A model of minimal probabilistic belief revision

Citation for published version (APA):

Perea y Monsuwé, A. (2005). *A model of minimal probabilistic belief revision*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 035
<https://doi.org/10.26481/umamet.2005035>

Document status and date:

Published: 01/01/2005

DOI:

[10.26481/umamet.2005035](https://doi.org/10.26481/umamet.2005035)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

A Model of Minimal Probabilistic Belief Revision*

Andrés Perea[†]
Maastricht University

This Version: September 2005

Abstract

A probabilistic belief revision function assigns to every initial probabilistic belief and every observable event some revised probabilistic belief that only attaches positive probability to states in this event. We propose three axioms for belief revision functions: (1) *linearity*, meaning that if the decision maker observes that the true state is in $\{a, b\}$, and hence state c is impossible, then the proportions of c 's initial probability that are shifted to a and b , respectively, should be independent of c 's initial probability; (2) *transitivity*, stating that if the decision maker deems belief β equally similar to states a and b , and deems β equally similar to states b and c , then he should deem β equally similar to states a and c ; (3) *information-order independence*, stating that the way in which information is received should not matter for the eventual revised belief. We show that a belief revision function satisfies the three axioms above if and only if there is some linear one-to-one function φ , transforming the belief simplex into a polytope that is closed under orthogonal projections, such that the belief revision function satisfies minimal belief revision with respect to φ . By the latter, we mean that the decision maker, when having initial belief β_1 and observing the event E , always chooses the revised belief β_2 that attaches positive probability only to states in E and for which $\varphi(\beta_2)$ has minimal Euclidean distance to $\varphi(\beta_1)$.

Keywords: Belief revision, probabilistic beliefs, non-Bayesian updating.

JEL Classification: C73, D81, D83.

1 Introduction

1.1 Motivation and Outline

Beliefs and belief revision play a fundamental role in decision making, both on a professional and an informal level. For instance, a physician who is uncertain about the precise disease of a patient may prescribe a treatment on the basis of the subjective probabilities, or beliefs, he assigns to each of the possible diseases. Upon observing new symptoms, the physician may exclude some of these diseases, and may redistribute the probabilities assigned to the excluded diseases among the diseases he still deems possible. For the eventual treatment to be prescribed,

*First of all, I would like to thank Tim Schulteis for the many stimulating conversations we had about this project. I also thank Luis Gonzalez, Werner Güth, Peter Wakker and Anthony Ziegelmeyer for helpful discussions.

[†]Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: a.perea@ke.unimaas.nl, Web: <http://www.personeel.unimaas.nl/a.perea/>

Figure 1: An introductory example

it is therefore crucial how the physician redistributes these probabilities, that is, how he revises his beliefs upon receiving this new information. In this paper we shall focus on exactly this problem, namely how decision makers may revise their probabilistic beliefs upon acquiring new information.

In most decision theoretic and game theoretic models, it is assumed that the decision maker (or player) revises his subjective beliefs by Bayesian updating whenever the newly acquired information does not rule out all states to which he previously assigned positive probabilities. Whereas this assumption is completely natural if the uncertain events have well-known objective probabilities, we shall argue in the next few paragraphs that Bayesian updating may not be the only reasonable way of revising beliefs when no such objective probabilities exist.

As to clarify this point, consider the game tree in Figure 1. We analyze the situation from player 2's viewpoint. Assume that player 2 believes that player 1 is rational, but that player 2 has uncertainty about player 1's preference relation over his actions. More precisely, player 2 believes that one of the following three states is the true state:

- state A : $a \succ c \succ b$;
- state B : $b \succ c \succ a$;
- state C : $c \succ a \succ b$.

Hence, A is the state where player 1 strictly prefers a over c and strictly prefers c over b . Similarly for states B and C . Suppose first that player 1 is drawn randomly from a large population of rational players of which it is *known* that exactly one third has preference relation A , one third has preference relation B and one third has preference relation C . As such, player 2 will assign probability one third to states A, B and C before the game starts, since these probabilities coincide with the relative frequencies of the states in the population. When player 2 has to choose at his information set, he knows that player 1 has not chosen c , and hence must conclude that state C is no longer possible. Since the new relative frequencies of A and B , conditional on the event that C does not occur, are both one half, player 2's revised probabilities assigned to A and B should be one half, as Bayesian updating suggests.

Consider next the situation where no such population exists, but where player 2 assigns *subjective* probability one to state C before the game starts. Upon observing that his information set is reached, he must conclude that C is not the true state, and hence must revise his subjective belief. The idea of minimal belief revision (see Schulte (2002), for instance, for an excellent

discussion of this idea and an overview of the literature) suggests that the new belief should be “as similar as possible” to the previous belief while explaining the newly observed event that player 1 has not chosen c . Among the states A and B that could explain this event, state A seems to be the state that, intuitively, is most similar to C . Namely, preference relations A and C only differ on one pairwise ranking of actions, $\{a, c\}$, whereas B and C differ on two pairwise rankings, $\{a, b\}$ and $\{b, c\}$. As such, the concept of minimal belief revision suggests that player 2’s revised belief should attach a higher probability to A , the “more similar state”, than to B , the “less similar state”.

