented by it. Take some belief revision function br that satisfies LIN, TRA and IOI.

We shall prove that there is some linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$, closed under orthogonal projections, such that $br = br_\varphi$.

The outline of this proof is as follows: Take a belief revision function br that satisfies the three axioms. For a given function φ , let $d_\varphi(a, b) := \|\varphi([a]) - \varphi([b])\|$ be the induced distance between states a and b . In Section 4.2.1, we present a system of equations for d_φ that is shown to be necessary and sufficient for $br = br_\varphi$. In Section 4.2.2 we construct a distance function d that satisfies this system of equations. In Section 4.2.3, finally, we prove that there is some linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^m$, closed under orthogonal projections, with $d_\varphi = d$. Hence, for this particular φ we would have that $br = br_\varphi$, which would complete the proof. In this section, we give for each lemma an intuitive argument that is easy to understand, and at the same time reveals the main idea behind the proof. The formal algebraic proofs are included in the Appendix. The philosophy is to first provide the reader with a picture of how the proof works, before confronting the reader with all the technical details.

4.2.1. Necessary and sufficient conditions for $br = br_\varphi$

In this part we derive a system of equations for d_φ that is both necessary and sufficient for $br = br_\varphi$. We proceed in two steps. As a first step, we show that br and br_φ coincide if and only if they coincide on every triangle $\{a, b, c\}$. Hence, a belief revision function satisfying the axioms LIN, TRA and IOI is completely determined by its behavior on triangles. In the second step, we provide a system of equations for d_φ that is necessary and sufficient for the event that br and br_φ coincide on every triangle. Combined with the first step, this system is also necessary and sufficient for $br = br_\varphi$.

In order to prove the first step we need the following lemma.

LEMMA 4.1. *If $br([a]|\{b, c\}) = [b]$, then $br([c]|\{a, b\}) = [b]$ and $br([b]|\{a, c\}) \notin \{[a], [c]\}$.*




Figure 6. A geometrical argument for Lemma 4.1.

Intuitive argument. Consider the second triangle of Figure 6 as an illustration. Suppose that $br([a]|\{b, c\}) = [b]$, and let

$$A := br([c]|\{a, b\}), \quad B := \frac{1}{2}[a] + \frac{1}{2}[c], \quad C := br(B|\{b, c\}) \text{ and} \\ D := br(B|\{a, b\}).$$

By LIN, line BC must be parallel to line ab , and hence $C = \frac{1}{2}[b] + \frac{1}{2}[c]$. But then, $br(B|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$ and $br(B|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$. By TRA, it must then hold that $D = \frac{1}{2}[a] + \frac{1}{2}[b]$. Since, by LIN, cA must be parallel to BD , it follows that $A = [b]$, as was to show.

Hence, we know that $br([a]|\{b, c\}) = [b]$ implies $br([c]|\{a, b\}) = [b]$. We now show that, under this assumption, $br([b]|\{a, c\}) \notin \{[a], [c]\}$. Suppose, on the contrary, that $br([b]|\{a, c\}) = [a]$. Consider the first triangle of Figure 6 as an illustration. Let

$$A := \frac{1}{2}[b] + \frac{1}{2}[c], \quad B := br(A|\{a, c\}) \text{ and } C := br(A|\{a, b\}).$$

By LIN, AB is parallel to ba , and hence $B = \frac{1}{2}[a] + \frac{1}{2}[c]$. Then, $br(A|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$ and $br(A|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$. By TRA, it follows that $C = \frac{1}{2}[a] + \frac{1}{2}[b]$. But then, AC is not parallel to cb , which contradicts LIN. Similarly, $br([b]|\{a, c\}) = [c]$ would lead to a contradiction as well. \square

We now prove that it is sufficient to check that br and br_φ agree on every triangle.




Figure 7. Belief revision on triangles is decisive.

LEMMA 4.2. (Belief revision on triangles is decisive). *If $br([c]|{a, b}) = br_\varphi([c]|{a, b})$ for all $a, b, c \in X$, then $br = br_\varphi$.*

Intuitive argument. We provide an argument for the case of four states. Consider the state space $X = \{a, b, c, d\}$. By LIN of br and br_φ , it follows immediately that $br(\beta|\{x, y\}) = br_\varphi(\beta|\{x, y\})$ for all $x, y \in X$. Consider the belief simplex in Figure 7 for the state space $X = \{a, b, c, d\}$. We show that the belief revision $br([d]|{a, b, c})$ is completely determined by the belief revisions $br([d]|{a, b})$, $br([d]|{b, c})$, $br([a]|{b, c})$ and $br([c]|{a, b})$. Let the beliefs A, \dots, F be as depicted in this figure. Hence, we must prove that $br([d]|{a, b, c})$ is completely determined by A, B, E and F . The line CA denotes all the beliefs in $B(X|\{a, b, c\})$ that, upon observing $\{a, b\}$, are mapped to A . Similarly, the line DB denotes all the beliefs in $B(X|\{a, b, c\})$ that, upon observing $\{b, c\}$, are mapped to B . By LIN, CA is parallel to cE , and DB is parallel to aF . By Lemma 4.1, we know that it cannot be the case that $E = [a]$ and $F = [c]$. As such, the lines cE and aF cannot coincide, implying that the lines CA and DB cannot be parallel. Hence, CA and DB have at most one intersection point. On the other hand, IOI of br




Figure 8. Necessary conditions for $br = br_\varphi$.

implies that $br(br([d]\{a, b, c\})\{a, b\}) = br([d]\{a, b\}) = A$ and, similarly, $br(br([d]\{a, b, c\})\{b, c\}) = br([d]\{b, c\}) = B$. Hence, $br([d]\{a, b, c\})$ should be on CA and DB . Since we have seen that CA and DB intersect at most once, $br([d]\{a, b, c\})$ is completely determined by A, B, C and D . By LIN, C and D are completely determined by E and F . Hence, $br([d]\{a, b, c\})$ is completely determined by A, B, E and F .

Since the same holds for br_φ , it follows that if br and br_φ coincide on the triangles $\{a, b, d\}, \{b, c, d\}$ and $\{a, b, c\}$, then $br([d]\{a, b, c\}) = br_\varphi([d]\{a, b, c\})$, and hence, by LIN of br and br_φ , $br(\beta\{a, b, c\}) = br_\varphi(\beta\{a, b, c\})$ for every β . By repeating this argument for $\{a, b, d\}, \{a, c, d\}$ and $\{b, c, d\}$, we may conclude the following: if br and br_φ coincide on every triangle, then $br = br_\varphi$, as was to show. \square

We are now ready to derive necessary and sufficient conditions for $br = br_\varphi$. Assume for the moment that $br = br_\varphi$. Take some states $a, b, c \in X$, let $\beta_1 := br([a]\{b, c\})$ and let $\beta_2 := br([c]\{a, b\})$. By definition of br_φ , it must then hold that

$$\varphi(\beta_1) - \varphi([a]) \perp \varphi([b]) - \varphi([c]) \quad \text{and} \quad \varphi(\beta_2) - \varphi([c]) \perp \varphi([a]) - \varphi([b]).$$

Figure 8 provides an illustration of these facts. Let the angle α be as in Figure 8. Assume first that α is less than 90° . Then,




Figure 9. Necessary conditions for $br = br_\varphi$.

$\beta_1(c) > 0$, and

$$\cos \alpha = \frac{\beta_1(c) d_\varphi(b, c)}{d_\varphi(a, b)} = \frac{\beta_2(a) d_\varphi(a, b)}{d_\varphi(b, c)}$$

which implies

$$\frac{d_\varphi(b, c)}{d_\varphi(a, b)} = \sqrt{\frac{\beta_2(a)}{\beta_1(c)}}. \tag{3}$$

If the angle α is exactly 90° , then $\beta_1(c) = \beta_2(a) = 0$, and (3) would not be well-defined. In that case, consider the angle γ (less than 90°) in Figure 9. Let $\beta_3 := br([b]|\{a, c\})$. Then, $\beta_3(c) > 0$ and

$$\cos \gamma = \frac{\beta_3(c) d_\varphi(a, c)}{d_\varphi(a, b)} = \frac{d_\varphi(a, b)}{d_\varphi(a, c)}$$

which implies that

$$\frac{d_\varphi(a, c)}{d_\varphi(a, b)} = \sqrt{\frac{1}{\beta_3(c)}}.$$

Since $d_\varphi(b, c) = \sqrt{d_\varphi(a, c)^2 - d_\varphi(a, b)^2}$, it follows that

$$\frac{d_\varphi(b, c)}{d_\varphi(a, b)} = \sqrt{\frac{1}{\beta_3(c)} - 1} = \sqrt{\frac{\beta_3(a)}{\beta_3(c)}}. \tag{4}$$

The necessary conditions (3) and (4) for $br = br_\varphi$ lead us to the following definition: For every three states a, b, c define

$$\lambda_{br}(a, b, c) := \begin{cases} \sqrt{\frac{br([c]|\{a, b\})(a)}{br([a]|\{b, c\})(c)}}, & \text{if } br([a]|\{b, c\})(c) > 0 \\ \sqrt{\frac{br([b]|\{a, c\})(a)}{br([b]|\{a, c\})(c)}}, & \text{if } br([a]|\{b, c\})(c) = 0 \end{cases} \tag{5}$$

By Lemma 4.1 we know that $br([a]|\{b, c\})(c) = 0$ implies $br([b]|\{a, c\})(c) > 0$, and hence (5) is well-defined. Note also that $\lambda_{br}(a, b, c) > 0$ for all a, b, c . Assume, namely, that $br([c]|\{a, b\})(a) = 0$. Then, by Lemma 4.1, $br([a]|\{b, c\})(c) = 0$ and $br([b]|\{a, c\})(a) > 0$, which means that $\lambda_{br}(a, b, c) > 0$. On the other hand, if $br([b]|\{a, c\})(a) = 0$, then, by the same lemma, $br([a]|\{b, c\})(c) = 1$ and $br([c]|\{a, b\})(a) > 0$, and hence $\lambda_{br}(a, b, c) > 0$.

By our insights above, we know that the system

$$\frac{d_\varphi(b, c)}{d_\varphi(a, b)} = \lambda_{br}(a, b, c) \text{ for all pairwise different } a, b, c \in X \tag{6}$$

provides a set of necessary conditions for $br = br_\varphi$. The following lemma states that this system is also sufficient for $br = br_\varphi$.

LEMMA 4.3. (Necessary and sufficient conditions for $br = br_\varphi$). *The belief revision functions br and br_φ coincide if and only if λ_{br} and d_φ satisfy system (6).*




Figure 10. System (6) is sufficient for $br = br_\varphi$.

Intuitive argument. From above, it should be clear that $br = br_\varphi$ only if λ_{br} and d_φ satisfy (6). Now, suppose that λ_{br} and d_φ satisfy (6). Choose some states $a, b, c \in X$. We prove that, for fixed φ , the system (6) completely determines the belief revision function br on $\{a, b, c\}$. Since, clearly, λ_{br_φ} and d_φ satisfy the system (6) as well, it would follow that br and br_φ agree on $\{a, b, c\}$. As this would hold for every a, b, c , Lemma 4.2 would imply that $br = br_\varphi$.

