Sten Lindström and Wlodzimierz Rabinowicz

ON PROBABILISTIC REPRESENTATION OF NON-PROBABILISTIC BELIEF REVISION

Contents

1. Introduction

In this paper, we study connections between two kinds of modellings of belief states and belief revision, one probabilistic and the other non-probabilistic.[1] In the non-probabilistic modelling, a belief state of an ideally rational agent is represented by a "belief set" — a set of propositions that is consistent and closed under logical consequence. Correspondingly, belief revision is interpreted as an operation that maps belief sets into belief sets. In the probabilistic case, belief states are modelled by probability functions and belief revision by mappings from probability functions to probability functions. In the following, we shall reserve the term "belief revision" for the non-probabilistic modelling. Revision of probability functions will be referred to as "probability revision".

It might seem that a natural way of connecting the two approaches is to think of probability revision as the more fundamental notion and view belief revision as being somehow derived from probability revision. In particular, it is tempting to identify the belief set of an agent with the "top" of his probability function, i.e., with the set of all propositions which are assigned the probability one by that function.

It is not equally clear, however, how to define in a unique way a belief revision operation from a probability revision operation. At first sight, there seems to be no problem here: in order to define the revision $A \oplus x$ of a belief set A with a proposition x, take the probability function associated with A, revise it with x and let $A \oplus x$ be the top of the result. However, which probability function should we choose? In general, for any given A, there are many different probability functions that have A as their top and which therefore may be associated with A. If P and Q are two such probability functions, there is no guarantee that their revisions with x will determine the same belief set.

In the following we are going to discuss five different ways to deal with this Non-Uniqueness Problem, each of which is leading to problems of its own.[2] In order to reach uniqueness, we can do one of the following: (i) impose a suitable condition on probability revision, (ii) let the revised belief set be the common part of the different possible candidates to this title, (iii) introduce an explicit mapping that associates a specific probability function with each belief set, (iv) let belief revision be a relation rather than a function, or, finally, we can (v) introduce the notion of a belief state as a new primitive and define both belief revision and probability revision in terms of a revision operation on belief states.

2. Logical Preliminaries

In both kinds of modellings of belief and belief revision, we assume a set S of propositions as given. In addition to S, we consider an operation Cn that takes sets of propositions to sets of propositions. Cn is assumed to be a *consequence operation*, i.e., it satisfies the following conditions for any sets A and B of propositions:

(1) $A \subseteq Cn(A)$ *(Inclusion)*

(2) $Cn(A) = Cn(Cn(A))$ *(Iteration)*

(3) If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$ *(Monotonicity).*

In addition, we assume:

(4) $Cn(A) = \bigcup\{Cn(B): B \subseteq A \text{ and } B \text{ is finite}\}$ *(Finiteness)*

(5) $\bot \notin Cn(\varnothing)$ *(Consistency)*

(6) $Cn(\{\bot\}) = S$

(7) If $Cn(A \cup \{x \to \bot\}) = S$, then $x \in Cn(A)$

(8) $x \to y \in Cn(A)$ if and only if $y \in Cn(A \cup \{x\})$,

where $\bot$ *(absurdity)* belongs to S, and $\to$ is a binary operation on S *(the classical conditional).* (7) is the classical rule of *Reductio ad Absurdum* and (8) combines the *Deduction Theorem* with a form of *Modus Ponens.* The operations $\neg$, $\wedge$, $\vee$, are defined in terms of $\bot$ and $\to$ in the usual manner. In particular, $\neg x =df (x \to \bot)$.

It follows from the conditions above that:

(9) $A \subseteq Cn(B)$ iff $Cn(A) \subseteq Cn(B)$

(10) $Cn(A \cup B) = Cn(A \cup Cn(B)) = Cn(Cn(A) \cup Cn(B))$

(11) If $x \in Cn(A)$ and $x \to y \in Cn(A)$, then $y \in Cn(A)$

(12) If x is a classical tautology, then $x \in Cn(A)$.

A set A of propositions is *inconsistent* if $Cn(A) = S$; and it is *consistent* otherwise. A proposition x is said to be *consistent (inconsistent)* iff {x} is consistent (inconsistent). We denote the set of all consistent propositions by Con. By a *theory* we understand a set of propositions which is closed under Cn, i.e., A is a theory if $Cn(A) = A$ or equivalently if there exists a B such that $Cn(B) = A$. A *belief set* is a consistent theory, i.e., A is a belief set if $A \neq S$ and $A = Cn(A)$. We let $\mathbb{T}$ and $\mathbb{K}$ be the sets of all theories and belief sets, respectively. That is, $\mathbb{T} = \mathbb{K} \cup \{S\}$.

It is easily seen that the set $\mathbb{K}$ of all belief sets satisfies the following conditions:

(i) $\mathbb{K} \neq \varnothing$;

(ii) If $\mathbb{F} \subseteq \mathbb{K}$ and $\mathbb{F} \neq \varnothing$, then $\bigcap\mathbb{F} \in \mathbb{K}$;

(iii) if $\mathbb{F} \subseteq \mathbb{K}$ and $\mathbb{F}$ is directed, then $\bigcup\mathbb{F} \in \mathbb{K}$. A family $\mathbb{F}$ of subsets of S is *directed* if $\mathbb{F} \neq \varnothing$ and if:

$$A, B \in \mathbb{F} \Rightarrow (\exists C \in \mathbb{F})(A \subseteq C \ \&\ B \subseteq C);$$

(iv) for each $A \in \mathbb{K}$, $\bot \notin A$;

(v) for each $A \in \mathbb{K}$ and each $x, y \in S$, $x \to y \in A$ iff for all $B \in \mathbb{K}$, if $A \subseteq B$ and $x \in B$, then $y \in B$.

(vi) for each $A \in \mathbb{K}$, if there is no $B \in \mathbb{K}$ such that $A \subseteq B$ and $x \to \bot \in B$, then $x \in A$.

The conditions (i) - (vi) give an *intrinsic characterization* of the set $\mathbb{K}$.[3] That is, there is a natural 1-1 correspondence between consequence relations Cn satisfying conditions (1) - (8) and sets of propositions satisfying (i) - (vi). Starting out from the assumption that $\mathbb{K}$ is a family of subsets of S satisfying (i) - (vi), we can define the notions of *consistency* and *logical consequence* as follows: A set A of propositions is consistent iff for some $B \in \mathbb{K}$, $A \subseteq B$. The operation Cn: Pow(S) $\to$ Pow(S), is defined in terms of $\mathbb{K}$ by letting

$$Cn(A) = \bigcap\{B: B \in \mathbb{K} \text{ and } A \subseteq B\},$$

for each $A \subseteq S$. It can then be verified that Cn satisfies conditions (1) - (8) and that $A \in \mathbb{K}$ iff $A \neq S$ and $Cn(A) = A$.

3. Belief Revision and Probability Revision

The simplest way of adding a new proposition x to a belief set A is *expansion*, whereby x is added set-theoretically to A and the result is closed under Cn. Clearly, this way of revising beliefs is possible only if $A \cup \{x\}$ is consistent, since otherwise the result is not a belief set. Therefore we define $A + x$, the *expansion of* A *with* x, to be $Cn(A \cup \{x\})$, if $A \cup \{x\}$ is consistent; and undefined otherwise. That is, we let expansion be a *partial* function from $\mathbb{K} \times$ Con to $\mathbb{K}$.

Belief revision, $\oplus$, is an operation which to every belief set A and every consistent proposition x assigns a belief set $A \oplus x$, where the latter set may be interpreted as A revised by (the addition of) x (as a sole piece of new information).[4] If $A \cup \{x\}$ is consistent, then $A \oplus x$ is just the expansion of A with x. However, $A \oplus x$ differs from $A + x$ in being defined also for consistent propositions x that are inconsistent with A. In the latter case, we may think of $A \oplus x$ as the result of first modifying A to obtain a belief set B which is consistent with x and then expanding B with x. We have the following formal definition:[5]

Definition. A *belief revision operation* is a function $\oplus: \mathbb{K} \times \text{Con} \to \mathbb{K}$ satisfying the following axioms for all belief sets A and all consistent propositions x, y:

($\oplus$1) $x \in A \oplus x$.

($\oplus$2) If $A \cup \{x\}$ is consistent, then $A \oplus x = A + x$.[6]

($\oplus$3) If $Cn(\{x\}) = Cn(\{y\})$, then $A \oplus x = A \oplus y$.

($\oplus$4) If $(A \oplus x) \cup \{y\}$ is consistent, then

$A \oplus (x \wedge y) = (A \oplus x) + y$.

For future reference, we shall speak of axiom ($\oplus$4) as *Revision by Conjunction*.