A similar line of reasoning can be adopted, however, if player 2 initially assigns subjective probability one third to states A, B and C , instead of initially assigning subjective belief one to C . In this case, player 2, upon observing that his information set is reached, faces the task of redistributing the weight one third, initially assigned to the excluded state C , among the non-excluded states A and B . Within the spirit of minimal belief revision as discussed above, player 2 should then transfer more than fifty percent of this weight to state A . As a consequence, minimal belief revision leads player 2’s revised belief to assign a higher probability to A than to B , and hence would lead to a different revised belief than obtained by Bayesian updating. The crucial difference with the first scenario discussed above, in which the state was randomly drawn from a population with known proportions, is that in the first scenario the event of reaching player 2’s information set does not change player 2’s belief in the objective a-priori probabilities of states A, B and C . Therefore, player 2 has no reason to change his belief about the relative likelihoods of A and B , and hence Bayesian updating is the only sensible way of revising the beliefs. In the last scenario, however, no such objective a-priori probabilities exist. The initial subjective probabilities assigned to states A, B and C only reflect player 2’s initial personal opinion about the relative likelihoods of the states. In particular, upon observing that player 1 has not chosen c , he may change his opinion about the relative likelihoods of A and B , and the concept of minimal belief revision indeed suggests player 2 to do so.

In this paper we propose a model of probabilistic belief revision that incorporates this idea of minimal belief revision. To this purpose, we first define the concept of a *belief revision function* which is a function that assigns to every possible initial belief and every possible event some revised belief attaching probability zero to all states outside this event. We then require that the decision maker, upon observing that the real state is in $\{a, b\}$ and state c is impossible, should redistribute the weight initially attached to c among the non-excluded states in exactly the same way as he *would* have done if he initially assigned weight one to c . As an illustration, suppose that in the example above player 2 initially assigns subjective probability one to C , and upon observing that his information set has been reached assigns probability α to A and probability $1 - \alpha$ to B . These revised probabilities intuitively reflect the similarity between C and A , as opposed to the similarity between C and B , as perceived by player 2: The higher α , the higher the perceived similarity between C and A compared to the perceived similarity between C and B . We then require that, if player 2 would initially assign probability $\beta(C) < 1$ to C , he should redistribute the weight $\beta(C)$ among the states A and B using these same proportions α and $1 - \alpha$. That is, the revised probabilities of A and B should be $\beta(A) + \alpha\beta(C)$ and $\beta(B) + (1 - \alpha)\beta(C)$, respectively, where $\beta(A)$ and $\beta(B)$ are the initial probabilities of A and B . Throughout the paper we shall refer to this condition as *linearity*, as it implies that the revised belief depends linearly upon the initial belief.

The second condition we impose on a belief revision function states that, whenever the decision maker deems belief β “equally similar” to states x and y , and deems β “equally similar” to states y and z , he should deem β “equally similar” to states x and z . Here, by “ β equally similar to states x and y ” we mean that the decision maker, when having initial belief β and upon observing the event that only states x and y are possible, assigns equal probabilities to x and y . The second condition thus states that the “equally-similar-to relation” should be transitive. For this reason, we call this condition *transitivity*.

The third and final condition we impose states that it should not matter for the revised belief whether the acquired information is received at once or stepwise, and in the latter case it should not matter in which order the various pieces of information are received. This condition is called *information-order independence*.

Our main theorem shows that a belief revision function is linear, transitive and information-order independent if and only if it satisfies minimal belief revision with respect to some linear one-to-one function φ transforming the simplex of beliefs into some polytope that is closed under orthogonal projections. By minimal belief revision with respect to φ we mean that the decision maker, when having initial belief β_1 and observing event E , should choose the revised belief β_2 such that $\varphi(\beta_2)$ has the least Euclidean distance to $\varphi(\beta_1)$ amongst the beliefs that are compatible with E . Here, the vector $\varphi(\beta)$ may be interpreted as a list of characteristics that describes the belief β . As such, stating that the belief revision function satisfies minimal belief revision with respect to φ means that the decision maker always selects the revised belief for which the list of characteristics is as close as possible to the list of characteristics of the initial belief. By “closed under orthogonal projections” we mean that for every point in the polytope, its orthogonal projection on each of the faces is contained in the polytope.