Consider the belief revisions

$$A := br([a] \{b, c\}), \quad B := br([b] \{a, c\}) \quad \text{and} \quad C := br([c] \{a, b\}),$$

as depicted in Figure 10. We show that A, B and C are completely determined by (6). Consider the triangle ABC and the angles α, β and γ as shown in the same figure. Let the function φ be fixed. Then, the equation

$$\frac{d_\varphi(b, c)}{d_\varphi(a, b)} = \lambda_{br}(a, b, c)$$

determines the ratio between $br([c] \{a, b\})(a)$ and $br([a] \{b, c\})(c)$. As such, it determines the ratio between the lengths of the line segments bC and bA , and thereby determines the angle γ . Similarly, the equations

$$\frac{d_\varphi(c, a)}{d_\varphi(b, c)} = \lambda_{br}(b, c, a) \text{ and } \frac{d_\varphi(a, b)}{d_\varphi(c, a)} = \lambda_{br}(c, a, b)$$

determine the angles α and β , respectively. However, there is only one triangle ABC with A on bc , B on ac and C on ab , inducing exactly these angles α, β and γ . Hence, A, B and C are completely determined by (6). As such, the behavior of br on $\{a, b, c\}$ is completely determined by (6), as was to show. \square

4.2.2. Existence of distance function d solving system (6)

In this part we prove that there is some symmetric⁴ distance function d , assigning to each pair a, b of states some positive number $d(a, b)$, that satisfies the system of equations (6). That is, we show that for the given belief revision function br there is some symmetric d with

$$\frac{d(b, c)}{d(a, b)} = \lambda_{br}(a, b, c)$$

for all pairwise different $a, b, c \in X$. We show this result in three steps. In step 1, we prove that the system (6) admits a solution d if and only if “every cycle of λ_{br} ’s has product 1”. Below, we explain what we exactly mean by this. In step 2, we show that it is sufficient to check that every cycle of three λ_{br} ’s has product 1. In step 3 we show that, indeed, every cycle of three λ_{br} ’s has product 1.

We first explain what we mean by a cycle of λ_{br} ’s. Consider two numbers $\lambda_{br}(a_1, b_1, c_1)$ and $\lambda_{br}(a_2, b_2, c_2)$, where (a_1, b_1, c_1) and (a_2, b_2, c_2) are ordered triples of states. We say that $\lambda_{br}(a_1, b_1, c_1)$ and $\lambda_{br}(a_2, b_2, c_2)$ are *adjacent* if $\{a_2, b_2\} = \{b_1, c_1\}$. Hence, $(a_2, b_2) = (b_1, c_1)$ or $(a_2, b_2) = (c_1, b_1)$. A *cycle of λ_{br} ’s* is a sequence

$$(\lambda_{br}(a_1, b_1, c_1), \lambda_{br}(a_2, b_2, c_2), \dots, \lambda_{br}(a_K, b_K, c_K))$$

of finite length such that $\lambda_{br}(a_k, b_k, c_k)$ and $\lambda_{br}(a_{k+1}, b_{k+1}, c_{k+1})$ are adjacent for all $k \in \{1, \dots, K - 1\}$, and $\lambda_{br}(a_K, b_K, c_K)$ is

adjacent to $\lambda_{br}(a_1, b_1, c_1)$. The *product* of this cycle is defined as $\lambda_{br}(a_1, b_1, c_1) \cdot \lambda_{br}(a_2, b_2, c_2) \cdots \lambda_{br}(a_K, b_K, c_K)$.

LEMMA 4.4. (Every cycle of λ_{br} 's must have product 1). *The system of equations (6) admits a solution d if and only if every cycle of λ_{br} 's has product 1.*

Intuitive argument. Obviously, (6) has a solution d only if every cycle of λ_{br} 's has product 1. Assume, now, that every cycle of λ_{br} 's has product 1. We show, for the case of four states, that there is a solution d to (6). Let $X = \{a, b, c, d\}$, and define the distance function d by

$$\begin{aligned} d(a, b) &:= 1, \\ d(a, c) &:= \lambda_{br}(b, a, c), \quad d(a, d) := \lambda_{br}(b, a, d), \\ d(b, c) &:= \lambda_{br}(a, b, c), \quad d(b, d) := \lambda_{br}(a, b, d), \\ d(c, d) &:= \lambda_{br}(a, b, c)\lambda_{br}(b, c, d). \end{aligned}$$

Then, it may easily be checked that d and λ_{br} satisfy (6). For instance,

$$\begin{aligned} \frac{d(d, c)}{d(a, d)} &= \frac{\lambda_{br}(a, b, c)\lambda_{br}(b, c, d)}{\lambda_{br}(b, a, d)} \\ &= \lambda_{br}(a, b, c)\lambda_{br}(b, c, d)\lambda_{br}(d, a, b) \\ &= \frac{1}{\lambda_{br}(c, d, a)}\lambda_{br}(a, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, a) \\ &\quad \lambda_{br}(d, a, b) \\ &= \frac{1}{\lambda_{br}(c, d, a)} = \lambda_{br}(a, d, c). \end{aligned}$$

Here, the second and the fifth equality follow from the fact that $1/\lambda_{br}(x, y, z) = \lambda_{br}(z, y, x)$. The fourth equality follows from the assumption that the product of the four λ_{br} 's is 1, since this sequence is a cycle of λ_{br} 's. Similarly, one can verify that the other equations in (6) are satisfied. \square

We next show that it is sufficient to check for the products of 3-cycles. A 3-cycle of λ_{br} 's is simply a cycle containing 3




Figure 11. Triangle-3-cycles and star-3-cycles.

λ_{br} 's. It is easily seen that there exist two types of 3-cycles:

$$(\lambda_{br}(a, b, c), \lambda_{br}(b, c, a), \lambda_{br}(c, a, b)) \text{ and } (\lambda_{br}(b, a, c), \lambda_{br}(c, a, d), \lambda_{br}(d, a, b)).$$

We refer to these two types as *triangle-3-cycles* and *star-3-cycles*, respectively. See Figure 11 for an illustration. From this picture, it also becomes clear why we have chosen these names.

LEMMA 4.5. (Checking for 3-cycles is sufficient). *Every cycle of λ_{br} 's has product 1 if and only if every 3-cycle of λ_{br} 's has product 1.*

Intuitive argument. We illustrate this lemma by means of the following example. Consider the state space $X = \{a, b, c, d\}$, and assume that every 3-cycle of λ_{br} 's has product 1. We show that the 4-cycle

$$(\lambda_{br}(a, b, c), \lambda_{br}(b, c, d), \lambda_{br}(c, d, a), \lambda_{br}(d, a, b))$$

has product 1. Since every star-3-cycle has product 1, we have that

$$\lambda_{br}(a, b, c)\lambda_{br}(c, b, d)\lambda_{br}(d, b, a) = 1 \text{ and} \\ \lambda_{br}(c, d, a)\lambda_{br}(a, d, b)\lambda_{br}(b, d, c) = 1,$$




Figure 12. Decomposition of cycle into 3-cycles.

or, equivalently,

$$\lambda_{br}(a, b, c) = \lambda_{br}(a, b, d)\lambda_{br}(d, b, c) \text{ and } \lambda_{br}(c, d, a) \\ = \lambda_{br}(c, d, b)\lambda_{br}(b, d, a).$$

We thus obtain

$$\lambda_{br}(a, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, a)\lambda_{br}(d, a, b) \\ = \lambda_{br}(a, b, d)\lambda_{br}(d, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, b)\lambda_{br}(b, d, a) \\ \times \lambda_{br}(d, a, b) \\ = [\lambda_{br}(a, b, d)\lambda_{br}(b, d, a)\lambda_{br}(d, a, b)] \\ \times [\lambda_{br}(d, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, b)] = 1,$$

where the latter sequence consists of two triangle-3-cycles for which the product, by assumption, is 1. Hence, we have decomposed the 4-cycle into two triangle-3-cycles, using the fact that every star-3-cycle has product 1. See Figure 12 for an illustration of this method. The formal proof is based on exactly this method. \square

We now prove that, indeed, every 3-cycle of λ_{br} 's has product 1. This would eventually imply that there exists a distance function d such that d and λ_{br} satisfy the system (6).

LEMMA 4.6. *Every triangle-3-cycle of λ_{br} 's has product 1.*

Intuitive argument. Consider a triangle $\{a, b, c\}$. We show that $\lambda_{br}(a, b, c)\lambda_{br}(b, c, a)\lambda_{br}(c, a, b) = 1$. Let $A := br([a]|\{b, c\})$, $B := br([b]|\{a, c\})$ and $C := br([c]|\{a, b\})$. By A_b we denote the probability that A assigns to state b . Similarly, we define




Figure 13. Every triangle-3-cycle of λ_{br} 's has product 1.

A_c, B_a, B_c, C_a and C_b . We focus on the case where A and C are in the interior of the line segments bc and ab , respectively, as illustrated in Figure 13. Then, by Lemma 4.1, also B must be in the interior of the line segment ac , and hence

$$\lambda_{br}(a, b, c) = \sqrt{\frac{C_a}{A_c}}, \quad \lambda_{br}(b, c, a) = \sqrt{\frac{A_b}{B_a}}, \quad \lambda_{br}(c, a, b) = \sqrt{\frac{B_c}{C_b}}.$$

Showing that $\lambda_{br}(a, b, c)\lambda_{br}(b, c, a)\lambda_{br}(c, a, b) = 1$ thus amounts to proving that

$$A_b B_c C_a = A_c B_a C_b.$$

By Ceva's theorem, this is equivalent to showing that the three lines aA, bB and cC are concurrent.⁵ More precisely, Ceva's theorem states that the three lines aA, bB and cC are concurrent if and only if

$$\|A - [c]\| \|B - [a]\| \|C - [b]\| = \|A - [b]\| \|B - [c]\| \|C - [a]\|. \quad (7)$$

However, since

$$A_b = \frac{\|A - [c]\|}{\|[b] - [c]\|}, \quad A_c = \frac{\|A - [b]\|}{\|[b] - [c]\|},$$

and similarly for B_c, B_a, C_a and C_b , (7) is equivalent to $A_b B_c C_a = A_c B_a C_b$.

We now show that aA, bB and cC are concurrent. Note that the lines aA and cC intersect at some point D in




Figure 14. Every star-3-cycle of λ_{br} 's has product 1.

the triangle (see Figure 13). That is, $br(D|\{b, c\}) = A$ and $br(D|\{a, b\}) = C$. Let $a' := \frac{1}{2}[b] + \frac{1}{2}[c]$, $b' := \frac{1}{2}[a] + \frac{1}{2}[c]$ and $c' := \frac{1}{2}[a] + \frac{1}{2}[b]$, and let D' be the “image” of D in the triangle $a'b'c'$. That is,

$$D' = D_a a' + D_b b' + D_c c'.$$

Similarly, let A' and C' be the images of A and C , as depicted in the second triangle of Figure 13. Then, the lines $A'a'$ and $C'c'$ intersect at the point D' . Since the lines $A'a'$ and $C'c'$ are parallel to the lines aA and cC , respectively, we may conclude that $br(A'|\{b, c\}) = a'$ and $br(C'|\{a, b\}) = c'$. Hence, $br(D'|\{b, c\}) = a'$ and $br(D'|\{a, b\}) = c'$. By TRA, it follows that $br(D'|\{a, c\}) = b'$. Let B' be the point on $a'c'$ such that $B'b'$ contains D' . Then, $br(B'|\{a, c\}) = b'$, which implies that the line $B'b'$ is parallel to the line bB , and hence B' must be the image of B . Since the lines $A'a'$, $B'b'$ and $C'c'$ are concurrent, and A', a', B', b', C', c' are the images of A, a, B, b, C, c , respectively, it follows that the lines aA, bB and cC must be concurrent, which was to prove. \square

LEMMA 4.7. Every star-3-cycle of λ_{br} 's has product 1.