Let $\mathbb{P}$ be the set of all possible (monadic) *probability functions* on S. Each such function P in $\mathbb{P}$ assigns real numbers between 0 and 1 to the propositions in S and satisfies the standard probability axioms, i.e.,

(P1) $P(T) = 1$, where $T =_{df} \neg\bot$.

(P2) $P(x) \geq 0$.

(P3) If $Cn(\{x,y\}) = S$, then $P(x \vee y) = P(x) + P(y)$.

(P4) If $Cn(\{x\}) = Cn(\{y\})$, then $P(x) = P(y)$.

It can be easily shown that the set of all propositions that are assigned the value 1 by a probability function P is a belief set. We shall refer to this set as t(P) — the *top* of P, or as the belief set *associated* with P.

If $P(x) > 0$, then we can define the *conditionalization* of P by x, $P + x$. The probability function $P + x$ assigns to each proposition y the ratio:

$$(P + x)(y) = \frac{P(x \wedge y)}{P(x)} \cdot$$

Probability revision is an operation $*$ which to every (monadic) probability function P and every consistent proposition x assigns the "revised" probability function $P * x$. Intuitively, if P specifies the "unconditional" subjective probabilities of an agent, then for all x and y, $(P * x)(y)$ is the corresponding probability of y on the condition that x holds. Thus, what we have here is a form of conditionalization, but a "non-standard" one. Unlike P + x, P ∗ x is thought of as being non-trivially defined for consistent propositions x even when P(x) = 0. The formal definition is as follows:[7]

Definition. A *probability revision operation* is a function $*: \mathbb{P} \times \text{Con} \rightarrow \mathbb{P}$ satisfying the following axioms for all probability functions P and all consistent propositions x, y:

(∗1) (P ∗ x)(x) = 1.

(∗2) If P(x) > 0, then P ∗ x = P + x.

(∗3) If Cn({x}) = Cn({y}), then P ∗ x = P ∗ y.

(∗4) If (P ∗ x)(y) > 0, then P ∗ (x ∧ y) = (P ∗ x) + y.

In addition to monadic probability functions, we are also going to consider dyadic ones. One type of dyadic probability functions are the *standard conditional probability functions* obtained from monadic ones by conditionalization, i.e., for each monadic probability function P, there is a unique dyadic function P(—/...) defined by:

$$P(x/y) = \begin{cases} (P + y)(x) & \text{in case } P(y) > 0 \\ \text{undefined otherwise.} \end{cases}$$

Note that for each P(—/...), the corresponding monadic probability function can be recovered by: P(x) = P(x/T), for all x in S. Another kind of dyadic probability functions are the *non-standard conditional probability functions* or *Popper functions.* Formally, we define:[8]

Definition. A *Popper function* is a mapping Π from S × Con into [0,1] such that the following conditions hold whenever the equations are defined:

(Π1) Π(x/x) = 1.

(Π2) For constant y, Π(x/y) is a probability function.

(Π3) If Cn({x}) = Cn({y}), then Π(z/x) = Π(z/y).

(Π4) $\Pi(x \wedge y/z) = \Pi(x/z)\cdot\Pi(y/x \wedge z)$.

A Popper function Π differs from a standard conditional probability function in being defined also for those y ∈ Con for which Π(y/T) = 0. Given any Popper function Π and any x ∈ Con, we can define a monadic probability function P_X by the condition:

$$P_x(y) = \Pi(y/x).$$

Let P be the monadic probability function $P_T = \Pi(—/T)$. We then have:

(1) $P_x(x) = 1$.

(2) If $P(x) > 0$, then $P_x = P + x$.

(3) If $Cn(\{x\}) = Cn(\{y\})$, then $P_x = P_y$.

(4) If $P_x(y) > 0$, then $P_{x \wedge y} = P_x + y$.

Conversely, let P be a fixed (monadic) probability function and assume that there is an operation which assigns a probability function P_x to every consistent proposition x in such a way that conditions (1) - (4) are satisfied. Then, we can define a dyadic probability function Π by the condition:

$$\Pi(y/x) = P_x(y).$$

It is easily seen that Π is a Popper function and that $\Pi(—/T) = P$.

Thus there is a close relationship between Popper functions and probability revision on monadic functions. However, a single Popper function Π only gives us a method of revising one particular monadic probability function, namely Π(—/T), with different consistent propositions x. It does not give us any possibility of making repeated, *iterated* probability revisions.[9] This limitation can be overcome, however, if we assume that each monadic probability function P is correlated with a Popper function s(P) such that s(P)(__/T) = P. In view of conditions (1) - (4) above, we can then define a probability revision operation $*: \mathbb{P} \times Con \rightarrow \mathbb{P}$ as follows:

$$(P * x)(y) = s(P)(y/x).$$

Conversely, given a probability revision operation $*$, we can of course correlate a Popper function s(P) with each monadic probability function via the definition:

$$s(P)(y/x)) = (P * x)(y).$$[10]

Hence, there is a natural one-to-one correspondence between probability revision operations and assignments of Popper functions to the monadic probability functions. In other words, a probability revision operation is essentially an assignment to each monadic probability function P of a Popper function s(P) representing a possible method of revising P in the light of new information.

One might object to our discussion so far that it is not reasonable to identify an agent's set of beliefs with the *top* of his probability function. It is more plausible — one might say — to identify the agent's beliefs with the propositions to which he assigns a *sufficiently high* probability, say, higher than 1 - ε for some suitably small ε.

However, if P is a probability function, then the *ε-top* of P, i.e., the set of all propositions x such that P(x) > 1 - ε, may not be closed under Cn and therefore may not be a belief set. As is well-known, the conjunction of two sufficiently probable propositions may be insufficiently probable.

On the other hand, some philosophers may welcome this lack of logical closure because they have always thought that belief sets, partly due to the agent's lack of logical omniscience, should be seen as something "partial" or "gappy". However, it seems to us that, for a friend of gaps, ε-tops are not gappy enough. While not being "globally" closed under logical consequence, they exhibit at least a "local" closure. Clearly, for each x and y, if x is sufficiently probable and y is a logical consequence of x, then y must also be sufficiently probable. On the other hand, as is well-known, the ε-top of a perfectly consistent probability function may turn out to be inconsistent. For example, the following propositions may all be sufficiently probable:

$$\neg Ga_1, \neg Ga_2, \ldots, \neg Ga_n, Ga_1 \vee Ga_2 \vee, \ldots, \vee Ga_n.$$

This is the well-known *lottery paradox.* Hence, of the following conditions, that jointly characterize our notion of a belief set, ε-tops satisfy the first, while the other two may fail:

(1) if $x \in A$ and $y \in Cn(\{x\})$, then $y \in A$;

(2) if $x, y \in A$, then $x \wedge y \in A$;

(3) A is consistent.

For these reasons, it seems to us that a friend of partial belief sets should instead consider another approach to the problem of probabilistic representation. He can continue to look upon belief sets as tops of probability functions. But, at the same time, he should allow *partial* probability functions. (p is a partial probability function iff, for some (total) probability function P, p is included in P.) Then, partial belief sets may be seen as the tops of partial probability functions. If he, in addition, is willing to allow for *inconsistent* as well as partial belief sets, he should include among the (total) probability functions also the *absurd* probability function which assigns 1 to every proposition. Clearly, one might look upon belief revision as an operation on partial (and/or inconsistent) belief sets.[11] Then it becomes possible to inquire how such a belief revision operation relates to probability revision seen as an operation on partial probability functions. However, in the present paper we shall not pursue this line of inquiry.

4. The Principle of Top Equivalence

Peter Gärdenfors has suggested that we might interpret belief revision in terms of probability revision: if A is the belief set associated with a probability function P, then we might define $A \oplus x$ as the belief set associated with $P * x$. Given this interpretation, he then proves that the axioms for belief revision follow from the axioms for probability revision.[12]

However, Gärdenfors does not mention that this probabilistic representation of belief revision confronts the following difficulty (the Non-Uniqueness Problem): as easily seen, each belief set A is associated with many different probability functions: P, Q, R, etc.; therefore, it does not seem to be possible to determine belief revision uniquely in terms of probability revision — we have no means of determining whether $A \oplus x$ is the belief set associated with $P * x$, or with $Q * x$, or with ....