As to illustrate how the function φ may be chosen naturally in particular examples, consider again the game tree of Figure 1. Let $[A]$, $[B]$ and $[C]$ be the beliefs that assign probability one to the states A , B and C , respectively. A possible, natural way to choose the list of characteristics $\varphi([A])$ for belief $[A]$ would be to identify $[A]$ with a Boolean vector of size 3 stating for each of the pairwise rankings $a \succ b$, $a \succ c$ and $b \succ c$ whether it is true or false. By substituting ‘1’ for ‘true’ and ‘0’ for false, one would thus define $\varphi([A]) = (1, 1, 0)$. Similarly, $\varphi([B]) = (0, 0, 1)$ and $\varphi([C]) = (1, 0, 0)$. By linearity of φ , we must define $\varphi(\beta) = \beta(A)\varphi([A]) + \beta(B)\varphi([B]) + \beta(C)\varphi([C])$ for every belief β . Hence, the list of characteristics $\varphi(\beta)$ specifies for each of the three pairwise rankings above the probability β attaches to the event of this ranking being true, and this list could then be used as a criterion to define the distance between two beliefs. In particular, if the decision maker initially holds the belief $\beta_1 = [C]$ and observes the event that c has not been chosen, then the revised belief β_2 whose list of characteristics $\varphi(\beta_2)$ is closest to $\varphi(\beta_1)$ is given by $\beta_2 = (\frac{2}{3}, \frac{1}{3}, 0)$ assigning probability $2/3$ to A and probability $1/3$ to B . As a consequence, the belief revision function that satisfies minimal belief revision with respect to φ should prescribe the revised belief $(\frac{2}{3}, \frac{1}{3}, 0)$ in this situation, reflecting the fact that the decision maker deems C more similar to A than to B .

The outline of this 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 as discussed above. In Section 3 we present the main theorem of this paper. In Section 4 we provide the reader with an overview of the proof of the main theorem. We present the lemmas that are used to show the main result, and present for each lemma an intuitive, often

geometrical, argument that reveals the main ideas behind the formal proof. The full algebraic proofs are presented in the appendix. The reason for this procedure is that the proof is rather long, and therefore we prefer to give the reader an overview of the main steps first, before confronting him or her with all the technical details. We conclude in Section 5 with a discussion of the main theorem and its proof.

1.2 Related Literature

The problem of belief revision has been studied in many different areas, including logic, philosophy, computer science, decision theory and game theory. The study of this problem in the former three areas has led to a separate field, known as *belief revision theory*. Our approach is similar to most investigations in belief revision theory in the sense that we put restrictions directly on the belief revision function, rather than embedding the belief revision function into one-person or multi-person decision problems, as is common in decision theory and game theory. Within the area of belief revision theory, our representation theorem is similar, in spirit, to Grove's representation theorem of (non-probabilistic) belief revision functions (Grove (1988)). Grove has shown that a belief revision function satisfies the AGM-axioms (Alchourrón and Makinson (1982), Gärdenfors (1988)) if and only if for every initial belief, consisting of some subset of states¹, there exists some plausibility ranking on the set of states such that for every piece of new information, the revised belief contains all those states that are most plausible among the states that are compatible with the new information. In our representation theorem, the plausibility of each revised belief β_2 is measured according to the Euclidean distance between $\varphi(\beta_2)$ and $\varphi(\beta_1)$, where β_1 is the initial belief. Within this interpretation, the decision maker in our model always chooses the revised belief β_2 that is most plausible, given the initial belief β_1 , among all beliefs that are compatible with the new information. Our representation theorem may thus be interpreted as a probabilistic variant of Grove's representation theorem.

In decision theory, the common approach to belief revision is to extract the decision maker's belief revision function from his preferences over acts, or menus of acts, in a dynamic decision problem. Many papers focus on the relationship between belief revision on the one hand, and *dynamic consistency* of preferences, or the violation thereof, on the other hand. Intuitively, dynamic consistency means that, upon observing some non-null event, the decision maker should rank two acts in the same way as he would have done initially, before observing this event. 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 belief revision functions studied in our paper does not satisfy Bayesian updating, the two results above imply

¹The interpretation is that the decision maker initially believes that one of the states in the subset is the true state, without exactly knowing which one (if the subset contains more than one state, of course).

that embedding these belief revision functions 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 belief revision function that satisfies our three axioms, 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*. By the latter, we mean the condition that a vector of *ad interim* optimal actions for a player always induces an *ex ante* optimal strategy. That is, if a player, at each of his information sets, chooses an action that is optimal given his conditional belief about the opponents' strategies, and that correctly anticipates on his own behavior at future information sets, then this will lead to a strategy that is *ex ante* optimal from each of his information sets onwards. It may be verified that dynamic consistency of preferences implies the one-deviation property. Hendon, Jacobsen and Sloth (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 belief revision functions satisfying our three axioms 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 beliefs 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 conditional belief at h_2 , and subsequently should choose the action at h_1 that is optimal given his conditional belief at h_1 about the opponents' choices, and given his own future choice at h_2 .