Intuitive argument. Let $a, b, c, d \in X$. We show that $\lambda_{br}(a, d, b)\lambda_{br}(b, d, c)\lambda_{br}(c, d, a) = 1$. Assume that the beliefs A, B, \dots, P are as depicted in Figure 14. Here, $A := br([a]|\{b, c\})$, and similarly for B, C, \dots, I . We define $J := br([d]|\{a, b, c\})$. By K, L and M we denote the mapping of J on $\{a, b, d\}, \{b, c, d\}$ and $\{a, c, d\}$, respectively. Finally, N, O and P are the mappings of J on $\{a, d\}, \{b, d\}$ and $\{c, d\}$, respectively. By IOI of br , the mappings of K and M on $\{a, d\}$ are both equal to N . Similarly for O and P . By definition,

$$\lambda_{br}(a, d, b) = \sqrt{\frac{B_a}{A_b}}, \quad \lambda_{br}(b, d, c) = \sqrt{\frac{E_b}{D_c}}, \quad \lambda_{br}(c, d, a) = \sqrt{\frac{H_c}{G_a}},$$

where B_a is the probability that B assigns to a , and so on. Hence, showing that

$$\lambda_{br}(a, d, b)\lambda_{br}(b, d, c)\lambda_{br}(c, d, a) = 1$$

is equivalent to proving that

$$B_a E_b H_c = A_b D_c G_a.$$

Note that, by Lemma 4.6, the lines aA, bB and dC are concurrent. The same holds for the lines bD, cE and dF and for the lines cG, aH and dI . Note also that, by IOI,

$$br(K|\{a, b\}) = br(J|\{a, b\}) = br([d]|\{a, b\}) = C,$$

and hence K lies on the line dC . Similarly, L and M lie on the lines dF and dI , respectively. Since we have seen that $br(K|\{a, d\}) = N$, it follows by LIN that line KN is parallel to line bB . Similarly for the lines KO, LO, LP, MP and MN .

But then, it is easily seen from the figure that

$$\frac{B_a}{A_b} = \frac{N_a}{O_b}, \quad \frac{E_b}{D_c} = \frac{O_b}{P_c}, \quad \frac{H_c}{G_a} = \frac{P_c}{N_a},$$

which implies that

$$\frac{B_a}{A_b} \frac{E_b}{D_c} \frac{H_c}{G_a} = 1.$$

Hence, $B_a E_b H_c = A_b D_c G_a$, which was to show. \square

Combining Lemmas 4.4, 4.5, 4.6, and 4.7 thus leads to the following conclusion.

COROLLARY 4.8. *There is a symmetric distance function d such that d and λ_{br} satisfy system (6).*

4.2.3. *Existence of function φ with $br_\varphi = br$*

By Corollary 4.8 we know that there is some distance function d such that d and λ_{br} solve (6). We shall now explicitly construct a linear one-to-one mapping $\varphi: B(X) \rightarrow \mathbb{R}^{n-1}$, closed under orthogonal projections, for which the induced distance function d_φ coincides with d . As before, n denotes the number of states. Then, d_φ and λ_{br} would satisfy (6). By Lemma 4.3 it would then follow that $br_\varphi = br$, and hence br would be represented by φ . This would thus complete the proof of our Theorem 3.3.

We first show the reader how we construct φ for the case of three and four states, respectively, and provide for both cases a geometrical argument as to why the induced distance function coincides with d . We then provide a general formula for φ , and show that $d_\varphi = d$.

Case of three states. Let $X = \{a, b, c\}$, and let the distance function d be such that d and λ_{br} satisfy (6). We construct $\varphi: B(X) \rightarrow \mathbb{R}^2$ as follows. We choose

$$\varphi([a]) := (0, 0) \text{ and } \varphi([b]) := (d(a, b), 0).$$

Since φ must be linear, we define

$$\varphi(\beta) := \beta(a)\varphi([a]) + \beta(b)\varphi([b])$$

for all $\beta \in B(X|\{a, b\})$. Now, let $A := br([c]|\{a, b\})$ and $B := br(A|\{a, c\})$, as depicted in Figure 15. Since φ is to be constructed such that $br = br_\varphi$, the vector $\varphi(A) - \varphi([c])$ must be orthogonal to $\varphi([b]) - \varphi([a])$, and hence

$$\varphi([c]) = \varphi(A) + h(0, 1)$$

for some $h > 0$. The question is how to choose h . Let α be the angle as depicted in Figure 15. Since $\varphi(B) - \varphi(A)$ must be




Figure 15. Construction of φ for three states.

orthogonal to $\varphi([c]) - \varphi([a])$, we have

$$\cos \alpha = \frac{h}{\|\varphi([c]) - \varphi([a])\|} = \frac{B_a \|\varphi([c]) - \varphi([a])\|}{h}.$$

Since, moreover, we want to construct φ such that $d_\varphi = d$, we may substitute $\|\varphi([c]) - \varphi([a])\| = d(a, c)$, and obtain

$$h = \sqrt{B_a} d(a, c). \tag{8}$$

Hence,

$$\varphi([c]) := \varphi(A) + \sqrt{B_a} d(a, c)(0, 1).$$

By defining

$$\varphi(\beta) := \beta(a)\varphi([a]) + \beta(b)\varphi([b]) + \beta(c)\varphi([c])$$

for all $\beta \in B(X)$, the construction of φ is complete.

We shall now show that $d_\varphi = d$. Clearly, $d_\varphi(a, b) = d(a, b)$. By construction,

$$\begin{aligned} \|\varphi([c]) - \varphi([a])\|^2 &= \|\varphi([c]) - \varphi(A)\|^2 + \|\varphi(A) - \varphi([a])\|^2 \\ &= B_a d^2(a, c) + \|\varphi(A) - \varphi([a])\|^2 \\ &= d^2(a, c) - B_c d^2(a, c) + \|\varphi(A) - \varphi([a])\|^2. \end{aligned} \tag{9}$$

We show that $B_c d^2(a, c) = \|\varphi(A) - \varphi([a])\|^2$, which would imply that $\|\varphi([c]) - \varphi([a])\| = d(a, c)$. By LIN of br , we have that $B = A_a[a] + A_b C$ and hence $B_c = A_b C_c$. As such,

$$B_c d^2(a, c) = A_b C_c d^2(a, c).$$

Since d and λ_{br} satisfy (6), it follows that $d(a, c)/d(a, b) = \lambda_{br}(b, a, c)$. By definition,

$$\lambda_{br}(b, a, c) = \sqrt{\frac{A_b}{C_c}},$$

which implies that $C_c d^2(a, c) = A_b d^2(a, b)$. Consequently,

$$B_c d^2(a, c) = A_b C_c d^2(a, c) = A_b^2 d^2(a, b) = \|\varphi(A) - \varphi([a])\|^2.$$

Hence, by (9), $\|\varphi([c]) - \varphi([a])\| = d(a, c)$ and hence $d_\varphi(a, c) = d(a, c)$.

We now prove that $\|\varphi([c]) - \varphi([b])\| = d(b, c)$. From Figure 15, it is clear that $B_a/D_b = C_a/E_b$. Moreover, since d and λ_{br} satisfy (6), we have that $d(b, c)/d(a, c) = \lambda_{br}(a, c, b)$. By definition,

$$\lambda_{br}(a, c, b) = \sqrt{\frac{C_a}{E_b}} = \sqrt{\frac{B_a}{D_b}},$$

which yields $B_a/D_b = d^2(b, c)/d^2(a, c)$. Hence, $\sqrt{B_a}d(a, c) = \sqrt{D_b}d(b, c)$. By (8) we may conclude that

$$h = \sqrt{D_b}d(b, c).$$

By using a similar argument as we used above for $d_\varphi(a, c)$, we may then show that $d_\varphi(b, c) = d(b, c)$. Hence, $d_\varphi = d$, as was to show. \square

Case of four states. Let $X = \{a, b, c, d\}$, and let d be such that d and λ_{br} satisfy (6). We construct $\varphi: B(X) \rightarrow \mathbb{R}^3$ as follows. From the case of three states, we know how to construct $\varphi([a]), \varphi([b]), \varphi([c]) \in \mathbb{R}^3$, with the last coordinate of each of these vectors being zero, such that d_φ and d coincide on $B(X|\{a, b, c\})$. We show how to construct $\varphi([d])$. Let




Figure 16. Construction of φ for four states.

$A := br([d]|\{a, b, c\})$ and $B := br(A|\{a, d\})$, as depicted in Figure 16. As φ has to be constructed such that $br_\varphi = br$, it must be the case that $\varphi(A) - \varphi([d])$ is orthogonal to the triangle spanned by $\varphi([a]), \varphi([b])$ and $\varphi([c])$. Hence,

$$\varphi([d]) = \varphi(A) + h(0, 0, 1)$$

for some $h > 0$. It remains to determine h . Let α be the angle as depicted in Figure 16. Since $\varphi(B) - \varphi(A)$ is orthogonal to $\varphi([d]) - \varphi([a])$, we have

$$\cos \alpha = \frac{h}{\|\varphi([d]) - \varphi([a])\|} = \frac{B_a \|\varphi([d]) - \varphi([a])\|}{h},$$

hence we must set

$$h = \sqrt{B_a} \|\varphi([d]) - \varphi([a])\| = \sqrt{B_a} d(a, d), \tag{10}$$

as we wish to achieve that $\|\varphi([d]) - \varphi([a])\| = d(a, d)$. By defining

$$\varphi(\beta) = \beta(a)\varphi([a]) + \beta(b)\varphi([b]) + \beta(c)\varphi([c]) + \beta(d)\varphi([d])$$

for all $\beta \in B(X)$, the construction of φ is complete.

We now show that $d_\varphi = d$. By the case of three states, we already know that d_φ and d coincide on $\{a, b, c\}$. We first show




Figure 17. Construction of φ for four states.

that $d_\varphi(a, d) = d(a, d)$. Let the beliefs C, \dots, H be as depicted in Figure 17. By construction, and by (10),

$$\begin{aligned} \|\varphi([d]) - \varphi([a])\|^2 &= h^2 + \|\varphi(A) - \varphi([a])\|^2 = B_a d^2(a, d) \\ &\quad + \|\varphi(A) - \varphi([a])\|^2 \\ &= d^2(a, d) - B_d d^2(a, d) + \|\varphi(A) - \varphi([a])\|^2. \end{aligned} \tag{11}$$

We show that $B_d d^2(a, d) = \|\varphi(A) - \varphi([a])\|^2$. By LIN of br , we have that $B = A_a[a] + A_b D + A_c H$, and hence $B_d = A_b D_d + A_c H_d$. As such,

$$B_d d^2(a, d) = A_b D_d d^2(a, d) + A_c H_d d^2(a, d). \tag{12}$$

Since d and λ_{br} satisfy (6), we have that $d(a, d)/d(a, b) = \lambda_{br}(b, a, d)$. By definition,

$$\lambda_{br}(b, a, d) = \sqrt{\frac{C_b}{D_d}},$$

which implies that $D_d d^2(a, d) = C_b d^2(a, b)$. Similarly, $d(a, d)/d(a, c) = \lambda_{br}(c, a, d)$. By definition,

$$\lambda_{br}(c, a, d) = \sqrt{\frac{G_c}{H_d}},$$

implying that $H_d d^2(a, d) = G_c d^2(a, c)$. Substituting these identities in (12) yields

$$\begin{aligned} B_d d^2(a, d) &= A_b C_b d^2(a, b) + A_c G_c d^2(a, c) \\ &= A_b C_b \|\varphi([b]) - \varphi([a])\|^2 + A_c G_c \|\varphi([c]) - \varphi([a])\|^2. \end{aligned} \quad (13)$$