Let us say that two (monadic) probability functions P and Q are *top-equivalent*, if they have the same top, that is, if their associated belief sets coincide. The Non-Uniqueness Problem arises because different top-equivalent probability functions, P, Q, etc., each with the same A as its top, may give rise to different x-revisions $P * x$, $Q * x$, etc. However, the problem would be solved if such x-revisions of top-equivalent probability functions were themselves required to be top-equivalent. In that case, the identification of $A \oplus x$ with their common top would be unproblematic. Hence, we are led to consider the following *Principle of Top Equivalence* as a possible additional axiom on probability revision:

(TE) For all P, Q $\in \mathbb{P}$, and x $\in$ Con,

$$\text{if } t(P) = t(Q), \text{ then } t(P * x) = t(Q * x).$$

The relation of top equivalence (written: $\approx$) is of course an equivalence relation in the set $\mathbb{P}$. The principle of Top Equivalence says that $\approx$ also is a *congruence relation* in $\mathbb{P}$ with respect to the operation $*$, i.e., that for all P, Q $\in \mathbb{P}$, and x $\in$ Con, if $P \approx Q$, then $P * x \approx Q * x$.

It is easily seen that the Principle of Top Equivalence is not a logical consequence of the axioms ($*$1) - ($*$4) for probability revision. To prove this, associate with each probability function P a well-ordering $\mathbf{O}_P$ of all probability functions, with P as its first member. (Here we use the Axiom of Choice twice: first to prove the existence of the relevant well-orderings and then to associate a particular well-ordering with each P.) For each P and each consistent proposition x, let P^x be the first probability function in $\mathbf{O}_P$ that gives positive probability to x. We then define $P * x$ by letting $P * x = P^x + x$. It is then possible to prove that the operation $*$ so defined satisfies ($*$1) - ($*$4). The only troublesome axiom is:

($*$4) If $(P * x)(y) > 0$, then $P * (x \wedge y) = (P * x) + y$.

To prove this, we assume that $(P * x)(y) > 0$. We observe that the following conditions then obtain:

(i) If $P^x(x \wedge y) > 0$, then $P^x + (x \wedge y) = (P^x + x) + y$;

(ii) $P^x(x \wedge y) = P^x(x) \times P^x(y/x) = P^x(x) \times (P * x)(y) > 0$;

(iii) $P^x = P^{(x \wedge y)}$.

Hence,

$$(P * x) + y = (P^x + x) + y = P^x + (x \wedge y) = P^{(x \wedge y)} + (x \wedge y) = P * (x \wedge y).$$

Assume now that P and Q are top-equivalent probability functions and let x and y be propositions such that $P(x) = Q(x) = 0$ and neither y nor $\neg y$ belongs to $Cn(\{x\})$. We can always choose the orderings $\mathbf{O}_P$ and $\mathbf{O}_Q$ in such a way that the first probability function in $\mathbf{O}_P$ that assigns positive probability to x, assigns probability 1 to y, while the corresponding first function in $\mathbf{O}_Q$ assigns to y a probability <1. Then $(P * x)(y) = 1$, while $(Q * x)(y) < 1$. In this way we can make sure that the Top Equivalence principle fails.

It should be noted that Top Equivalence is an essentially different kind of condition from the standard axioms on probability revision. The latter principles are "local" — they consider revisions of one probability function at a time. Each of them is of the form: "For any probability function P, ... ". Top Equivalence, on the other hand, is a "global" condition on probability revision. As is clear from its form, "For all probability functions P and Q, ... ", it is a condition on the relationship between revisions of *different* probability functions.

Now we might ask whether there are any good reasons to accept the Principle of Top Equivalence. And what about reasons against it? Notice that Top Equivalence follows from the following *Principle of Monotonicity:*

(M) For all $P, Q \in \mathbb{P}$, and $x \in$ Con, if $t(P) \subseteq t(Q)$, then $t(P * x) \subseteq t(Q * x)$.

However, as Gärdenfors has proved, this principle cannot, "on pain of triviality", be added to (∗1) - (∗4) as a new axiom.[13] More precisely, the principle (M) is inconsistent with (∗1) and the following consequence of (∗2):

(P) If $P(x) > 0$, then $t(P) \subseteq t(P * x)$ *(Preservation),*

together with the assumption of *non-triviality*, namely that Con contains at least three pairwise incompatible propositions.

One might suspect that the Principle of Top Equivalence will give rise to a similar result. In fact, however, no such triviality result is forthcoming. In personal communication, Gärdenfors has suggested how to prove this claim. The following argument is a modification of his proposal. Consider any well-ordering **O** of the probability functions in $\mathbb{P}$. Now, define a probability revision operation ∗ as follows: for any P and $x \in$ Con,

$$P * x = \begin{cases} P + x \text{ in case } P(x) > 0 \\ \mathbf{O}^x + x \text{ if } P(x) = 0. \end{cases}$$

where $\mathbf{O}^x$ is the first probability function Q in the well-ordering **O** such that Q(x) > 0. It is then straightforward to verify that ∗ satisfies the axioms (∗1) - (∗4). (For (∗4), see the corresponding proof in the construction above.) Also note that this definition immediately implies that ∗ satisfies Top Equivalence: If P and Q are top-equivalent and P(x) > 0, then the top-equivalence of P ∗ x and Q ∗ x follows from the properties of conditionalization. And if P(x) = 0, then P ∗ x simply coincides with Q ∗ x given the present definition of ∗. Since the definition does not depend on any special assumptions concerning the consequence operation Cn or the cardinality of S, adding Top Equivalence as a further axiom on probability revision does not lead to any triviality result.

Consider now the following principle of *Restricted Monotonicity:*

$$\text{(RM)} \qquad \frac{t(P) \subseteq t(Q) \text{ and either } P(x) = 0 \text{ or } Q(x) > 0}{t(P * x) \subseteq t(Q * x)} \; .$$

That is, if the top of P is included in the top of Q *and* it is not the case that both P(x) > 0 and Q(x) = 0, then the top of P ∗ x is included in the top of Q ∗ x.[14] The analogue for belief revision of this principle would be:

$$\text{(RMB)} \qquad \frac{A \subseteq B \text{ and either } \neg x \in A \text{ or } \neg x \notin B}{A \oplus x \subseteq B \oplus x} \; ,$$

i.e., if A is included in B *and* it is not the case that x is consistent with A and inconsistent with B, then A ⊕ x is included in B ⊕ x.[15]

Restricted Monotonicity, although weaker than Monotonicity, still entails Top Equivalence. Furthermore, Gärdenfors' proof that the addition of Monotonicity leads to triviality does not extend to Restricted Monotonicity. In fact, it is easy to prove that the construction of ∗ described above validates not only Top Equivalence but also Restricted Monotonicity: Namely, assume that t(P) ⊆ t(Q) and that either P(x) = 0 or Q(x) > 0. Then there are two possibilities: (i) P(x) = 0 and Q(x) = 0; and (ii) P(x) > 0 and Q(x) > 0. In case (ii), the definition of ∗ yields: P ∗ x = Q ∗ x. In case (ii), t(P ∗ x) ⊆ t(Q ∗ x) follows by the properties of conditionalization.

Although the Principle of Restricted Monotonicity — unlike the unrestricted principle (M) — does not lead to triviality, it seems unacceptable on intuitive grounds. Thus, to give an example, let x, y, and z be the propositions: **Tweety is a penguin**, **Tweety can fly**, and **Penguins cannot fly**, respectively. Suppose that P describes the probabilities of someone who is certain of $\neg$x and of y, but who lacks any opinion about z, i.e., P(x) = 0, P(y) = 1 and P(z) = 0.5. It seems that if such a person were to learn x, he would still keep y among the propositions he is certain of, i.e., (P $*$ x)(y) = 1. By the axioms of probability revision, P $*$ z = P + z. Hence, by the definition of conditionalization, t(P) $\subseteq$ t(P $*$ z). It follows that (P $*$ z)(x) = 0. From P(x) = (P $*$ z)(x) = 0 and t(P) $\subseteq$ t(P $*$ z), it follows by Restricted Monotonicity that t(P $*$ x) $\subseteq$ t((P $*$ z) $*$ x). Hence, if Restricted Monotonicity were a valid principle, the proposition y, being a member of t(P $*$ x), would also belong to t((P $*$ z) $*$ x). But in fact, the opposite seems to be the case: someone who originally believes that Tweety is not a penguin ($\neg$x), that Tweety can fly (y), and that penguins cannot fly (z), might very well give up his belief in Tweety's ability to fly upon learning that Tweety in fact *is* a penguin.

Now, one might ask: Is it possible to construct a comparable *intuitive* counterexample to the principle of Top Equivalence? David Makinson (personal communication) has suggested the following example:

> "Consider two detectives, both very clever and eminently rational, one of whom has a penchant for the heuristic "cherchez la femme", and the other for the heuristic "look for the one who made a profit". Suppose that in a particular case, at time t, there is not much that either is certain of, and on that little the two detectives entirely agree. But suppose that at time t', one of those certainties is overthrown. Who knows whether our two detectives will be certain about exactly the same things now?"