Within the existing literature on belief revision, the paper that perhaps comes closest to ours is Majumdar (2004). Similar to our model, it imposes several axioms on probabilistic belief revision functions. It should be mentioned that Majumdar's definition of a belief revision function is more restrictive, as it is only defined whenever the observed event does not rule out all states that are attached positive probability by the initial belief. In our model, a belief revision

function is also defined for these cases. The crucial difference between the two models lies in the belief revision functions that are characterized by the axioms: While our axioms characterize a class of belief revision functions that does *not* satisfy Bayesian updating, Majumdar’s axioms characterize *exactly* Bayesian updating. As such, it may be interesting to compare these two sets of axioms here. Majumdar’s axioms are: (1) *path independence*, which is identical to our axiom of information-order independence; (2) *symmetry*, stating that if the initial probabilities are permuted, the revised probabilities should be permuted by the same permutation; (3) *continuity*, stating that the revised belief should depend continuously upon the initial belief; (4) *monotonicity*, stating that the revised probability on a state should be as least as large as the initial probability on a state; and (5) *no mistake hypothesis*, stating that if some state initially has probability zero, then the revised probability on that state should remain zero. The main result in this paper is to show that a belief revision function satisfies the axioms (1) - (5) if and only if it satisfies Bayesian updating.

On the other hand, our axioms of linearity, transitivity and information-order independence characterize belief revision functions that satisfy minimal belief revision with respect to some function φ that is linear, one-to-one, and closed under orthogonal projections. It is easily seen that each such function satisfies path independence, continuity and monotonicity, but not symmetry and the no mistake hypothesis. In the example of Figure 1, for instance, take the function φ we have constructed above, and assume that the decision maker initially assigns equal weight to states A and B . If the decision maker observes that state C is impossible, and if the belief revision function satisfies minimal belief revision with respect to φ , then more weight should be shifted towards A (the state more similar to C) than to B (the state less similar to C), and hence the revised probability on A will be larger than the revised probability on B . Symmetry, on the other hand, implies that the revised probabilities on A and B should be equal. Hence, symmetry is violated. In order to see that in this example the no mistake hypothesis is violated also, assume that the decision maker initially assigns probability zero to A and probabilities $1/2$ to both B and C . If he observes that C is ruled out, and satisfies minimal belief revision with respect to φ , then he assigns positive probability to A in the revised belief. The reason is that he shifts a strictly positive part of the weight $1/2$, initially assigned to C , towards A , the state more similar to C . As a consequence, minimal belief revision with respect to φ implies that a decision maker may initially believe that some state has probability zero, but believe later on that this same state has positive probability. This makes sense, since assigning probability zero to a state does not necessarily mean that one deems this state completely impossible. It only means that one deems this state considerably less plausible than some of the other states. If a decision maker observes that some of the states that he deemed plausible is no longer possible, than he may transfer a part of its weight towards a state that is similar to the excluded state, but that he deemed implausible before. Thus, a subjectively implausible state may be turned into a subjective plausible state whenever it is similar to a state that is excluded by the newly acquired information, and that was deemed subjectively plausible before.

Reversely, it is easily seen that Bayesian updating satisfies our axioms of transitivity and information-order independence, but violates our axiom of linearity. Hence, the axioms of symmetry and no mistake hypothesis in Majumdar’s model, and the axiom of linearity in our model, are precisely the axioms that distinguish Bayesian updating from the notion of minimal belief revision discussed here.

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 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, within the spirit of minimal belief revision, 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, minimal belief revision suggests that he should shift the proportion α of $\beta(c)$ to a , and shift the proportion $1 - \alpha$ of $\beta(c)$ to b . That is, the proportions 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

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

In particular, linearity implies that $br(\beta|\{a, b\}) = \beta$ whenever β assigns positive probability only to states in $\{a, b\}$. That is, if the initial belief β is already in accordance with the event $\{a, b\}$, the belief should not be changed. Mathematically speaking, the axiom of linearity implies that, for some fixed event $E = \{a, b\}$, the revised belief is a linear function of the initial belief. The first triangle in Figure 2 illustrates this fact for the state space $X = \{a, b, c\}$, the fixed

Figure 2: Linear versus Bayesian belief revision

event $E = \{a, b\}$ and various initial beliefs in $B(X)$. This is an example of a decision maker who deems c more similar to a than to b , and who therefore always shifts more weight towards a than towards b . As a comparison, the second triangle illustrates how Bayesian belief revision operates for the same event E and the same initial beliefs. Hence, geometrically speaking, the crucial difference between linear belief revision and Bayesian belief revision is that the former always revises beliefs in the same direction, namely towards the state that is perceived most similar to the excluded state c , while the direction of revision in Bayesian belief revision is completely determined by the initial relative weights assigned to the non-excluded states a and b , and therefore varies for different initial beliefs. As a consequence, the revised beliefs in a linear belief revision function are Lipschitz-continuous with respect to the initial belief, meaning that small changes in the initial belief always lead to small changes in the revised belief, while the revised belief in Bayesian belief revision may change dramatically in response to small changes in the initial belief close to $[c]$.