By IOI of br , we have that $C = br(A|\{a, b\})$. From the case of three states we know that br coincides with br_φ on $\{a, b, c\}$. Hence, $\varphi(C) - \varphi([a])$ is the orthogonal projection of $\varphi(A) - \varphi([a])$ on $\varphi([b]) - \varphi([a])$, which means that

$$\varphi(C) - \varphi([a]) = \frac{\langle \varphi(A) - \varphi([a]), \varphi([b]) - \varphi([a]) \rangle}{\|\varphi([b]) - \varphi([a])\|^2} (\varphi([b]) - \varphi([a])),$$

where $\langle \cdot \rangle$ denotes the dot product between vectors. Hence,

$$\|\varphi(C) - \varphi([a])\| = \frac{\langle \varphi(A) - \varphi([a]), \varphi([b]) - \varphi([a]) \rangle}{\|\varphi([b]) - \varphi([a])\|}. \quad (14)$$

Since

$$C_b \|\varphi([b]) - \varphi([a])\| = \|\varphi(C) - \varphi([a])\|,$$

it follows that

$$\begin{aligned} C_b \|\varphi([b]) - \varphi([a])\|^2 &= \|\varphi(C) - \varphi([a])\| \|\varphi([b]) - \varphi([a])\| \\ &= \langle \varphi(A) - \varphi([a]), \varphi([b]) - \varphi([a]) \rangle, \end{aligned}$$

where the last equality follows from (14). Similarly, we may conclude that

$$G_c \|\varphi([c]) - \varphi([a])\|^2 = \langle \varphi(A) - \varphi([a]), \varphi([c]) - \varphi([a]) \rangle.$$

Together with (13), this yields

$$\begin{aligned}
 B_d d^2(a, d) &= A_b \langle \varphi(A) - \varphi([a]), \varphi([b]) - \varphi([a]) \rangle \\
 &\quad + A_c \langle \varphi(A) - \varphi([a]), \varphi([c]) - \varphi([a]) \rangle \\
 &= A_a \langle \varphi(A) - \varphi([a]), \varphi([a]) - \varphi([a]) \rangle \\
 &\quad + A_b \langle \varphi(A) - \varphi([a]), \varphi([b]) - \varphi([a]) \rangle \\
 &\quad + A_c \langle \varphi(A) - \varphi([a]), \varphi([c]) - \varphi([a]) \rangle \\
 &= \langle \varphi(A) - \varphi([a]), (A_a \varphi([a]) + A_b \varphi([b]) \\
 &\quad + A_c \varphi([c])) - \varphi([a]) \rangle \\
 &= \langle \varphi(A) - \varphi([a]), \varphi(A) - \varphi([a]) \rangle = \|\varphi(A) - \varphi([a])\|^2.
 \end{aligned}$$

Substituting this result in (11) leads to the conclusion that

$$\|\varphi([d]) - \varphi([a])\|^2 = d^2(a, d),$$

which implies that $d_\varphi(a, d) = d(a, d)$. We now prove that $d_\varphi(b, d) = d(b, d)$. Let $I := br(A|\{b, d\})$ and $J := br([a]|\{b, d\})$, as depicted in Figure 17. From the discussion of Figure 14 earlier in this paper we know that $B_a/I_b = D_a/J_b$. Since $d(b, d)/d(a, d) = \lambda_{br}(a, d, b)$, and

$$\lambda_{br}(a, d, b) = \sqrt{\frac{D_a}{J_b}}$$

it follows that $D_a/J_b = d^2(b, d)/d^2(a, d)$, and hence $B_a/I_b = d^2(b, d)/d^2(a, d)$. As such, $\sqrt{B_a}d(a, d) = \sqrt{I_b}d(b, d)$. Since we have seen in (10) that $h = \sqrt{B_a}d(a, d)$, it follows that

$$h = \sqrt{I_b}d(b, d).$$

But then, by using the same argument as above, one can show that $d_\varphi(b, d) = d(b, d)$. The proof for $d_\varphi(c, d) = d(c, d)$ is similar. Hence, we may conclude that $d_\varphi = d$, as was to show. \square

The case of n states. The cases for three and four states already suggest how the general construction of the function φ for the case of n states should look like. Let $X = \{x_1, \dots, x_n\}$. We define a linear one-to-one function $\varphi : B(X) \rightarrow \mathbb{R}^{n-1}$ as

follows: Let $\mathbf{0}$ denote the zero vector in \mathbb{R}^{n-1} , and let $\{\mathbf{e}_1, \dots, \mathbf{e}_{n-1}\}$ denote the canonical basis for \mathbb{R}^{n-1} . Define

$$\varphi([x_1]) := \mathbf{0}, \quad \varphi([x_2]) := d(x_1, x_2)\mathbf{e}_1,$$

and

$$\varphi(\beta) := \beta(x_1)\varphi([x_1]) + \beta(x_2)\varphi([x_2])$$

for all $\beta \in B(X|\{x_1, x_2\})$. For every $k \in \{3, \dots, n\}$, let $X_{k-1} = \{x_1, \dots, x_{k-1}\}$, and recursively define

$$\begin{aligned} \varphi([x_k]) := & \varphi(br([x_k]|X_{k-1})) \\ & + \sqrt{br(br([x_k]|X_{k-1})|\{x_1, x_k\})(x_1)} \cdot d(x_1, x_k)\mathbf{e}_{k-1}, \end{aligned}$$

and

$$\varphi(\beta) := \sum_{i=1}^k \beta(x_i)\varphi([x_i])$$

for all $\beta \in B(X|\{x_1, \dots, x_k\})$. In the following lemma, we show that this function φ has all the desired properties.

LEMMA 4.9. (*Existence of φ with $d_\varphi = d$*). *Let φ be the function defined above. Then, φ is linear, one-to-one and closed under orthogonal projections, and $d_\varphi = d$.*

Intuitive argument. We first show, by induction on k , that d_φ and d coincide on $\{x_1, \dots, x_k\}$ for every $k \in \{3, \dots, n\}$. For $k = 3$, this follows from our argument above for three states. Suppose that $k \geq 4$ and that d_φ and d coincide on $\{x_1, \dots, x_{k-1}\}$. We show that $d_\varphi(x_1, x_k) = d(x_1, x_k)$. Define $A := br([x_k]|X_{k-1})$ and $B := br(A|\{x_1, x_k\})$. Then, we have that

$$\varphi([x_k]) = \varphi(A) + \sqrt{B_{x_1}}d(x_1, x_k)\mathbf{e}_{k-1}.$$

Clearly,

$$\begin{aligned} \|\varphi([x_k]) - \varphi([x_1])\|^2 &= \|\varphi([x_k]) - \varphi(A)\|^2 + \|\varphi(A) - \varphi([x_1])\|^2 \\ &= B_{x_1}d^2(x_1, x_k) + \|\varphi(A) - \varphi([x_1])\|^2 \\ &= d^2(x_1, x_k) - B_{x_k}d^2(x_1, x_k) + \|\varphi(A) \\ &\quad - \varphi([x_1])\|^2. \end{aligned}$$

By using the same techniques as in our proof for four states, one can show that

$$B_{x_k}d^2(x_1, x_k) = \|\varphi(A) - \varphi([x_1])\|^2,$$

and hence $\|\varphi([x_k]) - \varphi([x_1])\|^2 = d^2(x_1, x_k)$, implying that $d_\varphi(x_1, x_k) = d(x_1, x_k)$.

We now show that $d_\varphi(x_i, x_k) = d(x_i, x_k)$ for every $i \in \{2, \dots, k-1\}$. Let $C := br(A|\{x_i, x_k\})$. Then, using the same method as in our argument for four states, one can show that

$$\sqrt{B_{x_1}}d(x_1, x_k) = \sqrt{C_{x_i}}d(x_i, x_k)$$

and hence

$$\varphi([x_k]) = \varphi(A) + \sqrt{C_{x_i}}d(x_i, x_k)\mathbf{e}_{k-1}.$$

Using the same argument as above, one can then show that $d_\varphi(x_i, x_k) = d(x_i, x_k)$. By induction, it then follows that $d_\varphi = d$.

Since it is clear, by construction, that φ is linear and one-to-one, it only remains to show that φ is closed under orthogonal projections. It is sufficient to show that for every cornerpoint $\varphi([x_k])$ of the polytope $\varphi(B(X))$, and for every face $\varphi(B(X|E))$ with $E \subseteq X$, the orthogonal projection of $\varphi([x_k])$ on $\varphi(B(X|E))$ lies in $\varphi(B(X|E))$.

By construction of φ , the orthogonal projection of $\varphi([x_k])$ on $\varphi(B(X|\{x_1, \dots, x_{k-1}\}))$ is equal to $\varphi(br([x_k]|\{x_1, \dots, x_{k-1}\}))$, which lies in $\varphi(B(X|\{x_1, \dots, x_{k-1}\}))$.

Now, choose some arbitrary state $y_k \in X$, and some arbitrary event $E = \{y_1, \dots, y_{k-1}\} \subseteq X$ containing $k-1$ states, but not containing y_k . We show that the orthogonal projection of $\varphi([y_k])$ on $\varphi(B(X|E))$ is contained in $\varphi(B(X|E))$. Let $X \setminus (E \cup \{y_k\}) = \{y_{k+1}, \dots, y_n\}$. Suppose that we would apply the algorithm above, used to compute φ , not with respect to the order x_1, x_2, \dots, x_n , but with respect to the order y_1, y_2, \dots, y_n . This would yield some other function, say $\tilde{\varphi}$, and hence some other polytope, $\tilde{\varphi}(B(X))$. However, by using the same argument as above, one could then still prove that $\|\tilde{\varphi}([a]) - \tilde{\varphi}([b])\| = d(a, b)$ for all $a, b \in X$. That is, $\|\tilde{\varphi}([a]) - \tilde{\varphi}([b])\| = \|\varphi([a]) - \varphi([b])\|$ for all $a, b \in X$, and hence the two polytopes $\varphi(B(X))$ and

$\tilde{\varphi}(B(X))$ are isomorphic. By construction, the orthogonal projection of $\tilde{\varphi}([y_k])$ on $\tilde{\varphi}(B(X|\{y_1, \dots, y_{k-1}\}))$ is equal to $\tilde{\varphi}(br([y_k]|\{y_1, \dots, y_{k-1}\}))$, which lies in $\tilde{\varphi}(B(X|\{y_1, \dots, y_{k-1}\}))$. Hence, the orthogonal projection of $\tilde{\varphi}([y_k])$ on $\tilde{\varphi}(B(X|E))$ is contained in $\tilde{\varphi}(B(X|E))$. Since $\varphi(B(X))$ is isomorphic to $\tilde{\varphi}(B(X))$, it follows that also the orthogonal projection of $\varphi([y_k])$ on $\varphi(B(X|E))$ is contained in $\varphi(B(X|E))$. Since this holds for all $y_k \in X$ and all $E \subseteq X \setminus \{y_k\}$, we may conclude that φ is closed under orthogonal projections. This completes the argument for this lemma. \square

We are now fully equipped to prove the representation theorem. Let br be a belief revision function that satisfies the axioms LIN, TRA and IOI. Then, by Corollary 4.8, there is some distance function d such that d and λ_{br} satisfy the system (6) of equations. Moreover, by Lemma 4.9, there is some linear one-to-one function $\varphi: B(X) \rightarrow \mathbb{R}^{n-1}$, closed under orthogonal projections, with $d_\varphi = d$. Hence, d_φ and λ_{br} satisfy (6), which, by Lemma 4.3, implies that $br = br_\varphi$. Hence, br is represented by the function φ , which is linear, one-to-one and closed under orthogonal projections. This completes the proof of the representation theorem.