The idea seems to be that two persons who are certain of exactly the same things (and, what is important, who presumably have the same principles for probability revision) may still differ very dramatically in their non-extreme probability assignments: one is nearly certain that "there is a woman behind it", while the other is nearly sure that it is instead the profit motive that is the explanation. In such a situation, it seems possible, Makinson argues, that they might reach different certainties after being confronted with new information that contradicts what they have previously been certain of.

Consider a somewhat different example. Suppose that there are two persons, Castor and Pollux, with the same policies for belief revision. Assume that their belief sets coincide and contain, in particular, two logically independent propositions x and y. Assume also that there are propositions r and s that strongly support x and y, respectively. Castor assigns high probability to r and low to s, while Pollux does the opposite. Thus, while they have the same set of certainties, they drastically differ with respect to some of their non-extreme probability assignments. Suppose now that each of them receives the same proposition, $\neg(x \wedge y)$, as the sole piece of new information. Each of them, therefore, has to remove x or y, or both of these propositions, from the original belief set. Assuming that they want to keep as much as possible of their initial beliefs, it seems quite possible that Castor, whose initial probability for the x-supporting proposition r was high, will continue to believe x, while Pollux, for the analogous reasons, will keep to y. If this is possible, and it seems to be, then we have a counterexample to Top Equivalence. Hence, this appealing and elegant principle must probably be abandoned.

5. The Common Part Approach

If we do not want to assume Top Equivalence, we might try to explore another way of dealing with the Non-Uniqueness Problem: at least *prima facie*, it seems possible to interpret the revised belief set as the *common top part* of the different revised probability functions. To put it more precisely, we might define, for all $A \in \mathbb{K}$ and $x \in$ Con,

$$A \oplus x = \bigcap\{B \in \mathbb{K}: (\exists P \in \mathbb{P})(A = t(P) \text{ and } B = t(P * x))\}.$$

That is, we let $A \oplus x$ be the intersection — the common part — of all the belief sets B such that, for some probability function P, A is the top of P and B is the top of $P * x$. Thus, instead of trying to choose between different probability functions with which A is associated, we just take what is common to all of them.

Why the intersection and not the union? There are obvious intuitive reasons for this manoeuvre — belief revision by intersection is much more cautious than revision by unions. But the formal reasons are even stronger: The intersection of a set of belief sets is consistent and logically closed and therefore is itself a belief set. The union, on the other hand, may be neither consistent nor logically closed.

However, this definition of belief revision, while in many respects attractive, is not quite satisfactory. To take the intuitive criticism first, the "common part"-definition breaks the basic connection between the belief sets and the probabilities which is fundamental to the whole probabilistic interpretation of the belief set model. Intuitively speaking, according to this "Bayesian" interpretation, a person's belief set is the set of all propositions to which he assigns probability 1. Now, let us consider several persons who have the same belief set A and the same principles for probability revision, but whose probability assignments otherwise may differ. Suppose now that each of them receives an information x and revises his probabilities accordingly. Then, *insofar* as we accept the "common part"-definition of $A \oplus x$, it may well happen that each of these persons will have some beliefs that do not belong to $A \oplus x$. Thus, while A constitutes the prior belief set for each of these persons, $A \oplus x$ might not be the total posterior belief set for anyone of them.

In fact, given the "common part"-definition, there might not exist *any* probability function P with A as its top, for which $A \oplus x = t(P * x)$. To believe that such a function will always exist is to accept the following *intersection principle:*[16]

(IP) $\forall A \in \mathbb{K}, \forall x \in \text{Con}, \exists Q \in \mathbb{P},$

$$t(Q) = A \;\&\; t(Q * x) = \bigcap\{t(P * x)): t(P) = A\}.$$

To put it somewhat differently, for every consistent x, there exists a probability function Q such that $t(Q) = A$ and, for every P, if $t(P) = A$, then $t(Q * x)$ is included in $t(P * x)$. (IP) does not follow from the standard axioms for probability revision. (Note that (IP), just like Top Equivalence , is a "global" condition, while all the standard axioms are "local".) On the other hand, (IP) would follow from Top Equivalence but not vice versa.

Moving now to a formal criticism, it is easy to see that the "common part"-definition validates all of the Gärdenfors axioms on belief revision with one exception: Revision by Conjunction, the principle according to which the consistent expansion of $A \oplus x$ with y equals $A \oplus (x \wedge y)$, fails to be valid. Here follows a very rough explanation of this failure.

Suppose that propositions y and z are mutually consistent but that each of them is inconsistent with a belief set A. Now, assume that, for all P which have A as their top, $P * (y \vee z)$ assigns 1 to z and that some of these $(y \vee z)$-revised functions assign positive probability to y. Finally, assume that there are functions P with A as their top such that their revisions by y do not assign 1 to z.

Given these assumptions and our present "common part"-definition of belief revision, it immediately follows that y is consistent with $A \oplus (y \vee z)$ and that $y \wedge z$ belongs to $(A \oplus (y \vee z)) + y$ but not to $A \oplus ((y \vee z) \wedge y)$, (note that the latter set equals $A \oplus y$). Thus, letting x be the disjunction $y \vee z$, we get here a counterexample to Revision by Conjunction.

Note that this axiom for belief revision fails to be valid even though its analogue for probability revision,

(∗4) If $(P * x)(y) > 0$, then $P * (x \wedge y) = (P * x) + y$,

has been assumed to hold. Thus, the "common part"-definition introduces an unattractive dissimilarity between the two kinds of revision.

If we would like to validate Revision by Conjunction, we could do this by adding to the standard axioms on probability revision the following *Strong Intersection Principle:*

(SIP) $\forall A \in \mathbb{K}, \exists Q \in \mathbb{P}, \forall x \in \text{Con},$

$$t(Q) = A \ \& \ t(Q * x) = \bigcap\{t(P * x)): t(P) = A\}.$$

Note that this principle differs from (IP) only with respect to the order of the quantifiers $\exists Q$ and $\forall x$. According to (SIP), there exists a probability function Q such that, for every consistent x, $t(Q) = A$ and, for every P, if $t(P) = A$, then $t(Q * x)$ is included in $t(P * x)$.

To prove Revision by Conjunction, note first that (SIP) together with the "common part"- definition of $\oplus$ imply the following: for every A, there is some Q such that, for all x and y,

(1) $A \oplus x = t(Q * x)$ and $A \oplus (x \wedge y) = t(Q * (x \wedge y))$.

From (∗4), and assuming that $(A \oplus x) \cup \{y\}$ is consistent, i.e., that $(Q * x)(y) > 0$, we have

(2) $(Q * x) + y = Q * (x \wedge y)$.

Hence,

(3) $t((Q * x) + y) = t(Q * (x \wedge y))$.

From ordinary probability theory, we get:

(4) $t((Q * x) + y) = t(Q * x) + y$.

From (1), (3) and (4) we finally obtain,

(5) $(A \oplus x) + y = A \oplus (x \wedge y)$,

provided that $(A \oplus x) \cup \{y\}$ is consistent. Q. E. D.

It is easy to see that (SIP), just like (IP), follows from Top Equivalence but not vice versa. The intuitive idea behind this principle seems to be as follows: The probability revision operation $*$ is such that, for every belief set, there is a probability function with this belief set as its top which is, so to say, maximally "cautious" insofar as its behavior under revisions is concerned: whatever proposition we revise it with, the resulting probability function gives rise to a minimal belief set. Thus, (SIP) is a kind of cautiousness condition on $*$.

An even more cautious revision operation would always allow genuinely minimal belief sets after revision, i.e., belief sets that include nothing but the new information. We could express this idea as the following *Principle of Caution:*

(PC) $\forall A \in \mathbb{K}, \exists Q \in \mathbb{P}$ such that $t(Q) = A$ and $\forall x \in$ Con,

if x is inconsistent with A, then $t(Q * x) = Cn(\{x\})$.[17]

Clearly, (PC) is stronger than (SIP) and it neither entails nor is entailed by Top Equivalence.

If we add (PC) or (SIP) to the standard axioms on probability revision, the "common part"-definition of $\oplus$ will solve the Non-Uniqueness Problem. However, the intuitive criticism of this definition still remains unanswered and, in addition, neither (PC) nor (SIP) seems to be especially plausible.

6. Base Functions

When confronted with the Non-Uniqueness Problem, we might simply bite the bullet and admit that the probabilistic interpretation of belief revision presupposes a *choice:* we have to choose a particular mapping b from the set of belief sets $\mathbb{K}$ to the set of probability functions $\mathbb{P}$, and then define belief revision in terms of the probability revision operation $*$ *together with* the function b. Thus, by a *base function,* we shall understand a function b from $\mathbb{K}$ into $\mathbb{P}$ such that for every A in $\mathbb{K}$, b(A) is one of the probability functions that have A as their top. The existence of base functions is ensured by the Axiom of Choice.