The second axiom, transitivity, states that, whenever a belief β is perceived “equally similar” to states a and b , and is perceived “equally similar” to states b and c , then the belief β should be perceived “equally similar” to states 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 probability one beliefs $[a]$ and $[b]$. In other words, the belief β is perceived equally similar to the states 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 $br(\beta|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$, it should also hold that $br(\beta|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$.*

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 3 provides an 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

Figure 3: Transitivity

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 4 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

Intuitively, a belief revision function br satisfies *minimal belief revision* if for every initial belief β_1 and every event E , the function br selects the revised belief β_2 in $B(X|E)$ that is “most similar” to β_1 . A possible way to formalize the phrase “most similar” would be to identify each belief β with some vector $\varphi(\beta)$, and to require that $\varphi(\beta_2)$ should have minimal Euclidean distance to $\varphi(\beta_1)$ amongst all beliefs in $B(X|E)$. Formally, let $\mathbb{R}^m$ be some Euclidean space with $m \in \mathbb{N}$, and let φ be some one-to-one function from $B(X)$ to $\mathbb{R}^m$. For every two vectors $v, w \in \mathbb{R}^m$, we denote by $\|v - w\|$ the Euclidean distance between v and w .

Definition 3.1 (*Minimal belief revision*) We say that the belief revision function br satisfies *minimal belief revision with respect to φ* if for every initial belief $\beta_1 \in B(X)$ and every event $E \subseteq X$:

$$\|\varphi(br(\beta_1|E)) - \varphi(\beta_1)\| \leq \|\varphi(\beta_2) - \varphi(\beta_1)\|$$

for all $\beta_2 \in B(X|E)$.

Figure 4: Information-order independence

The main theorem in this paper states that a belief revision function br is linear, transitive and information-order independent if and only if there is some *linear* one-to-one function φ , *closed under orthogonal projections*, such that br satisfies minimal belief revision with respect to φ . We first formalize what we mean by “linear” and “closed under orthogonal projections”.

Definition 3.2 (*Linear function*) A 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 . In the example of the introduction, for instance, the vector $\varphi([A])$ describes for each of the pairwise rankings whether this ranking is true or false at A , and can thus be seen as a vector of characteristics for state A . Similarly for states B and C . 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 $\{\varphi([x]) \mid x \in X\}$. In the introductory example, the vector $\varphi(\beta)$ specifies for each pairwise ranking the probability that β assigns to the event of this ranking being true, and hence represents the expected vector of characteristics under β .

If the function φ is linear and one-to-one, it is easily seen that φ transforms the original belief simplex $B(X)$ into a polytope $\varphi(B(X))$ with the same dimension, and cornerpoints $\{\varphi([x]) \mid x \in X\}$. The following condition states that for every point and every face in this polytope, the orthogonal projection of this point on the face is contained in the polytope.

Definition 3.3 (*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 the face $\varphi(B(X|E))$. Note that this orthogonal projection is always unique. We are now ready to present the main theorem of this paper.

Theorem 3.4 (*Representation Theorem*). *Let X be a finite set of 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 $\mathbb{R}^m$ and a linear one-to-one function $\varphi : B(X) \rightarrow \mathbb{R}^m$, closed under orthogonal projections, such that br satisfies minimal belief revision with respect to φ .*

For the remainder of this paper, whenever we say that φ *represents* the belief revision function br , we mean that br satisfies minimal belief revision with respect to φ .

4 Proof of the Representation Theorem

4.1 Minimal Belief Revision implies Axioms

The easy part is to show the implication from (2) to (1). Assume that br satisfies minimal belief revision (MBR) with respect to 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 $\beta_2 = br(\beta_1|\{a, b\})$. By MBR,

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

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 $\beta_1^{rev} := br(\beta_1|\{a, b\})$, $\beta_2^{rev} := br(\beta_2|\{a, b\})$, $\beta^{rev} := br(\beta|\{a, b\})$. We show that $\beta^{rev} = \alpha\beta_1^{rev} + (1 - \alpha)\beta_2^{rev}$. By MBR,

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

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

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

Similarly,

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

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

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

which, by MBR, implies that $\beta^{rev} = \tilde{\beta}^{rev}$. Hence, $\beta^{rev} = \alpha\beta_1^{rev} + (1 - \alpha)\beta_2^{rev}$. 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 $br(\beta|\{a, c\}) = \frac{1}{2}[a] + \frac{1}{2}[c]$. By MBR and the fact that φ is closed under orthogonal projections, this is equivalent to stating that

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

This, in turn, implies that

$$\|\varphi(\beta) - \varphi(a)\| = \|\varphi(\beta) - \varphi(b)\| \text{ and } \|\varphi(\beta) - \varphi(a)\| = \|\varphi(\beta) - \varphi(c)\|.$$

As such,

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

yielding

$$\varphi(\tfrac{1}{2}[b] + \tfrac{1}{2}[c]) - \varphi(\beta) \perp \varphi([b]) - \varphi([c])$$

and hence, by MBR, $br(\beta|\{b, c\}) = \frac{1}{2}[b] + \frac{1}{2}[c]$. 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. By MBR, 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

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

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 Axioms imply Minimal Belief Revision

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 $br([b]|\{a, c\}) \notin \{[a], [c]\}$.*

Intuitive argument. Consider Figure 5 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 Figure 6 as an illustration. Let