5. DISCUSSION

5.1. *Imaging versus Bayesian updating*

For some reason, the idea of imaging has not received much attention in decision theory and game theory. A possible explanation is that imaging would lead to a violation of *dynamic consistency*, a behavioral principle that plays a prominent role in both areas. In order to explain dynamic consistency, consider three possible states, a, b and c , and a decision maker who must choose between two acts, f and g . The final outcomes of f and g depend on the state. Suppose that there are three possible outcomes, x, y and z , and that the outcomes of the acts are given by the following table:

	<i>a</i>	<i>b</i>	<i>c</i>
<i>f</i>	<i>x</i>	<i>y</i>	<i>z</i>
<i>g</i>	<i>y</i>	<i>x</i>	<i>z</i>

Hence, *f* and *g* deliver the same outcome if the state would be *c*. Now, assume that the decision maker initially prefers *f* over *g*, indicating that he initially deems *a* or *b* possible. Dynamic consistency states that, if the decision maker would come to observe that the true state must be in $\{a, b\}$, he should still prefer *f* over *g*.

The connection between dynamic consistency and Bayesian updating has been studied in several decision theoretic models. Ghirardato (2002) studies a model in which the decision maker, for any possible event, holds a conditional preference relation over Savage-acts (Savage, 1954), and imposes axioms which guarantee that these conditional preferences can be represented by subjective expected utility functions. The paper shows that the conditional preferences are dynamically consistent if and only if the decision maker uses the same utility function for every observable event, and the induced belief revision function satisfies Bayesian updating. Epstein and Schneider (2003) prove that a similar result is true for conditional preferences over Anscombe-Aumann acts (Anscombe and Aumann, 1963), although studying a broader framework in which conditional preferences may, but need not, be of the subjective expected utility type.

Since the class of general imaging rules does not satisfy Bayesian updating, the two results above imply that embedding imaging rules into a dynamic decision problem with expected utility preferences necessarily leads to dynamically inconsistent preferences. Consequently, if a subjective expected utility maximizer in a dynamic decision problem uses a general imaging rule to revise his probabilistic beliefs, he should anticipate on the fact that his preferences over acts may change in the future due to new information about the state. A similar phenomenon is studied in Epstein (2005), who proposes an alternative model in which the decision maker's preferences over acts change over time due to non-Bayesian

belief revision. Epstein's model, which is based upon Gul and Pesendorfer (2001), considers a decision maker who has preferences over *menus* of acts, rather than over single acts alone. When choosing a menu of acts, the decision maker should anticipate on the fact that his preferences over acts within the menu may change in the future, when receiving new information about the state. Epstein then imposes axioms on the decision maker's preference relation over menus of acts which allow the decision maker to use a belief revision function that differs from Bayesian updating. More precisely, a preference relation that satisfies the axioms induces a belief revision function which, upon receiving new information about the state, generates a revised belief that can be written as a convex combination of the Bayesian update of the initial belief and some other revised belief which may significantly differ from this Bayesian update.

In the game-theoretic literature on belief revision, an important role is played by the relationship between belief revision and the *one-deviation property*. In order to explain the latter, consider a dynamic game in which a player must choose an action at two subsequent information sets, h_1 and h_2 . Suppose that the player chooses action a_1 at h_1 and action a_2 at h_2 . Assume, moreover, that a_1 is optimal at h_1 given his belief at h_1 about the opponents' actions and given his own future choice a_2 , and assume that a_2 is optimal at h_2 given his revised belief about the opponents' actions at h_2 . The one-deviation property states that in this case, the stream of actions (a_1, a_2) should be optimal given the player's belief at h_1 about the opponents' actions. That is, there should be no other stream of actions (a'_1, a'_2) that, given the player's belief at h_1 about the opponents' actions, would lead to a higher expected utility.

It may be verified that dynamic consistency of preferences implies the one-deviation property. Hendon et al. (1996) prove, within an equilibrium framework, that every pre-consistent belief revision function satisfies the one-deviation principle, whereas Perea (2002) shows that an appropriate weakening of pre-consistency, termed updating consistency, is

not only a sufficient condition, but also a necessary condition for the one-deviation principle. Both conditions, pre-consistency and updating consistency, are closely related to Bayesian updating. Since general imaging rules will typically not satisfy updating consistency when incorporated in a dynamic game, it follows that our model of belief revision is in conflict with the one-deviation property. The reason is that within our model, a player who must choose an action at the present information set h_1 and at some future information set h_2 , cannot evaluate the optimality of his actions at h_2 with his present belief at h_1 , since his conditional preferences at h_2 are in conflict with his initial preferences at h_1 . Rather, at h_1 he should correctly anticipate on the action he would choose at h_2 , given his future revised belief at h_2 , and subsequently should choose the action at h_1 that is optimal given his belief at h_1 about the opponents' actions, and given his own future choice at h_2 .

Summarizing, Bayesian updating is, roughly speaking, equivalent to dynamic consistency in decision problems and to the one-deviation property in dynamic games. Hence, if one insists on the latter two principles, one should accept Bayesian updating as the only reasonable candidate for belief revision. This may be the reason that decision theory and game theory have almost exclusively concentrated on Bayesian updating so far. However, it is not obvious that dynamic consistency should always be satisfied. The key idea in dynamic consistency is that the decision maker should be able to evaluate future choice problems, which would occur after receiving some new information, by beliefs that are held before receiving this new information. But why should this be the case? It seems natural to evaluate every separate choice problem by the revised belief you hold at the time you have to make the choice, and not by the beliefs you held before. It also seems natural to anticipate on future revised beliefs, and accordingly on future choices, when evaluating a choice problem today. Hence, a decision maker may analyze a dynamic choice problem in a perfectly sensible way while violating the dynamic consistency principle. I therefore believe that the

idea of imaging certainly deserves a place in the theory on dynamic choice problems and dynamic games.

In non-probabilistic belief revision theory, there is a similar discussion on the difference between the AGM axioms for belief revision (Alchourrón and Makinson, 1982; Gärdenfors, 1988) and the Katsuno–Mendelzon axioms for belief update (Katsuno and Mendelzon, 1992). The AGM model may be viewed as a non-probabilistic version of Bayesian updating, while the Katsuno–Mendelzon model represents a non-probabilistic version of imaging. For an intriguing comparison of both models, the reader is referred to Katsuno and Mendelzon (1992).

5.2. Independence of the axioms

It can easily be shown that the three axioms LIN, TRA and IOI are independent.

1. Let $X_1 = \{a, b, c\}$, and let br_1 be such that $br_1([a]|\{b, c\}) = [b]$, $br_1([b]|\{a, c\}) = br_1([c]|\{a, b\}) = [a]$, and $br_1(\beta|\{x, y\})$ is the revised belief obtained by Bayesian updating whenever β assigns positive probability to x or y , for all $x, y \in X_1$. Then, br_1 satisfies TRA and IOI, but not LIN.
2. Let $X_2 = \{a, b, c\}$, and let br_2 be such that $br_2([a]|\{b, c\}) = [b]$, $br_2([b]|\{a, c\}) = br_2([c]|\{a, b\}) = [a]$, and $br_2(\beta|\{x, y\}) = \beta(x)[x] + \beta(y)[y] + \beta(z)br_2([z]|\{x, y\})$ for all beliefs β , and all $x, y, z \in X_2$. Then, br_2 satisfies LIN and IOI, but not TRA, since $br_2(\frac{1}{2}[a] + \frac{1}{2}[c]|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$, $br_2(\frac{1}{2}[a] + \frac{1}{2}[c]|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$, whereas $br_2(\frac{1}{2}[a] + \frac{1}{2}[c]|\{a, b\}) = [a] \neq \frac{1}{2}[a] + \frac{1}{2}[b]$.
3. Let $X_3 = \{a, b, c, d\}$, and let br_3 be the linear belief revision function generated by $br_3([x]|\{y, z\}) = \frac{1}{2}[y] + \frac{1}{2}[z]$ for all $x, y, z \in X_3$, $br_3([d]|\{a, b, c\}) = [a]$ and $br_3([x]|X_3 \setminus \{x\}) = [d]$ for every $x \in X_3 \setminus \{d\}$. Then, br_3 satisfies LIN and TRA, but not IOI, since $br_3([d]|\{a, b\}) = \frac{1}{2}[a] + \frac{1}{2}[b]$, whereas

$$\begin{aligned} br_3(br_3([d]|\{a, b, c\})|\{a, b\}) &= br_3([a]|\{a, b\}) \\ &= [a] \neq br_3([d]|\{a, b\}). \end{aligned}$$

5.3. Uniqueness of the function φ

By the representation theorem we know that for every belief revision function br satisfying LIN, TRA and IOI there is some linear one-to-one function φ , closed under orthogonal projections, such that $br = br_\varphi$. Moreover, the proof of the representation theorem provides an algorithm to compute one such φ . Namely, the proof of Lemma 4.4 shows how for a given br we may compute a distance function d such that d and λ_{br} satisfy (6). In Section 4.2.3 we show how for such a distance function d we may subsequently compute a linear one-to-one function φ , closed under orthogonal projections, with $d_\varphi = d$. Hence, $br = br_\varphi$ for this particular φ .

We now focus on the question to what extent this function φ is unique. That is, how much freedom do we have when choosing a function φ such that $br = br_\varphi$? Assume that br satisfies the three axioms, and that $\varphi, \tilde{\varphi}$ are two different such functions with $br = br_\varphi = br_{\tilde{\varphi}}$. Then, by Lemma 4.3, it follows that d_φ and λ_{br} satisfy (6), but also $d_{\tilde{\varphi}}$ and λ_{br} satisfy (6). By the structure of the system (6), this implies that there is some scalar $\alpha > 0$ such that $d_{\tilde{\varphi}}(a, b) = \alpha d_\varphi(a, b)$ for all $a, b \in X$. Hence, $\|\tilde{\varphi}([a]) - \tilde{\varphi}([b])\| = \alpha \|\varphi([a]) - \varphi([b])\|$ for all $a, b \in X$, which means that the polytopes $\varphi(B(X))$ and $\tilde{\varphi}(B(X))$ are isomorphic. We thus arrive at the following conclusion:

LEMMA 5.1. (Uniqueness of φ). *Let br be a belief revision function satisfying LIN, TRA, and IOI. Let $\varphi, \tilde{\varphi}$ be two linear one-to-one functions, closed under orthogonal projections, such that $br = br_\varphi = br_{\tilde{\varphi}}$. Then, there is some $\alpha > 0$ such that $\|\tilde{\varphi}([a]) - \tilde{\varphi}([b])\| = \alpha \|\varphi([a]) - \varphi([b])\|$ for all $a, b \in X$.*

Obviously, the other direction is also true: If $br = br_\varphi$ and $\tilde{\varphi}$ is such that $\|\tilde{\varphi}([a]) - \tilde{\varphi}([b])\| = \alpha \|\varphi([a]) - \varphi([b])\|$ for all $a, b \in X$, then $br = br_{\tilde{\varphi}}$. In other words, if we have found one φ with $br = br_\varphi$, then we know how to generate all other $\tilde{\varphi}$ with $br = br_{\tilde{\varphi}}$.