Now, in terms of a base function b, it becomes easy to define belief revision from probability revision: to revise A by a consistent proposition x, go first to the probability function b(A), revise this function by x, and then, finally, go to the top of this revision. In other words, given a base function b, we define $\oplus$ to be the unique function from $\mathbb{K} \times$ Con to $\mathbb{K}$ satisfying for every A in $\mathbb{K}$ and x in Con the condition:

(DEF) $$(A \oplus x) = t(b(A) * x).$$

It is easy to check that, whatever base function b has been chosen, the so defined operation $\oplus$ satisfies the Gärdenfors axioms ($\oplus$1) - ($\oplus$4) for belief revision.

Note that a similar approach may be used for defining a particular probability revision operation on monadic probability functions, starting out from the set of (dyadic) Popper functions instead. As we have done in section 3 above, we may choose an arbitrary mapping, s, from monadic probability functions to the dyadic ones. As we know, a base function assigns to each belief set a probability function that has this belief set as its top. Analogously, we demand from s that it should assign to each monadic P a dyadic Π that has P as its monadic part:

$$\text{for all } x,\ \Pi(x/T) = P(x).$$

Now it becomes easy to define a probability revision operation on monadic probability functions:

for all P and x, and for every y,

$$(P * x)(y) = s(P)(y/x).$$

Then, by using two arbitrary embeddings, s and b, we can first move from the set of dyadic Popper functions to a specific revision operation on the set of monadic probability functions and then, from that revision operation, to a revision operation on belief sets.

Even though defining $\oplus$ in terms of $*$ and a particular base function validates all the standard axioms on $\oplus$, certain base functions still seem to be *inappropriate* as belief revision generators. They are just too "arbitrary" in their assignments and this arbitrariness shows up under *iterated* revisions. Thus, consider what happens if b is our chosen base function and we apply (DEF) in order to revise a given A by a sequence $x_1,..., x_n$ of consistent propositions, one by one. One would want the result of this process to coincide with the top of the probability function obtained by stepwise revisions of the probability function b(A) with the same sequence $x_1,..., x_n$:

(T) $$(...((b(A) * x_1) * x_2)...) * x_n \approx b((...((A \oplus x_1) \oplus x_2)...) \oplus x_n),$$

for all $A \in \mathbb{K}$ and $x_1,..., x_n \in$ Con. ($\approx$ is the relation of top equivalence.)

Metaphorically, one might express this adequacy condition on base functions as follows: iterated belief revision should keep track of iterated probability revision. If a base function b satisfies this *Tracking Condition,* then we shall say that b is *appropriate.* It is easy to check that, in the absence of Top Equivalence, some base functions may not be appropriate.[18]

As a matter of fact, there is also a stronger but simpler adequacy condition that we might want to impose on eligible base functions: we might want to use a base function that *commutes* with revision. That is, we might want it to be the case that, for all $A \in \mathbb{K}$ and $x \in$ Con,

(C) $$b(A \oplus x) = b(A) * x.$$

If b satisfies this *Commutativity Condition,* then we shall say that b is *perfectly appropriate.*

We have the following formal definitions, where $\oplus$ is a belief revision operation and $*$ a probability revision operation.

Definition.

(a) A *base function* is a function $b: \mathbb{K} \rightarrow \mathbb{P}$ such that, for every $A \in \mathbb{K}$, $t(b(A)) = A$.

(b) A base function b is said to be a base function *for* $\oplus$ *relative to* $*$, iff for all $A \in \mathbb{K}$ and $x \in$ Con,

(DEF) $$A \oplus x = t(b(A) * x).$$

That is, iff

$$b(A) * x \approx b(A \oplus x).$$

(c) A base function b is said to be *appropriate* for $\oplus$ relative to $*$, iff for all $A \in \mathbb{K}$ and $x_1,..., x_n \in$ Con,

(T) $$(...((b(A) * x_1) * x_2)...) * x_n \approx b((...((A \oplus x_1) \oplus x_2)...) \oplus x_n).$$

(d) A base function b is said to be *perfectly appropriate* for $\oplus$ relative to $*$, iff for all $A \in \mathbb{K}$, $x \in$ Con,

(C) $$b(A \oplus x) = b(A) * x.$$

(e) b is an *appropriate (perfectly appropriate)* base function relative to $*$ iff, relative to $*$, b is appropriate (perfectly appropriate) for the operation $\oplus$ which is defined from b and $*$ via (DEF).

Of course, every appropriate base function with respect $\oplus$ and $*$ is a base function for $\oplus$ relative to $*$. We also have the following:

Lemma 1.

(a) If $*$ satisfies Top Equivalence, then every base function is appropriate relative to $*$.

(b) If b is perfectly appropriate for $\oplus$ relative to $*$, then b is appropriate for $\oplus$ relative to $*$. Moreover, for all $A \in \mathbb{K}$ and $x_1,..., x_n \in$ Con,

$$(...((b(A) * x_1) * x_2)...) * x_n) = b(...((A \oplus x_1) \oplus x_2)...) \oplus x_n).$$

There are several questions concerning appropriate and perfectly appropriate base functions that we have not been able to answer:

(i) Let $*$ be any probability revision operation. Does it follow that there exists a base function b such that, relative to $*$, b is appropriate (perfectly appropriate) for the belief revision operation $\oplus$ defined via (DEF) in terms of $*$ and b? In other words, can we construct belief revision from probability revision in an appropriate (perfectly appropriate) way starting from any $*$ whatsoever?

The claim:

There exists an appropriate (perfectly appropriate) base function relative to $*$,

is essentially a claim about the probability revision operation only. It can be reformulated in such a way that it does not even implicitly involve the concept of belief revision. Thus, consider the following definitions:

A set $\mathbb{F}$ of probability functions is an *appropriate family* of such functions (with respect to a given $*$) iff

(a) every belief set is the top of some member of $\mathbb{F}$;

(b) $\mathbb{F}$ is closed under $*$, i.e., for every P in $\mathbb{F}$ and for every x, $P * x \in \mathbb{F}$; and

(c) $\mathbb{F}$ satisfies Top Equivalence, i.e., for every P and Q in $\mathbb{F}$ and for every x, if $P \approx Q$, then $P * x \approx Q * x$.

An appropriate family $\mathbb{F}$ shall be said to be *perfectly appropriate* iff every belief set is the top of *exactly* one member of $\mathbb{F}$. (Note that, given this strengthening of (a), clause (c) becomes trivially true and therefore redundant.)

Now it can be shown that the following equivalences obtain:

Lemma 2. There exists an appropriate (perfectly appropriate) base function relative to $*$ iff $\mathbb{P}$ includes an appropriate (perfectly appropriate) family relative to $*$. (This family is simply the range of the base function in question closed under $*$.)

Thus, when discussing the existence of appropriate (perfectly appropriate) base functions, one may exclusively concentrate on the probability revision operation.

(ii) If question (i) is answered in the negative, one can still ask whether there are any simple and natural conditions that $*$ should satisfy if the existence of appropriate (perfectly appropriate) base functions is to be guaranteed.

(iii) If with respect to a given $*$, there exist (perfectly) appropriate base functions, does it follow that there also exists such a function for every P in $\mathbb{P}$? That is, does it follow that, for every P, there exists an appropriate (perfectly appropriate) base function b such that b assigns P to the belief set associated with P? And if not, what further conditions on $*$ are needed to guarantee this result?

All of these questions have probability revision as their starting point. There is another series of questions starting from *belief* revision instead and concerning the possibility of its probabilistic representation:

(0) Let $\oplus$ be any belief revision operation. Does there always exist some probability revision operation $*$ and a function b: $\mathbb{K} \to \mathbb{P}$ such that b is a base function for $\oplus$ relative to $*$? It should be noted that if this question has an affirmative answer, then every condition on belief revision $\oplus$ that is a consequence of the axioms on $*$ together with (DEF) already follows from the belief revision axioms ($\oplus$1) - ($\oplus$4). This would mean then that the latter axiom set is *complete* with respect to the probabilistic interpretation.

(1) Consider any $\oplus$ such that, for some b and $*$, b is a base function for $\oplus$ relative to $*$. Does there exist some probability revision operation $*$ and a base function b such that b is appropriate (perfectly appropriate) for $\oplus$ relative to $*$? In other words, is $\oplus$ representable by a probability revision operation in a (perfectly) appropriate way?