Figure 5: A geometrical argument for Lemma 4.1

Figure 6: A geometrical argument for Lemma 4.1

Figure 7: Belief revision on triangles is decisive

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

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. $\diamond$

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

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 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

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

intersection point. On the other hand, IOI of br implies that $br(br([d]|\{a, b, c\})|\{a, b\}) = A$ and $br(br([d]|\{a, b, c\})|\{b, c\}) = B$. 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 $br([d]|\{a, b, c\}) = br_\varphi([d]|\{a, b, c\})$. By repeating this argument for the mappings of $[a]$, $[b]$ and $[c]$ on the opposite face of the simplex, and by using LIN of br and br_φ , we may conclude the following: if br and br_φ coincide on every triangle, then $br = br_\varphi$, as was to show. $\diamond$

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

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

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 degrees. Then, $\beta_1(c) > 0$, 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

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

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

If the angle α is exactly 90 degrees, 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 degrees) 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

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

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

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

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

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

$\lambda_{br}(a, b, c) > 0$. 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

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

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).*

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

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

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. $\diamond$

4.2.2 Existence of distance function d solving system (6)

In this part we prove that there is some symmetric² 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, $\lambda_{br}(a_2, b_2, c_2) = \lambda_{br}(b_1, c_1, c_2)$ or $\lambda_{br}(a_2, b_2, c_2) = \lambda_{br}(c_1, b_1, c_2)$. 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 $\lambda_{br}(a_1, b_1, c_1) \cdot \lambda_{br}(a_2, b_2, c_2) \cdot \dots \cdot \lambda_{br}(a_K, b_K, c_K)$.

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

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

²By symmetric, we mean that $d(a, b) = d(b, a)$ for all a, b .

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

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

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

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.

◇

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 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

$$\begin{aligned} & \lambda_{br}(a, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, a)\lambda_{br}(d, a, b) \\ &= \lambda_{br}(a, b, d)\lambda_{br}(d, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, b)\lambda_{br}(b, d, a)\lambda_{br}(d, a, b) \\ &= [\lambda_{br}(a, b, d)\lambda_{br}(b, d, a)\lambda_{br}(d, a, b)] [\lambda_{br}(d, b, c)\lambda_{br}(b, c, d)\lambda_{br}(c, d, b)] = 1, \end{aligned}$$

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. $\diamond$

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 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

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

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

By Ceva's theorem, this is equivalent to showing that the three lines aA , bB and cC are concurrent³. 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

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

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 the triangle (see Figure 13). That is, $br(D|\{a, c\}) = A$ 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. $\diamond$

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

³That is, the three lines intersect at a single point.

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

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 definition,

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

where B_a is the probability that B assigns to a , and so on. Hence, we must prove 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. Moreover, by IOI,

$$br(K|\{a, d\}) = br(J|\{a, d\}) = N,$$

and hence, by LIN, 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. $\diamond$

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 satisfy minimal belief revision with respect to φ . This would thus complete the proof of our Theorem 3.4.

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 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

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

Figure 15: Construction of φ for three states

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,

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

By LIN of br , we have that $B = A_a[a] + A_b C$ and hence $B_c = A_b C_c$. As such,

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

Since d and λ_{br} satisfy (6), it follows that $d(a, c)/d(a, b) = \lambda_{br}(b, a, c)$, which implies that $C_c d^2(a, c) = A_b d^2(a, b)$. Consequently,

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

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 $d(b, c)/d(a, c) = \lambda_{br}(a, c, b)$, which yields $C_a/E_b = d^2(b, c)/d^2(a, c)$. Hence, $B_a/D_b = d^2(b, c)/d^2(a, c)$, which implies $\sqrt{B_a}d(a, c) = \sqrt{D_b}d(b, c)$. As such, we may conclude that

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

Figure 16: Construction of φ for four states

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. $\diamond$

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 $A := br([d]|\{a, b, c\})$ and $B := br(A|\{a, d\})$, as depicted in Figure 16. As φ has to be constructed such that br satisfies minimal belief revision with respect to φ , 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

$$h = \sqrt{B_a} \|\varphi([d]) - \varphi([a])\| = \sqrt{B_a} d(a, d), \quad (8)$$

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.

Figure 17: Construction of φ for four states

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

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

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,

$$B_d d^2(a, d) = A_b D_d d^2(a, d) + A_c H_d d^2(a, d).$$

Since d and λ_{br} satisfy (6), we have that $d(a, d)/d(a, b) = \lambda_{br}(b, a, d)$, which implies that $D_d d^2(a, d) = C_b d^2(a, b)$. Similarly, $d(a, d)/d(a, c) = \lambda_{br}(c, a, d)$, implying that $H_d d^2(a, d) = G_c d^2(a, c)$. Substituting these identities in the equation above yields

$$B_d d^2(a, d) = A_b C_b d^2(a, b) + A_c G_c d^2(a, c) = A_b C_b \|\varphi([b]) - \varphi([a])\|^2 + A_c G_c \|\varphi([c]) - \varphi([a])\|^2. \quad (10)$$