5.4. Representing states by Boolean vectors

From the main theorem we know that for a given belief revision function br satisfying the three axioms we may construct a function φ that represents br . In many practical examples it seems plausible that the decision maker reasons in the other direction: He first chooses a function φ , identifying each state with some appropriately chosen vector of characteristics, and then revises his beliefs according to br_φ . The question remains how to choose φ in an adequate manner. A possible way to do this would be as follows: One first makes a list of properties that states may have or not have, and that seem relevant for the decision problem at hand. Subsequently, one defines for each state x a Boolean vector $\varphi([x])$ of zeros and ones which specifies for each property whether it is “true” (1) or “false” (0) at state x . For every belief β , the vector $\varphi(\beta)$ would then be a vector of true-false-probabilities, specifying for each property the probability that it is deemed true in the belief β . For every two states, the induced distance $\|\varphi([a]) - \varphi([b])\|$ would then be equal to $\sqrt{dis(a, b)}$, where $dis(a, b)$ denotes the number of properties on which a and b disagree. This distance is also called the *Hamming distance*. By the proof of Lemma 4.3, the belief revision function br represented by φ would be such that

$$br([c]|\{a, b\})(a) = \frac{dis(a, b) + dis(b, c) - dis(a, c)}{2dis(a, b)}$$

for all states $a, b, c \in X$. (We must make sure that φ is closed under orthogonal projections such that this number is between 0 and 1). That is, upon observing that the real state is in $\{a, b\}$ and state c is impossible, more weight is shifted towards the state in $\{a, b\}$ that has minimal Hamming distance to c . More generally, the belief revision function br would select for every initial belief β_1 and every event E the revised belief $\beta_2 \in B(X|E)$ for which the vector of true-false-probabilities is as close as possible to the vector of true-false-probabilities for β_1 .

APPENDIX: ALGEBRAIC PROOFS

PROOF OF LEMMA 4.1 The proof of this lemma is basically a direct translation of the geometrical argument, as provided in Section 4, into formal algebraic statements. It is therefore omitted here. \square

For the proof of Lemma 4.2, we need the following property.

LEMMA 6.1. *Let br satisfy LIN, TRA, and IOI. Then, for every belief $\beta \in B(X)$ and every event $E \subseteq X$ it holds that*

$$br(\beta|E) = \sum_{x \in E} \beta(x)[x] + \sum_{x \in X \setminus E} \beta(x)br([x]|E). \quad (15)$$

This lemma thus states that a belief revision function satisfying our axioms also satisfies Gärdenfors' homomorphism condition.

PROOF OF LEMMA 6.1 We prove (15) by induction on $|E|$, where $|E|$ denotes the cardinality of E .

If $|E|=2$, (15) follows directly from LIN of br .

Now, take some E with $|E|=k > 2$, and suppose that (15) holds for all $\beta' \in B(X)$ and all E' with $|E'| < k$. Choose some $a, b \in E$, some belief $\beta \in B(X)$, and define

$$\begin{aligned} \beta_E &:= br(\beta|E), \beta_{E \setminus a} := br(\beta_E|E \setminus \{a\}) \text{ and} \\ \beta_{E \setminus b} &:= br(\beta_E|E \setminus \{b\}). \end{aligned}$$

By the induction assumption,

$$\begin{aligned} \beta_{E \setminus a} &= \beta_E(a)br([a]|E \setminus \{a\}) + \sum_{x \in E \setminus \{a\}} \beta_E(x)[x], \\ \beta_{E \setminus b} &= \beta_E(b)br([b]|E \setminus \{b\}) + \sum_{x \in E \setminus \{b\}} \beta_E(x)[x]. \end{aligned}$$

Since, by IOI, $\beta_{E \setminus a} = br(\beta|E \setminus \{a\})$ and $\beta_{E \setminus b} = br(\beta|E \setminus \{b\})$, we have

$$\begin{aligned} br(\beta|E \setminus \{a\}) &= \beta_E(a)br([a]|E \setminus \{a\}) + \sum_{x \in E \setminus \{a\}} \beta_E(x)[x], \\ br(\beta|E \setminus \{b\}) &= \beta_E(b)br([b]|E \setminus \{b\}) + \sum_{x \in E \setminus \{b\}} \beta_E(x)[x]. \end{aligned}$$

In particular,

$$\begin{aligned} br(\beta|E \setminus \{a\})(b) &= \beta_E(a)br([a]|E \setminus \{a\})(b) + \beta_E(b), \\ br(\beta|E \setminus \{b\})(a) &= \beta_E(b)br([b]|E \setminus \{b\})(a) + \beta_E(a), \end{aligned}$$

or, equivalently,

$$\begin{aligned} &\begin{bmatrix} br([a]|E \setminus \{a\})(b) & 1 \\ 1 & br([b]|E \setminus \{b\})(a) \end{bmatrix} \begin{bmatrix} \beta_E(a) \\ \beta_E(b) \end{bmatrix} \\ &= \begin{bmatrix} br(\beta|E \setminus \{a\})(b) \\ br(\beta|E \setminus \{b\})(a) \end{bmatrix}. \end{aligned} \quad (16)$$

The determinant of the matrix above is

$$br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a) - 1.$$

We show that this determinant is not zero. Assume, on the contrary, that the determinant would be zero. Then, $br([a]|E \setminus \{a\})(b) = 1$ and $br([b]|E \setminus \{b\})(a) = 1$, which means that $br([a]|E \setminus \{a\}) = [b]$ and $br([b]|E \setminus \{b\}) = [a]$. Choose some $c \in E \setminus \{a, b\}$. By IOI,

$$\begin{aligned} br([a]|\{b, c\}) &= br(br([a]|E \setminus \{a\})|\{b, c\}) = br([b]|\{b, c\}) = [b], \\ br([b]|\{a, c\}) &= br(br([b]|E \setminus \{b\})|\{a, c\}) = br([a]|\{a, c\}) = [a], \end{aligned}$$

which would contradict Lemma 4.1. Hence, the determinant is not zero. As such, the system (16) has a unique solution $(\beta_E(a), \beta_E(b))$ with, in particular,

$$\begin{aligned} br(\beta|E)(a) &= \beta_E(a) \\ &= \frac{br(\beta|E \setminus \{b\})(a) - br([b]|E \setminus \{b\})(a) \cdot br(\beta|E \setminus \{a\})(b)}{1 - br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a)}. \end{aligned} \quad (17)$$

By the induction assumption, we know that

$$\begin{aligned}
 br(\beta|E\setminus\{a\}) &= \sum_{x \in E\setminus\{a\}} \beta(x)[x] + \beta(a)br([a]|E\setminus\{a\}) \\
 &\quad + \sum_{x \in X\setminus E} \beta(x)br([x]|E\setminus\{a\}), \\
 br(\beta|E\setminus\{b\}) &= \sum_{x \in E\setminus\{b\}} \beta(x)[x] + \beta(b)br([b]|E\setminus\{b\}) \\
 &\quad + \sum_{x \in X\setminus E} \beta(x)br([x]|E\setminus\{b\}),
 \end{aligned}$$

which implies that

$$\begin{aligned}
 br(\beta|E\setminus\{a\})(b) &= \beta(b) + \beta(a)br([a]|E\setminus\{a\})(b) \\
 &\quad + \sum_{x \in X\setminus E} \beta(x)br([x]|E\setminus\{a\})(b), \\
 br(\beta|E\setminus\{b\})(a) &= \beta(a) + \beta(b)br([b]|E\setminus\{b\})(a) \\
 &\quad + \sum_{x \in X\setminus E} \beta(x)br([x]|E\setminus\{b\})(a).
 \end{aligned}$$

By substituting these two equations in (17), we eventually obtain

$$\begin{aligned}
 br(\beta|E)(a) &= \beta(a) + \sum_{x \in X\setminus E} \beta(x) \\
 &\quad \times \frac{br([x]|E\setminus\{b\})(a) - br([b]|E\setminus\{b\})(a) \cdot br([x]|E\setminus\{a\})(b)}{1 - br([a]|E\setminus\{a\})(b) \cdot br([b]|E\setminus\{b\})(a)} \\
 &= \beta(a) + \sum_{x \in X\setminus E} \beta(x)br([x]|E)(a), \tag{18}
 \end{aligned}$$

where the latter equality follows from applying (17) to the initial belief $\beta := [x]$. Since (18) holds for every $a \in E$, we may conclude that

$$br(\beta|E) = \sum_{x \in E} \beta(x)[x] + \sum_{x \in X\setminus E} \beta(x)br([x]|E),$$

which was to show. \square

PROOF OF LEMMA 4.2 Assume that br and br_φ coincide on every triangle. We prove that $br = br_\varphi$. To that purpose, we show that

$$br(\beta|E) = br_\varphi(\beta|E) \quad (19)$$

for every belief $\beta \in B(X)$ and every event $E \subseteq X$. We show (19) by induction on $|E|$.

If $|E|=2$, say $E=\{a, b\}$, (19) follows by the assumption that $br([x]|\{a, b\}) = br_\varphi([x]|\{a, b\})$ for every $x \in X \setminus \{a, b\}$, and the fact that br and br_φ satisfy LIN.

Now, let $|E|=k > 2$, and assume that (19) holds for every belief β' and every event E' with $|E'| < k$. Since br and br_φ satisfy LIN, TRA and IOI, we know, by the proof of Lemma 6.1, that both br and br_φ satisfy (17). That is,

$$br(\beta|E)(a) = \frac{br(\beta|E \setminus \{b\})(a) - br([b]|E \setminus \{b\})(a) \cdot br(\beta|E \setminus \{a\})(b)}{1 - br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a)},$$

$$br_\varphi(\beta|E)(a) = \frac{br_\varphi(\beta|E \setminus \{b\})(a) - br_\varphi([b]|E \setminus \{b\})(a) \cdot br_\varphi(\beta|E \setminus \{a\})(b)}{1 - br_\varphi([a]|E \setminus \{a\})(b) \cdot br_\varphi([b]|E \setminus \{b\})(a)}$$

for every $a, b \in E$. Since, by the induction assumption, $br(\beta|E \setminus \{a\}) = br_\varphi(\beta|E \setminus \{a\})$, $br(\beta|E \setminus \{b\}) = br_\varphi(\beta|E \setminus \{b\})$, $br([a]|E \setminus \{a\}) = br_\varphi([a]|E \setminus \{a\})$ and $br([b]|E \setminus \{b\}) = br_\varphi([b]|E \setminus \{b\})$ for every $a, b \in E$, it follows that $br(\beta|E)(a) = br_\varphi(\beta|E)(a)$ for every $a \in E$, and hence $br(\beta|E) = br_\varphi(\beta|E)$. This completes the proof. \square

PROOF OF LEMMA 4.3 We have already seen that br and d_φ must satisfy (6) if $br = br_\varphi$. Now, assume that λ_{br} and d_φ satisfy system (6). We show that this implies $br = br_\varphi$. In view of Lemma 4.2 it is sufficient to show that $br([c]|\{a, b\}) = br_\varphi([c]|\{a, b\})$ for every triple a, b, c of pairwise different states. Choose some arbitrary triple a, b, c of states, and define

$$\beta_1 := br([a]|\{b, c\}), \beta_2 := br([c]|\{a, b\}) \text{ and } \beta_3 := br([b]|\{a, c\}).$$

By Lemma 4.1, there are four possible cases to distinguish.