We can report a partial answer to these questions. Let $\mathbb{K}_f$ be the set of all finitary belief sets. (A is *finitary* iff, for some finite set X of propositions, A = Cn(X).) Let $\mathbb{P}_f$ be the set of probability functions P with finitary tops. We say that $\oplus$ is a *finitary belief revision operation*, if $\oplus$ is a function from $\mathbb{K}_f \times$ Con into $\mathbb{K}_f$ that satisfies ($\oplus$1) - ($\oplus$4). Similarly, we say that $*$ is a *finitary probability revision operation* iff $*$: $\mathbb{P}_f \times$ Con $\to \mathbb{P}_f$ and $*$ satisfies ($*$1) - ($*$4). The following notions: base function, a base function being (perfectly) appropriate, Top Equivalence, etc., are in an obvious way applicable to finitary belief revision and finitary probability revision. We then have the following *representation theorem* for finitary belief revision in terms of finitary probability revision.

Theorem. Let $\oplus$ be any finitary belief revision operation. Then, there exists a function $*$: $\mathbb{P}_f \times$ Con $\to \mathbb{P}_f$ and a function b: $\mathbb{K}_f \to \mathbb{P}_f$ such that $*$ is a finitary probability revision operation and b is a perfectly appropriate base function for $\oplus$ relative to $*$. Moreover, the $*$ in question satisfies the Principle of Top Equivalence.

Proof. Let $\oplus$ be any belief revision operation on the set $\mathbb{K}_f$ and let Q be any "regular" probability function. (Q is *regular* iff it assigns probability 1 only to logical truths.) For any $A \in \mathbb{K}_f$, let b(A) be Q + &A, where &A is the conjunction of some finite axiom set for A. (The choice of the axiom set is immaterial, since any two axiom sets for A are logically equivalent and therefore interchangeable in probabilistic contexts.) Clearly, since Q is regular, t(b(A)) = A. Thus, b is a base function. Now, we first prove the following:

Claim: If x is consistent with A, then b(A + x) = b(A) + x.

Proof of the claim: As we know, b(A) = Q + &A. We also have: A + x = Cn({&A $\wedge$ x}) and Q(&A $\wedge$ x) > 0. Thus:

$$b(A + x)(y) = \frac{Q((\&A \wedge x) \wedge y)}{Q(\&A \wedge x)} =$$

$$= \frac{Q(\&A \wedge (x \wedge y))}{Q(\&A)} \times \frac{Q(\&A)}{Q(\&A \wedge x)} = \frac{Q(x \wedge y/\&A)}{Q(x/\&A)} =$$

$$= \frac{b(A)(x \wedge y)}{b(A)(x)} = b(A)(y/x) = (b(A) + x)(y).$$

Next, we define $*$ as follows: For any P in $\mathbb{P}_f$ and any x $\in$ Con, (i) if P(x) > 0, P $*$ x = P + x; (ii) if P(x) = 0, then P $*$ x = b(t(P) $\oplus$ x). We must prove that $*$, so defined, satisfies the axioms ($*$1) - ($*$4). As usual, the only troublesome axiom is ($*$4). To prove this we assume that (P $*$ x)(y) > 0. We consider two cases:

Case 1: P(x) > 0. Then,

$$P(x \wedge y) = P(x) \times P(y/x) = P(x) \times (P * x)(y) > 0.$$

Hence,

$$P * (x \wedge y) = P + (x \wedge y) = (P + x) + y = (P * x) + y.$$

Case 2: P(x) = 0. Then, P $*$ (x $\wedge$ y) = b(t(P) $\oplus$ (x $\wedge$ y)). But (P $*$ x)(y) > 0, implies that y is consistent with t(P) $\oplus$ x. Hence,

$$t(P) \oplus (x \wedge y) = (t(P) \oplus x) + y \qquad \text{(axiom } (\oplus 4)).$$

Hence,

$$P * (x \wedge y) = b(t(P) \oplus (x \wedge y)) = b(t(P) \oplus x) + y) = b(t(P) \oplus x) + y = (P * x) + y.$$

Next, we prove that $*$ satisfies Top Equivalence. Assume that t(P) = t(Q). Again, we consider two cases:

Case 1: P(x) > 0. From this, together with t(P) = t(Q), it follows that also Q(x) > 0. Hence, we have both $P * x = P + x$ and $Q * x = Q + x$. However, $(P + x)(y) = 1$ iff $P(x \wedge y)/ P(x) = 1$ iff $P(x \rightarrow y) = 1$ iff $Q(x \rightarrow y) = 1$ iff $Q(x \wedge y)/Q(x) = 1$ iff $(Q + x)(y) = 1$. Thus, t(P) = t(Q) implies $t(P * x) = t(Q * x)$, in this case.

Case 2: P(x) = 0. Then, $P * x = b(t(P) \oplus x) = b(t(Q) \oplus x) = Q * x$. In consequence, $t(P * x) = t(Q * x)$.

Finally, we prove that b is perfectly appropriate, i.e., for all $A \in \mathbb{K}_f$ and $x \in$ Con:

$$b(A) * x = b(A \oplus x).$$

For the case when b(A)(x) > 0, this follows from the claim we proved above. Thus, let b(A)(x) = 0. Then, by the definition of $*$, $b(A) * x = b(t(b(A)) \oplus x) = b(A \oplus x)$.

Q. E. D.

(2) If, this partial result notwithstanding, question (1) is to be answered in the negative, what additional conditions should $\oplus$ satisfy in order to have a (perfectly) appropriate probabilistic representation?

(3) If a given $\oplus$ has a (perfectly) appropriate probabilistic representation $*$, does it follow that there also exists such a representation for every P in $\mathbb{P}$? That is, does it follow that, for every P, there exists a probability revision operation $*$ and a base function b such that b is appropriate (perfectly appropriate) for $\oplus$ relative to $*$ and b(t(P)) = P? And, if not, what further conditions on $\oplus$ are needed in order to guarantee this result?

Finally, there is a question of intuitive interpretation: Suppose we choose a particular base function b and then use (DEF) to define belief revision in terms of b and probability revision. What does this choice of a base function amount to in *intuitive terms?* Admittedly, this is a rather vague question, but even a vague answer would be welcome.

7. Belief Revision as a Relation

Belief revision, like all binary operations, may of course be viewed as a ternary relation. This way of looking at belief revision is natural if we think that the agent's policies for belief change may not always yield a *unique* belief set as the result of revising a given belief set A with a proposition x. Hence, we define:

Definition.

(a) A *belief revision relation* is a ternary relation $\mathbf{R} \subseteq \mathbb{K} \times \mathrm{Con} \times \mathbb{K}$ satisfying the following axioms for all belief sets A, B, C and all consistent propositions x, y:

(**R**0) $(\exists D \in \mathbb{K}) A\ \mathbf{R}_x\ D$. *(Seriality)*

(**R**1) If $A\ \mathbf{R}_x\ B$, then $x \in B$.

(**R**2) If $A \cup \{x\}$ is consistent and $A\ \mathbf{R}_x\ B$, then $B = A + x$.

(**R**3) If $Cn(\{x\}) = Cn(\{y\})$ and $A\ \mathbf{R}_x\ B$, then $A\ \mathbf{R}_y\ B$.

(**R**4) If $A\ \mathbf{R}_x\ B$, $B\ \mathbf{R}_y\ C$ and $B \cup \{y\}$ is consistent, then $A\ \mathbf{R}_{(x \wedge y)}\ C$.

(b) A belief revision relation **R** is said to be *functional* if, in addition to (**R**0) - (**R**4), it satisfies:

(**R**5) If $A\ \mathbf{R}_x\ B$ and $A\ \mathbf{R}_x\ C$, then $B = C$.

The intuitive reading of $A\ \mathbf{R}_x\ B$ is: B is a (possible) result (for a given agent) of revising A by (the addition of) x (as a sole piece of new information). The seriality axiom (**R**0) corresponds to the requirement that belief revision should be defined for all belief sets A and consistent propositions x. Axioms (**R**1) - (**R**4) are the relational counterparts to the axioms (⊕1) - (⊕4), respectively. There is a natural one-to-one correspondence between belief revision operations ⊕ and belief revision relations **R** that are functional:

Lemma. Every belief revision operation ⊕ determines a functional belief revision operation **R** via the condition:

$$A\ \mathbf{R}_x\ B \text{ iff } A \oplus x = B.$$

Conversely, given a functional belief revision relation **R**, the same condition determines a unique belief revision operation ⊕.

Consider now what happens to the Non-Uniqueness Problem when we adopt the relational approach to belief revision. There is no longer any problem of defining belief revision from probability revision. The natural definition is:

(D) A $\mathbf{R}_x$ B iff, for some P in $\mathbb{P}$, A = t(P) and B = t(P ∗ x).

The intuitive idea behind (D) could be expressed as follows. Suppose that we know only the agent's policy ∗ for revising probabilities and his belief set A (but not his entire probability function P). $\mathbf{R}_x$(A), i.e., {B: A $\mathbf{R}_x$ B}, is then the set of all conceivable alternative results of the revision of A with x — conceivable, that is, given what we know about the agent in question.