By IOI of br , we have that $C = br(A|\{a, b\})$. From the case of three states we know that br satisfies minimal belief revision with respect to 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,

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

which implies that

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

Similarly, we may conclude that

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

Together with (10), this yields

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

Substituting this result in (9) 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 $d(b, d)/d(a, d) = \lambda_{br}(a, d, b)$, 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 (8) 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. $\diamond$

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

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

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,

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

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

$$\varphi([x_k]) = \varphi(A) + \sqrt{C_{x_i}}d(x_i, x_k).$$

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 $\varphi(B(X|\{x_1, \dots, x_{k-1}\}))$. As such, we may conclude that for every $k \in \{3, \dots, n\}$ and every subset $E \subseteq \{x_1, \dots, x_{k-1}\}$, the orthogonal projection of $\varphi([x_k])$ on $\varphi(B(X|E))$ lies in $\varphi(B(X|E))$.

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. $\diamond$

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 satisfies minimal belief revision with respect to 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 Independence of the Axioms

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

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.

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 $br_2(\frac{1}{2}[a] + \frac{1}{2}[c]|\{a, b\}) = [a] \neq \frac{1}{2}[a] + \frac{1}{2}[b]$.

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

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

5.2 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 satisfies minimal belief revision with respect to φ . 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 satisfies minimal belief revision with respect to 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 satisfies minimal belief revision with respect to φ ? 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 $\alpha > 0$ with $d_{\tilde{\varphi}}(a, b) = \alpha d_\varphi(a, b)$ 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 satisfies minimal belief revision with respect to φ and $\tilde{\varphi}$. 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 satisfies minimal belief revision with respect to φ , 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 also satisfies minimal belief revision with respect to $\tilde{\varphi}$. In other words, if we have found one φ with $br = br_\varphi$, then we know how to generate all other $\tilde{\varphi}$ with $br = br_{\tilde{\varphi}}$.

5.3 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 such that br satisfies minimal belief revision with respect to φ . In many practical examples it seems plausible that the decision maker reasons in the other direction: he may first choose a function φ , measuring the similarity between states from his personal perspective, and then minimally revise his beliefs with respect to φ . That is, he first chooses φ , then br . 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 , like we did with the example in the introduction. 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 $\sqrt{\text{dis}(a, b)}$, where $\text{dis}(a, b)$ 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 that satisfies minimal belief revision with respect to φ would be such that

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

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 .

6 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. ■

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*

$$br(\beta|E) = \sum_{x \in E} \beta(x)[x] + \sum_{x \in X \setminus E} \beta(x)br([x]|E). \quad (11)$$

Proof of Lemma 6.1. We prove (11) by induction on $|E|$, where $|E|$ denotes the cardinality of E .

If $|E| = 2$, (11) follows directly from LIN of br .

Now, take some E with $|E| = k > 2$, and suppose that (11) 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

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

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,

$$\begin{bmatrix} br([a]|E \setminus \{a\})(b) & 1 \\ 1 & br([b]|E \setminus \{b\})(a) \end{bmatrix} \begin{bmatrix} \beta_E(a) \\ \beta_E(b) \end{bmatrix} = \begin{bmatrix} br(\beta|E \setminus \{a\})(b) \\ br(\beta|E \setminus \{b\})(a) \end{bmatrix}. \quad (12)$$

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 (12) has a unique solution $(\beta_E(a), \beta_E(b))$ with, in particular,

$$br(\beta|E)(a) = \beta_E(a) = \frac{br(\beta|E \setminus \{b\})(a) - br([b]|E \setminus \{b\})(a) \cdot br(\beta|E \setminus \{a\})(b)}{1 - br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a)}. \quad (13)$$

By the induction assumption, we know that

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

which implies that

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

By substituting these two equations in (13), we eventually obtain

$$\begin{aligned}
br(\beta|E)(a) &= \beta(a) + \sum_{x \in X \setminus E} \beta(x) \frac{br([x]|E \setminus \{b\})(a) - br([b]|E \setminus \{b\})(a) \cdot br([x]|E \setminus \{a\})(b)}{1 - br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a)} \\
&= \beta(a) + \sum_{x \in X \setminus E} \beta(x) br([x]|E)(a),
\end{aligned} \tag{14}$$

where the latter equality follows from applying (13) to the initial belief $[x]$. Since (14) 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. ■

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

$$br(\beta|E) = br_\varphi(\beta|E) \tag{15}$$

for every belief $\beta \in B(X)$ and every event $E \subseteq X$. We show (15) by induction on $|E|$.