Case 1. If $\beta_1 \notin \{[b], [c]\}$, $\beta_2 \notin \{[a], [b]\}$ and $\beta_3 \notin \{[a], [c]\}$. Then,

$$\lambda_{br}^2(a, b, c) = \frac{\beta_2(a)}{\beta_1(c)}, \quad \lambda_{br}^2(b, c, a) = \frac{\beta_1(b)}{\beta_3(a)} \text{ and}$$

$$\lambda_{br}^2(c, a, b) = \frac{\beta_3(c)}{\beta_2(b)}.$$

Since λ_{br} and d_φ satisfy (6), it follows that

$$\beta_1(c) d_\varphi^2(b, c) = \beta_2(a) d_\varphi^2(a, b),$$

$$\beta_3(a) d_\varphi^2(c, a) = \beta_1(b) d_\varphi^2(b, c),$$

$$\beta_2(b) d_\varphi^2(a, b) = \beta_3(c) d_\varphi^2(c, a).$$

It can easily be shown that this system, for fixed d_φ , has a unique solution $\beta_1, \beta_2, \beta_3$. In particular, this system implies that

$$br([c]|\{a, b\})(a) = \beta_2(a) = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)}.$$

Case 2. If $\beta_2 = [a]$, $\beta_3 = [a]$, and $\beta_1 \notin \{[b], [c]\}$. Then,

$$\lambda_{br}^2(a, b, c) = \frac{1}{\beta_1(c)}, \quad \lambda_{br}^2(b, c, a) = \beta_1(b) \text{ and } \lambda_{br}^2(c, a, b)$$

$$= \frac{\beta_1(c)}{\beta_1(b)}.$$

Since λ_{br} and d_φ satisfy (6), it follows that

$$\beta_1(c) d_\varphi^2(b, c) = d_\varphi^2(a, b),$$

$$d_\varphi^2(c, a) = \beta_1(b) d_\varphi^2(b, c),$$

$$\beta_1(b) d_\varphi^2(a, b) = \beta_1(c) d_\varphi^2(c, a).$$

Consequently,

$$d_\varphi^2(a, b) + d_\varphi^2(c, a) = \beta_1(c) d_\varphi^2(b, c) + \beta_1(b) d_\varphi^2(b, c) = d_\varphi^2(b, c).$$

Hence,

$$br([c]|\{a, b\})(a) = \beta_2(a) = 1 = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)}.$$

Case 3. If $\beta_2 = [b]$, $\beta_1 = [b]$ and $\beta_3 \notin \{[a], [c]\}$. By the same method as in Case 2, one can show that

$$d_\varphi^2(a, b) + d_\varphi^2(b, c) = d_\varphi^2(c, a)$$

and hence

$$br([c]|\{a, b\})(a) = \beta_2(a) = 0 = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)}.$$

Case 4. If $\beta_1 = [c]$, $\beta_3 = [c]$ and $\beta_2 \notin \{[a], [b]\}$. Then,

$$\begin{aligned} \lambda_{br}^2(a, b, c) &= \beta_2(a), \quad \lambda_{br}^2(b, c, a) = \frac{\beta_2(b)}{\beta_2(a)} \text{ and } \lambda_{br}^2(c, a, b) \\ &= \frac{1}{\beta_2(b)}. \end{aligned}$$

Since λ and d_φ satisfy (6), we have that

$$\begin{aligned} d_\varphi^2(b, c) &= \beta_2(a) d_\varphi^2(a, b), \\ \beta_2(a) d_\varphi^2(c, a) &= \beta_2(b) d_\varphi^2(b, c), \\ \beta_2(b) d_\varphi^2(a, b) &= d_\varphi^2(c, a). \end{aligned}$$

Hence,

$$\beta_2(a) = \frac{d_\varphi^2(b, c)}{d_\varphi^2(a, b)}$$

and

$$d_\varphi^2(b, c) + d_\varphi^2(c, a) = \beta_2(a) d_\varphi^2(a, b) + \beta_2(b) d_\varphi^2(a, b) = d_\varphi^2(a, b).$$

As such,

$$\begin{aligned} br([c]|\{a, b\})(a) &= \beta_2(a) = \frac{d_\varphi^2(b, c)}{d_\varphi^2(a, b)} \\ &= \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)}. \end{aligned}$$

Since these are all possible cases, we may conclude that, in general,

$$br([c]|\{a, b\})(a) = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)} \quad (20)$$

for all $a, b, c \in X$.

On the other hand, also λ_{br_φ} and d_φ satisfy (6), since this system provides necessary conditions for $br = br_\varphi$. By the same method as above, one can then show that

$$br_\varphi([c]|\{a, b\})(a) = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)}$$

for all $a, b, c \in X$. Hence, $br([c]|\{a, b\}) = br_\varphi([c]|\{a, b\})$ for all $a, b, c \in X$. But then, Lemma 4.2 guarantees that $br = br_\varphi$. \square

PROOF OF LEMMA 4.4 Obviously, if (6) has a solution d , then every cycle of λ_{br} 's must have product 1. Assume now that every cycle of λ_{br} 's has product 1, and let $X = \{x_1, x_2, \dots, x_n\}$. We define the distance function d by

$$\begin{aligned} d(x_1, x_2) &:= 1, \\ d(x_1, x_i) &:= \lambda_{br}(x_2, x_1, x_i) \text{ for all } i \geq 3, \\ d(x_2, x_i) &:= \lambda_{br}(x_1, x_2, x_i) \text{ for all } i \geq 3, \\ d(x_i, x_j) &:= \lambda_{br}(x_1, x_2, x_i)\lambda_{br}(x_2, x_i, x_j) \text{ for all } 3 \leq i < j. \end{aligned}$$

Finally, let $d(x_i, x_j) := d(x_j, x_i)$ whenever $i > j$. We prove that d satisfies (6). To that purpose, we show that

$$\frac{d(x_j, x_k)}{d(x_i, x_j)} = \lambda_{br}(x_i, x_j, x_k)$$

for all triples (x_i, x_j, x_k) . We must distinguish various cases, depending on whether some of the states in $\{x_i, x_j, x_k\}$ is x_1 or x_2 , whether $j < k$ or not, and whether $i < j$ or not. For the sake of brevity, we shall only deal with one case here, since the proofs for all other cases are similar. Consider the case

where $x_i, x_j, x_k \notin \{x_1, x_2\}$, and where $i < j < k$. Then,

$$\begin{aligned} \frac{d(x_j, x_k)}{d(x_i, x_j)} &= \frac{\lambda_{br}(x_1, x_2, x_j)\lambda_{br}(x_2, x_j, x_k)}{\lambda_{br}(x_1, x_2, x_i)\lambda_{br}(x_2, x_i, x_j)} \\ &= \lambda_{br}(x_1, x_2, x_j)\lambda_{br}(x_2, x_j, x_k)\lambda_{br}(x_i, x_2, x_1) \\ &\quad \times \lambda_{br}(x_j, x_i, x_2) \\ &= \frac{1}{\lambda_{br}(x_k, x_j, x_i)}\lambda_{br}(x_1, x_2, x_j)\lambda_{br}(x_2, x_j, x_k) \\ &\quad \times \lambda_{br}(x_k, x_j, x_i) \\ &\quad \times \lambda_{br}(x_j, x_i, x_2)\lambda_{br}(x_i, x_2, x_1) \\ &= \frac{1}{\lambda_{br}(x_k, x_j, x_i)} = \lambda_{br}(x_i, x_j, x_k). \end{aligned}$$

Here, the fourth equality follows from the assumption that the product of the 5 λ_{br} 's is 1, since this sequence is a cycle of λ_{br} 's. The other cases can be shown in a similar fashion. \square

PROOF OF LEMMA 4.5 Assume that every 3-cycle of λ_{br} 's has product 1. We show that every K -cycle of λ_{br} 's has product 1. We proceed by induction on K .

For $K = 3$, the statement follows trivially. Assume therefore that $K \geq 4$, and that the statement holds for all cycles of length less than K . Consider a K -cycle

$$C := (\lambda_{br}(a_1, b_1, c_1), \lambda_{br}(a_2, b_2, c_2), \dots, \lambda_{br}(a_K, b_K, c_K)).$$

We distinguish three possible cases.

Case 1. Suppose that $(a_2, b_2) = (b_1, c_1)$ and $(b_K, c_K) = (a_1, b_1)$. Then, since star-3-cycles have product 1,

$$\begin{aligned} \lambda_{br}(a_2, b_2, c_2) &= \lambda_{br}(b_1, c_1, c_2) = \lambda_{br}(b_1, c_1, a_1)\lambda_{br}(a_1, c_1, c_2), \\ \lambda_{br}(a_K, b_K, c_K) &= \lambda_{br}(a_K, a_1, b_1) = \lambda_{br}(a_K, a_1, c_1)\lambda_{br}(c_1, a_1, b_1). \end{aligned}$$

Consequently, the product of C is equal to

$$\begin{aligned} &[\lambda_{br}(a_K, b_K, c_K)\lambda_{br}(a_1, b_1, c_1)\lambda_{br}(a_2, b_2, c_2)] \\ &\quad \cdot [\lambda_{br}(a_3, b_3, c_3) \cdots \lambda_{br}(a_{K-1}, b_{K-1}, c_{K-1})] \\ &= [\lambda_{br}(a_K, a_1, c_1)\lambda_{br}(c_1, a_1, b_1)\lambda_{br}(a_1, b_1, c_1)\lambda_{br}(b_1, c_1, a_1) \\ &\quad \times \lambda_{br}(a_1, c_1, c_2)] \cdot \\ &\quad \cdot [\lambda_{br}(a_3, b_3, c_3) \cdots \lambda_{br}(a_{K-1}, b_{K-1}, c_{K-1})] \end{aligned}$$

$$\begin{aligned}
 &= [\lambda_{br}(c_1, a_1, b_1)\lambda_{br}(a_1, b_1, c_1)\lambda_{br}(b_1, c_1, a_1)] \cdot \\
 &\quad \cdot [\lambda_{br}(a_K, a_1, c_1)\lambda_{br}(a_1, c_1, c_2)\lambda_{br}(a_3, b_3, c_3) \cdots \cdots \\
 &\quad \times \lambda_{br}(a_{K-1}, b_{K-1}, c_{K-1})] \\
 &= [\lambda_{br}(c_1, a_1, b_1)\lambda_{br}(a_1, b_1, c_1)\lambda_{br}(b_1, c_1, a_1)] \cdot \\
 &\quad \cdot [\lambda_{br}(a_K, a_1, b_2)\lambda_{br}(a_1, b_2, c_2)\lambda_{br} \\
 &\quad (a_3, b_3, c_3) \cdots \cdots \lambda_{br}(a_{K-1}, b_{K-1}, c_{K-1})] = 1,
 \end{aligned}$$

since the latter sequence consists of a triangle-3-cycle and a cycle of length $K - 1$ for which the product, by the induction assumption, is 1.

Case 2. Suppose that $(a_2, b_2) = (c_1, b_1)$. Since star-3-cycles have product 1,

$$\begin{aligned}
 \lambda_{br}(a_1, b_1, c_1)\lambda_{br}(a_2, b_2, c_2) &= \lambda_{br}(a_1, b_1, c_1)\lambda_{br}(c_1, b_1, c_2) \\
 &= \lambda_{br}(a_1, b_1, c_2).
 \end{aligned}$$

Hence, the product of C is equal to

$$\begin{aligned}
 &\lambda_{br}(a_1, b_1, c_2)\lambda_{br}(a_3, b_3, c_3) \cdots \cdots \lambda_{br}(a_K, b_K, c_K) \\
 &= \lambda_{br}(a_1, b_2, c_2)\lambda_{br}(a_3, b_3, c_3) \cdots \cdots \lambda_{br}(a_K, b_K, c_K) = 1,
 \end{aligned}$$

since the latter sequence is a cycle of length $K - 1$ for which the product, by the induction assumption, is 1.

Case 3. Suppose that $(b_K, c_K) = (b_1, a_1)$. Similarly to case 2, we can show here that the product of C is 1.

By induction, the proof is complete. □

PROOF OF LEMMA 4.6 Let $a, b, c \in X$. We show that $\lambda_{br}(a, b, c)\lambda_{br}(b, c, a)\lambda_{br}(c, a, b) = 1$. Define

$$\beta_1 := br([a]|\{b, c\}), \quad \beta_2 := br([c]|\{a, b\}), \quad \beta_3 := br([b]|\{a, c\}).$$

We distinguish two cases.