Given this definition, the relation **R** satisfies the axioms (**R**0) - (**R**4). The Non-Uniqueness Problem disappears. What we have instead is simply the failure of the so defined belief revision relation to be functional. The functionality principle (**R**5) will be satisfied iff the probability revision operation satisfies Top Equivalence.[19]

On the other hand, in the absence of Top Equivalence, and, consequently, in the absence of the functionality principle, the relational approach sketched above is perhaps too weak in its expressive power. Let us explain.

Consider the relation of *iterated* belief revision, which we define as follows: B is a revision of A with $x_1,..., x_n$ (in that order) — in symbols, A $\mathbf{R}_{x_1,..., x_n}$ B — iff B is the top of some probability function P such that, if you first revise P with x_1, and then revise P ∗ x_1 with x_2, and so on until you come to x_n, you will thereby reach a probability function with B as its top. To put it more formally,

(D_{it}) A $\mathbf{R}_{x_1,..., x_n}$ B iff $(\exists P \in \mathbb{P})(A = t(P)\ \&\ B = t((...(P * x_1) * ...) * x_n))$.

Clearly, the relation of simple belief revision, $\mathbf{R}_x$, is just a special case of the iterated belief revision relation. At the same time, in the absence of Top Equivalence, the latter relation is not definable in terms of the former. In particular, appearances notwithstanding, the relation of iterated belief revision, $\mathbf{R}_{x_1,..., x_n}$, is stronger than the relative product of a series of simple belief revision transformations, $\mathbf{R}_{x_1}/\ldots/\mathbf{R}_{x_n}$. To illustrate, suppose that A $\mathbf{R}_{x_1}/\mathbf{R}_{x_2}$ B, i.e., that for some P and P', A = t(P), t(P $*$ x_1) = t(P') and B = t(P' $*$ x_2). Then there is no guarantee that there exists a single P for which A = t(P) and B = t((P $*$ x_1) $*$ x_2). That is, there is no guarantee that A $\mathbf{R}_{x_1,\ x_2}$ B. In other words, it may be impossible for an agent to reach B from A by a two-step revision of his probability function.

This suggests that a relational approach to belief revision may turn out to be more complex than one would have expected. If we do not want to lose expressive power, then, instead of starting out from the relation of simple belief revision, we may have to work with iterated belief revision as our fundamental relation.[20]

8. Belief States as a New Primitive

Another approach to the Non-Uniqueness Problem is to view it as a symptom that something is wrong with the very idea of belief revision as an operation on belief *sets*. Intuitively, belief revision is an operation on belief *states* and perhaps the source of the problem lies in the identification of belief states with sets. In particular, if we adopt a probabilistic view of belief, then it is quite clear that a belief state cannot be identified with any *set* of propositions, for example with the set of those propositions that are assigned the probability 1. Should we then identify a belief state with a probability function? Not necessarily. Perhaps a belief state is something that is even more finely structured.

To explore this idea, let us view belief states as "black boxes". That is, let us take the concept of a belief state as a primitive notion. Every belief state is then assumed to be associated with a belief set, the set of all propositions accepted (or believed) in that state, and with a probability function having the belief set in question as its top. Belief revision is now an operation which, when applied to a belief state s and a consistent proposition x, yields a new belief state s • x. For each belief state s, we let [s] be the associated belief set and <s> the associated probability function. We then get the following postulates for belief revision:

(B1) $x \in [s \bullet x]$.

(B2) If $[s] \cup \{x\}$ is consistent, then $[s \bullet x] = [s] + x$.

(B3) If $Cn(\{x\}) = Cn(\{y\})$, then $[s \bullet x] = [s \bullet y]$.

(B4) If $[s \bullet x] \cup \{y\}$ is consistent, then $[s \bullet (x \wedge y)] = [s \bullet x] + y$

The corresponding axioms for probability revision are:

(P1) $<s \bullet x>(x) = 1$.

(P2) If $<s>(x) > 0$, then $<s \bullet x> = <s> + x$.

(P3) If $Cn(\{x\}) = Cn(\{y\})$, then $<s \bullet x> = <s \bullet y>$.

(P4) If $<s \bullet x>(y) > 0$, then $<s \bullet (x \wedge y)> = <s \bullet x> + y$.

The postulates for belief revision can now be derived from the postulates for probability revision via the following *Belief-Probability Principle:*

(BP) $[s] = t(<s>)$.

The Non-Uniqueness Problem does not arise on this approach.

Consider now the following principles:

(1) If $[s_1] = [s_2]$, then $[s_1 \bullet x] = [s_2 \bullet x]$.

(2) If $<s_1> = <s_2>$, then $<s_1 \bullet x> = <s_2 \bullet x>$.

(3) If $t(<s_1>) = t(<s_2>)$, then $t(<s_1 \bullet x>) = t(<s_2 \bullet x>)$.

According to (1), if two belief states are characterized by the same belief set, then the same holds for their revisions with a given proposition. Analogously, (2) says that if two belief states are characterized by the same probability function, then the same holds for their revisions. (3) expresses Top Equivalence within the present framework. Clearly, given (BP), (1) and (3) are equivalent.

It is easy to see that the principle (1), if valid, would allow us to define belief revision as an operation on belief *sets*, just as in the original approach: for any belief set A, let s be an arbitrary belief state such that $A = [s]$ and define $A \oplus x$ as $[s \bullet x]$.

Analogously, if (2) were valid, we would be able to define a probability revision operation ∗ in terms of • : for any P, take an arbitrary s such that P = <s> and let P ∗ x be equal to <s • x>.

9. Concluding Remarks

In this paper, we have considered several possible reactions to the Non-Uniqueness Problem. The most direct response — the assumption of *Top Equivalence* — is not especially plausible on intuitive grounds. The *common part* approach has been shown to confront difficulties, both from a formal and from an intuitive point of view. In particular this approach does not work without the assumption of the Strong Intersection Principle. This principle, although being weaker than Top Equivalence, is still unintuitive.

The third approach, in terms of *base functions,* replaces the original problem of Non-Uniqueness with a problem of arbitrariness: why use this particular base function rather than some other one? On the other hand, the base function approach gives rise to several interesting mathematical questions concerning the representability of belief revision within the probabilistic framework. Seen from this perspective, base functions are just representational mappings from one kind of a model to another.

The *relational* approach and the approach that takes the notion of a *belief state* as a primitive have both the advantage of intuitive plausibility. The former abandons the idea that belief revision must always give a unique result, while the latter relinquishes the thought that belief revision operates on sets of beliefs. On the whole, one of these approaches, or their combination, seems to be the correct response to our original problem.

Notes

[1] We wish to thank Peter Gärdenfors, David Makinson, Howard Sobel and Paul Weirich for stimulating suggestions and criticisms of earlier drafts. We are also grateful for the very helpful comments we received from the participants in an English-Swedish philosophy symposium in London (Fall 87).

[2] The background for our discussion of the Non-Uniqueness Problem is the work by Alchourrón, Gärdenfors, and Makinson on the logic of belief revision and theory change. Our main source of reference is: Gärdenfors (1988). The locus classicus on non-probabilistic belief revision is Alchourrón, Gärdenfors, Makinson (1985). See also the references therein, especially Alchourrón and Makinson (1982) and Makinson (1985). Probability revision is discussed in Gärdenfors (1988) and (1986b).

[3] Of the above conditions, (vi) corresponding to *reductio ad absurdum* seems less intuitive than the others. Dropping (vi) would amount to assuming that the underlying logic is an extension of intuitionistic logic rather than of classical logic.

[4] There are problems with this interpretation having to do with the fact that normally, when we learn that x, we also learn *that* we have learned that x. Thus, very often, revising a belief set by adding x as a *sole* piece of new information may be simply impossible. To our knowledge, this difficulty was first noted by Bas van Fraassen in his review (1980) of Brian Ellis' **Rational Belief Systems**, where he illustrates it with an example due to Richmond Thomason. See also Paul Weirich (1983), Stalnaker (1984), chapter 6, and David Lewis (1986).

[5] Our treatment differs from that of Alchourrón, Gärdenfors, and Makinson in not letting $A \oplus x$ be defined for the case when A or x is inconsistent. In this way, we get a more natural correspondence between belief revision and probability revision. Any belief revision operation in our sense can be extended to a belief revision operation satisfying the Gärdenfors postulates by letting $A \oplus x = Cn(\{x\})$, if A or x is inconsistent.

[6] It should be noted that this axiom would be violated if x had not been the sole piece of new information that A is revised with. Cf. note 4.