If $|E| = 2$, say $E = \{a, b\}$, (15) 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 (15) 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 (13). That is,

$$\begin{aligned}
br(\beta|E)(a) &= \frac{br(\beta|E \setminus \{b\})(a) - br([b]|E \setminus \{b\})(a) \cdot br(\beta|E \setminus \{a\})(b)}{1 - br([a]|E \setminus \{a\})(b) \cdot br([b]|E \setminus \{b\})(a)}, \\
br_\varphi(\beta|E)(a) &= \frac{br_\varphi(\beta|E \setminus \{b\})(a) - br_\varphi([b]|E \setminus \{b\})(a) \cdot br_\varphi(\beta|E \setminus \{a\})(b)}{1 - br_\varphi([a]|E \setminus \{a\})(b) \cdot br_\varphi([b]|E \setminus \{b\})(a)}
\end{aligned}$$

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. ■

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

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

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

$$\begin{aligned} \beta_1(c) d_\varphi^2(b, c) &= \beta_2(a) d_\varphi^2(a, b), \\ \beta_3(a) d_\varphi^2(c, a) &= \beta_1(b) d_\varphi^2(b, c), \\ \beta_2(b) d_\varphi^2(a, b) &= \beta_3(c) d_\varphi^2(c, a). \end{aligned}$$

It may 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

$$\begin{aligned} \beta_1(c) d_\varphi^2(b, c) &= d_\varphi^2(a, b), \\ d_\varphi^2(c, a) &= \beta_1(b) d_\varphi^2(b, c), \\ \beta_1(b) d_\varphi^2(a, b) &= \beta_1(c) d_\varphi^2(c, a). \end{aligned}$$

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,

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

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,

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

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

$$br([c]|\{a, b\})(a) = \frac{d_\varphi^2(a, b) + d_\varphi^2(b, c) - d_\varphi^2(c, a)}{2d_\varphi^2(a, b)} \quad (16)$$

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 $br = br_\varphi$. ■

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,

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

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. ■

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,

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

Consequently, the product of C is equal to

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

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,

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

Hence, the product of C is equal to

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

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,

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

and hence we must show that

$$\beta_1(b)\beta_2(a)\beta_3(c) = \beta_1(c)\beta_2(b)\beta_3(a). \quad (17)$$

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 $\beta' := \frac{3}{2}\beta^* - \frac{1}{2}\beta$. 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

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

But then,

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

implying (17).

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

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

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. ■

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

$$\frac{N_a}{O_b} = \frac{C_a + C_bB_a}{C_b + C_aA_b}. \quad (18)$$

From (17) in the proof of Lemma 4.6, we know that

$$A_dB_aC_b = A_bB_dC_a. \quad (19)$$

Combining (18) and (19) yields

$$\begin{aligned} \frac{N_a}{O_b} &= \frac{C_a + C_bB_a}{C_b + C_aA_b} = \frac{A_bB_d(C_a + C_bB_a)}{A_bB_d(C_b + C_aA_b)} = \frac{A_dB_aC_b + A_bB_dC_bB_a}{A_bB_dC_b + A_dB_aC_bA_b} \\ &= \frac{B_a(A_d + A_bB_d)}{A_b(B_d + A_dB_a)} = \frac{B_a(A_d + A_bB_d)}{A_b(B_d + A_d(1 - B_d))} = \frac{B_a(A_d + A_bB_d)}{A_b(A_d + A_bB_d)} = \frac{B_a}{A_b}, \end{aligned}$$

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 (18) 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 $\lambda_{br}(c, d, a) = \sqrt{N_a/P_c}$. Consequently, $\lambda_{br}(a, d, b)\lambda_{br}(b, d, c)\lambda_{br}(c, d, a) = 1$, which completes the proof. ■

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. ■

References

- [1] Alchourrón, C.E. and D. Makinson (1982), The logic of theory change: Contraction functions and their associated revision function, *Theoria* **48**, 14-37.
- [2] Anscombe, F.J. and R. Aumann (1963), A definition of subjective probability, *Annals of Mathematical Statistics* **34**, 199-205.
- [3] Epstein, L.G. (2005), An axiomatic model of non-Bayesian updating, University of Rochester.
- [4] Epstein, L.G. and M. Schneider (2003), Recursive multiple-priors, *Journal of Economic Theory* **113**, 1-31.
- [5] Gärdenfors, P. (1988), *Knowledge in flux: modeling the dynamics of epistemic states*, Cambridge, MA: MIT Press.
- [6] Ghirardato, P. (2002), Revisiting Savage in a conditional world, *Economic Theory* **20**, 83-92.
- [7] Grove, A. (1988), Two modellings for theory change, *Journal of Philosophical Logic* **17**, 157-170.
- [8] Gul, F. and W. Pesendorfer (2001), Temptation and self-control, *Econometrica* **69**, 1403-1435.
- [9] Hendon, E., Jacobsen, J. and B. Sloth (1996), The one-shot deviation principle for sequential rationality, *Games and Economic Behavior* **12**, 274-282.
- [10] Majumdar, D. (2004), An axiomatic characterization of Bayes' rule, *Mathematical Social Sciences* **47**, 261-273.

- [11] Perea, A. (2002), A note on the one-deviation property in extensive form games, *Games and Economic Behavior* **40**, 322-338.
- [12] Savage, L.J. (1954), *The Foundations of Statistics*, Wiley, New York.
- [13] Schulte, O. (2002), Minimal belief change, Pareto-optimality and logical consequence, *Economic Theory* **19**, 105-144.