Case 1. Assume that $\beta_1 \notin \{[b], [c]\}$, $\beta_2 \notin \{[a], [b]\}$ and $\beta_3 \notin \{[a], [c]\}$. Then,

$$\begin{aligned}
 \lambda_{br}(a, b, c) &= \sqrt{\frac{\beta_2(a)}{\beta_1(c)}}, \quad \lambda_{br}(b, c, a) = \sqrt{\frac{\beta_1(b)}{\beta_3(a)}} \quad \text{and} \quad \lambda_{br}(c, a, b) \\
 &= \sqrt{\frac{\beta_3(c)}{\beta_2(b)}},
 \end{aligned}$$

and hence we must show that

$$\beta_1(b)\beta_2(a)\beta_3(c) = \beta_1(c)\beta_2(b)\beta_3(a). \quad (21)$$

Define the belief

$$\beta := \frac{1}{1 - \beta_1(c)\beta_2(a)} (\beta_1(b)\beta_2(a)[a] + \beta_1(b)\beta_2(b)[b] + \beta_1(c)\beta_2(b)[c]).$$

It may be verified, using LIN, that $br(\beta|\{b, c\}) = \beta_1$ and $br(\beta|\{a, b\}) = \beta_2$. We shall prove that $br(\beta|\{a, c\}) = \beta_3$.

Define $\beta^* := \frac{1}{3}[a] + \frac{1}{3}[b] + \frac{1}{3}[c]$, and let $\beta' := \frac{2}{3}\beta^* - \frac{1}{2}\beta$. Since $\beta^* = \frac{1}{3}\beta + \frac{2}{3}\beta'$, we know by LIN that

$$br(\beta^*|\{a, b\}) = \frac{1}{3}br(\beta|\{a, b\}) + \frac{2}{3}br(\beta'|\{a, b\}).$$

As, by LIN, $br(\beta^*|\{a, b\}) = \frac{1}{3}[a] + \frac{1}{3}[b] + \frac{1}{3}\beta_2$, and since we know that $br(\beta|\{a, b\}) = \beta_2$, it follows that

$$\frac{1}{3}\beta_2 + \frac{2}{3}br(\beta'|\{a, b\}) = \frac{1}{3}[a] + \frac{1}{3}[b] + \frac{1}{3}\beta_2,$$

implying that $br(\beta'|\{a, b\}) = \frac{1}{2}[a] + \frac{1}{2}[b]$. In a similar fashion, one can show that $br(\beta'|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$. By TRA, it then follows that $br(\beta'|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$. Since $\beta^* = \frac{1}{3}\beta + \frac{2}{3}\beta'$, LIN implies that

$$br(\beta^*|\{a, c\}) = \frac{1}{3}br(\beta|\{a, c\}) + \frac{2}{3}br(\beta'|\{a, c\}).$$

As $br(\beta^*|\{a, c\}) = \frac{1}{3}[a] + \frac{1}{3}[c] + \frac{1}{3}\beta_3$ and $br(\beta'|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$, it follows that

$$\frac{1}{3}[a] + \frac{1}{3}[c] + \frac{1}{3}\beta_3 = \frac{1}{3}br(\beta|\{a, c\}) + \frac{1}{3}[a] + \frac{1}{3}[c],$$

and hence $br(\beta|\{a, c\}) = \beta_3$. As such, $br(\beta|\{b, c\}) = \beta_1$, $br(\beta|\{a, b\}) = \beta_2$ and $br(\beta|\{a, c\}) = \beta_3$, which, by LIN, implies that

$$\frac{\beta_1(b)}{\beta_1(c)} = \frac{\beta(b)}{\beta(c)}, \quad \frac{\beta_2(a)}{\beta_2(b)} = \frac{\beta(a)}{\beta(b)}, \quad \frac{\beta_3(c)}{\beta_3(a)} = \frac{\beta(c)}{\beta(a)}.$$

But then,

$$\frac{\beta_1(b) \beta_2(a) \beta_3(c)}{\beta_1(c) \beta_2(b) \beta_3(a)} = 1,$$

implying (21).

Case 2. Assume that Case 1 does not hold. We may assume, without loss of generality, that $\beta_1 = [b]$. Then, by Lemma 4.1, we have that $\beta_2 = [b]$, and $\beta_3 \notin \{[a], [c]\}$, and hence

$$\begin{aligned} \lambda_{br}(a, b, c) &= \sqrt{\frac{\beta_3(a)}{\beta_3(c)}}, \quad \lambda_{br}(b, c, a) = \sqrt{\frac{1}{\beta_3(a)}} \quad \text{and} \quad \lambda_{br}(c, a, b) \\ &= \sqrt{\beta_3(c)}. \end{aligned}$$

It is clear that $\lambda_{br}(a, b, c)\lambda_{br}(b, c, a)\lambda_{br}(c, a, b) = 1$, and hence the proof is complete. \square

PROOF OF LEMMA 4.7 Let $a, b, c, d \in X$. We show that $\lambda_{br}(a, d, b)\lambda_{br}(b, d, c)\lambda_{br}(c, d, a) = 1$. Define the beliefs A, B, \dots, P as in Figure 14. We show that $\lambda_{br}(a, d, b) = \sqrt{N_a/O_b}$. We distinguish two cases.

Case 1. Suppose that $A \neq [d]$. Then, $\lambda_{br}(a, d, b) = \sqrt{B_a/A_b}$, and hence we must show that $N_a/O_b = B_a/A_b$. In the intuitive argument we have seen that, by IOI, $br(K|\{a, b\}) = C = br([d]|\{a, b\})$. Hence, by LIN, $K = (1 - \lambda)C + \lambda[d]$ for some $\lambda \in [0, 1]$. Moreover, we have seen in the intuitive argument that $N = br(K|\{a, d\})$. Consequently, by LIN,

$$N = (1 - \lambda)br(C|\{a, d\}) + \lambda[d] = (1 - \lambda)(C_a[a] + C_bB) + \lambda[d],$$

which implies that $N_a = (1 - \lambda)(C_a + C_bB_a)$. In a similar fashion, one can show that $O_b = (1 - \lambda)(C_b + C_aA_b)$, and hence

$$\frac{N_a}{O_b} = \frac{C_a + C_bB_a}{C_b + C_aA_b}. \tag{22}$$

From (21) in the proof of Lemma 4.6, we know that

$$A_dB_aC_b = A_bB_dC_a. \tag{23}$$

Combining (22) and (23) yields

$$\begin{aligned}
 \frac{N_a}{O_b} &= \frac{C_a + C_b B_a}{C_b + C_a A_b} = \frac{A_b B_d (C_a + C_b B_a)}{A_b B_d (C_b + C_a A_b)} \\
 &= \frac{A_d B_a C_b + A_b B_d C_b B_a}{A_b B_d C_b + A_d B_a C_b A_b} \\
 &= \frac{B_a (A_d + A_b B_d)}{A_b (B_d + A_d B_a)} = \frac{B_a (A_d + A_b B_d)}{A_b (B_d + A_d (1 - B_d))} \\
 &= \frac{B_a (A_d + A_b B_d)}{A_b (A_d + A_b B_d)} = \frac{B_a}{A_b},
 \end{aligned}$$

which was to show.

Case 2. Suppose that $A = [d]$. Then, by Lemma 4.1, $B = [d]$, and hence $\lambda_{br}(a, d, b) = \sqrt{C_a/C_b}$. Hence, we must show that $N_a/O_b = C_a/C_b$. We know from (22) that

$$\frac{N_a}{O_b} = \frac{C_a + C_b B_a}{C_b + C_a A_b}$$

which is equal to C_a/C_b , since $B_a = A_b = 0$.

Hence, we may conclude that, in general, $\lambda_{br}(a, d, b) = \sqrt{N_a/O_b}$. Similarly, we can show that $\lambda_{br}(b, d, c) = \sqrt{O_b/P_c}$ and $\lambda_{br}(c, d, a) = \sqrt{P_c/N_a}$. Consequently,

$$\lambda_{br}(a, d, b)\lambda_{br}(b, d, c)\lambda_{br}(c, d, a) = 1,$$

which completes the proof. \square

PROOF Of LEMMA 4.9 The proof for this lemma is basically a formal repetition of the intuitive argument presented in Section 4.2.3, and is therefore omitted. \square

ACKNOWLEDGMENTS

There are some people that I would like to thank for their contribution to this project: Peter Wakker for helpful preliminary discussions before the start of the project, Tim Schulte for many stimulating conversations when I was struggling with the proof of the main theorem, Oliver Schulte for teaching me many interesting aspects within belief revision theory,

Robert van Rooij and, again, Oliver Schulte for pointing at the connection between my project and the idea of imaging, János Flesch and, again, Tim Schulteis for carefully reading an earlier version of this manuscript, and, finally, Geir Asheim, Luis Gonzalez, Werner Güth and Anthony Ziegelmeyer for helpful discussions.

NOTES

1. See Lehmann et al. (2001) for an approach in which the similarity among states is represented by a *pseudo-distance*, which is not necessarily symmetric.
2. We borrow this term from Gärdenfors (1988).
3. The reader may wonder why it is not sufficient, in general, to simply require that all angles between cornerpoints in the polytope induced by φ are acute. Here is a counterexample: Take four states $\{a, b, c, d\}$ and choose $\varphi([a]) = (0, 0, 0)$, $\varphi([b]) = (2, 0, 0)$, $\varphi([c]) = (1, \sqrt{3}, 0)$ and $\varphi([d]) = (1, -\frac{1}{5}, 2)$. Then, all angles between cornerpoints in the induced polytope are acute, but the orthogonal projection of $\varphi([d])$ on the affine space spanned by $\varphi([a])$, $\varphi([b])$ and $\varphi([c])$ is not in the convex hull of these three points.
4. By symmetric, we mean that $d(a, b) = d(b, a)$ for all a, b .
5. That is, the three lines intersect at a single point.

REFERENCES

- Alchourrón, C.E. and Makinson, D. (1982), The logic of theory change: Contraction functions and their associated revision function, *Theoria* 48, 14–37.
- Anscombe, F.J. and Aumann, R. (1963), A definition of subjective probability, *Annals of Mathematical Statistics* 34, 199–205.
- Epstein, L.G. (2005), *An Axiomatic Model of non-Bayesian Updating*, University of Rochester.
- Epstein, L.G. and Schneider, M. (2003), Recursive multiple-priors, *Journal of Economic Theory* 113, 1–31.
- Gärdenfors, P. (1988), *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, Cambridge, MA.
- Ghirardato, P. (2002), Revisiting Savage in a conditional world, *Economic Theory* 20, 83–92.
- Gul, F. and Pesendorfer, W. (2001), Temptation and self-control, *Econometrica* 69, 1403–1435.

- Hendon, E., Jacobsen, J. and Sloth, B. (1996), The one-shot deviation principle for sequential rationality, *Games and Economic Behavior* 12, 274–282.
- Lehmann, D., Magidor, M. and Schlechta, K. (2001), Distance semantics for belief revision, *Journal of Symbolic Logic* 66, 295–317.
- Lewis, D.K. (1973), *Counterfactuals*, Blackwell, Oxford.
- Lewis, D.K. (1976), Probabilities of conditionals and conditional probabilities, *Philosophical Review* 85, 297–315.
- Katsuno, H. and Mendelzon, A.O. (1992), On the difference between updating a knowledge base and revising it, in Gärdenfors, P. (ed.), *Belief Revision*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, pp. 183–203.
- Majumdar, D. (2004), An axiomatic characterization of Bayes' rule, *Mathematical Social Sciences* 47, 261–273.
- Perea, A. (2002), A note on the one-deviation property in extensive form games, *Games and Economic Behavior* 40, 322–338.
- Savage, L.J. (1954), *The Foundations of Statistics*, Wiley, New York.
- Schulte, O. (2002), Minimal belief change, Pareto-optimality and logical consequence, *Economic Theory* 19, 105–144.

Address for correspondence: Andrés Perea, Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. Fax: +31-43-3884874; E-mail: a.perea@ke.unimaas.nl