[7] The postulates for probability revision are due to Gärdenfors (1986b). His treatment differs from ours in letting probability revision be defined also for inconsistent propositions x, in which case he lets $(P * x)$ be the absurd probability function $P_{\perp}$ that assigns the value 1 to every proposition.

[8] This definition of a Popper function is essentially due to van Fraassen (1976). Our Popper functions differ from his in being defined for propositions rather than events.

[9] While P_x has been defined as $\Pi(-/x)$, the result of an iterated revision, $(P_x)_y$, cannot be defined in terms of Π. In particular, it is impossible to identify $(P_x)_y$ with $\Pi(-/x \wedge y)$. In order to see this, note that the

latter expression is undefined when x and y are inconsistent with each other. But they may still be consistent taken individually, in which case $(P_x)_y$ should be defined.

[10] This method allows iterated revisions of *monadic* probability functions. But what if we want to revise dyadic Popper functions and do it repeatedly? In order to do that we may introduce a special revision operation on Popper functions that transforms every dyadic Π into another dyadic Popper function Π_x (for any x). However, it would then be natural to modify the concept of a *belief* set accordingly. The "top" of a dyadic probability function is not a set of beliefs but rather a set of *conditional* beliefs. A conditional belief x/y (x, *on the condition that* y) belongs to the top of Π iff Π(x/y) = 1. (Note that a conditional belief is not the same as a belief in a conditional. / is not a proposition-forming operation, but rather a relation between propositions. As a consequence, while conditionals can be iterated, for example there is a proposition such as $x \rightarrow (y \rightarrow z)$, there are no such things as (z/y)/x. Hence, there is no reason to fear that, once we allow conditional beliefs in addition to the simple ones, we shall have to introduce "conditionally conditional beliefs", etc..)

What axioms should we impose on the revision operation for Popper functions? As far as we can see there are two axioms that immediately come to mind:

(ΠR1) if $Cn(\{x\}) = Cn(\{y\})$, then $\Pi_x = \Pi_y$.

(ΠR2) $\Pi_x(\text{—}/T) = \Pi(\text{—}/x)$.

In the text above, we have shown how to define ∗ in terms of Popper functions and a mapping s as follows:

(1) $P * x = s(P)(\text{—}/x)$

Given a revision operation on Popper functions, there is an alternative definition of ∗:

(2) $P * x = s(P)_x(\text{—}/T)$.

(ΠR2) guarantees that the two definitions of ∗ determine the same operation. Also, given (2), one can derive the axioms (∗1) - (∗4) from (ΠR1) and (ΠR2) together with the axioms on Popper functions, (Π1) - (Π4).

Note that the axioms on conditional beliefs corresponding to (ΠR1) and (ΠR2) are:

(CBR1) If Cn({x}) = Cn({y}) and A and B are conditional belief sets, then $A_x = A_y$.

(CBR2) If A is a conditional belief set, then $y/T \in A_x$ iff $y/x \in A$.

It should be noted that (CRR2) is a version of the so-called Ramsey principle, see Ramsey (1950), Stalnaker (1970), Gärdenfors (1986a), (1987) and Lewis (1976), (1986). It is well-known that the Ramsey principle leads to triviality results if / is regarded as a propositional operator rather than as a relation between propositions. Further investigation of conditional belief sets and their revisions is an enterprise that must await another occasion.

[11] Cf., for example Sven Ove Hansson (1987).

[12] See his papers (1986a), esp. p. 88 f. and (1986b), esp. p. 33 f. See also Gärdenfors (1988), section 5.8.

[13] See Gärdenfors (1986a) where this triviality result is presented for the analogous principle of monotonicity for belief revision: If $A \subseteq B$, then $A \oplus x \subseteq B \oplus x$. In this paper Gärdenfors describes how his result relates to David Lewis' well-known triviality theorem in "Probabilities of Conditionals and Conditional Probabilities" (1976). See also Gärdenfors' (1987).

[14] The standard axioms for probability revision, in particular (∗2), only imply a weaker monotonicity principle:

$$\text{(WM)} \qquad \frac{t(P) \subseteq t(Q) \text{ and } Q(x) > 0}{t(P * x) \subseteq t(Q * x)} \ .$$

Clearly, given (WM), (RM) is equivalent to the simpler principle:

$$\text{(RM')} \qquad \frac{t(P) \subseteq t(Q) \text{ and } P(x) = 0}{t(P * x) \subseteq t(Q * x)} \ .$$

[15] By a reasoning analogous to that of the previous note, in the presence of the standard axioms for belief revision (in particular (⊕2)), this principle is equivalent to the simpler:

$$\text{(RM'B)} \qquad \frac{A \subseteq B \text{ and } \neg x \in A}{A \oplus x \subseteq B \oplus x} \ .$$

[16] This principle has been drawn to our attention by David Makinson.

[17] There is, of course, an even stronger principle of caution, namely,

Principle of Maximal Caution

For every P and x, if P(x) = 0, then t(P ∗ x) = Cn({x}).

It is easy to check that this principle is fully compatible with the standard axioms on ∗. *Mutatis mutandis*, the same applies to the corresponding principle for belief revision: for every A and x, if x is inconsistent with A, then A ⊕ x = Cn({x}).

[18] Note that the corresponding problem of inappropriateness does *not* arise for the embeddings s from monadic probability functions to the dyadic ones.

[19] Of course, it is also possible to view *probability* revision as a ternary relation. That is, we may introduce a relation **S** with the intuitive interpretation: P $\mathbf{S}_X$ Q iff the probability function Q is a (possible) result (for a given agent) of revising the probability function P by x (as a sole piece of new information). It is straightforward to formulate the appropriate axioms (**S**0) - (**S**4) for *probability revision relations.* In terms of a given such **S** we could then define a belief revision relation **R** by the condition: A $\mathbf{R}_X$ B iff for some P and Q, t(P) = A, t(Q) = B and P $\mathbf{S}_X$ Q. Given this change, the functionality principle (**R**5) will be satisfied iff the probability revision relation **S** satisfies the following generalization of Top Equivalence: For all P, P', Q, Q', if t(P) = t(P'), P $\mathbf{S}_X$ Q and P' $\mathbf{S}_X$ Q', then t(Q) = t(Q').

[20] But then, of course, we would have to determine the appropriate set of axioms for the iterated belief revision relation.

References

Alchourrón, C. E., P. Gärdenfors and D. Makinson (1985): "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", **The Journal of Symbolic Logic 50**: 510-530.

Alchourrón, C. E., D. Makinson (1982) "On the Logic of Theory Change: Contraction Functions and Their Associated Revision Functions", **Theoria 48**, 14-37.

Gärdenfors, P. (1986a) "Belief Revisions and the Ramsey Test for Conditionals", **Philosophical Review 95**: 81-93.

Gärdenfors, P. (1986b) "The Dynamics of Belief: Contractions and Revisions of Probability Functions", **Topoi 5**: 29-37.

Gärdenfors, P. (1987) "Variations on the Ramsey Test: More triviality results", **Studia Logica XLVI**, 321-327.

Gärdenfors, P. (1988) **Knowledge in Flux: Modelling the Dynamics of Epistemic States**, forthcoming in Bradford Books, MIT Press.

Hansson, S. O. (1987) "New operators for Theory Change", circulated manuscript, Department of Philosophy, Uppsala University, Uppsala.

Lewis, D. (1976) "Probabilities of Conditionals and Conditional Probabilities", **Philosophical Review 85**: 287-315. Reprinted in Lewis, D. **Philosophical Papers**, vol. 2, Oxford University Press, Oxford 1986.

Lewis, D. (1986) "Postscript to 'Probabilities of Conditionals and Conditional Probabilities'", **Philosophical Papers**, vol. 2, Oxford University Press, Oxford.

Makinson, D. (1985) "How to Give It Up: A Survey of Some Formal Aspects of the Logic of Theory Change", **Synthese 62**, 347-363.

Ramsey, F. P. (1950). "General Propositions and Causality", **Foundations of Mathematics and other Logical Essays**, ed. by R. B. Braithwaite, New York: Routledge and Kegan Paul, pp. 237-257.

Stalnaker, R. (1970) "Probability and Conditionals", **Philosophy of Science 37**, pp. 64-80.

Stalnaker, R. (1984). **Inquiry**, the MIT Press, Cambridge, Mass.

Van Fraassen, B. (1976) "Representation of Conditional Probabilities", **Journal of Philosophical Logic 5**: 417-430.

Van Fraassen, B. (1980) Review of Brian Ellis' **Rational Belief Systems**, in **Canadian Journal of Philosophy 10**, pp. 487-511.

Weirich, P. (1983) "Conditional Probabilities and Probabilities given Knowledge of a Condition", **Philosophy of Science 50**, pp. 82-95.