Probability, Rational Belief, and Belief Change

Charles G. Morgan
Department of Philosophy
University of Victoria
Victoria, B.C. V8W 3P4
CANADA
cmorgan@uvic.ca

Abstract

A simple model of rational belief holds that: (i) an instantaneous snapshot of an ideally rational belief system corresponds to a probability distribution; and (ii) rational belief change occurs by Bayesian conditionalization. But *a priori* probability distributions of the Kolmogorov sort cannot distinguish between propositions that are simply true from propositions that are necessarily true. Further, propositions accepted by Bayesian conditionalization become necessary truths of the updated distribution. Thus on this model, once you have accepted a proposition, it is impossible to change your mind. These problems are not avoided by Jeffrey conditionalization nor by adopting infinitesimal probability values.

In contrast, conditional probability distributions are able to distinguish propositions that are simply true from propositions that are necessarily true. However, Bayesian conditionalization as a model of belief change still makes the newly accepted proposition necessarily true, and hence immune to future revision. In this paper we develop a different revision scheme for conditional probability distributions that does permit one to accept a proposition with probability 1, but to subsequently change one's mind.

keywords: logic, probability, belief representation, belief change, belief revision

I. Belief Revision Using Kolmogorov and Bayes

Our most detailed theory about reasoning under conditions of uncertainty is formal probability theory. A simple model of rational belief holds that: (i) an instantaneous snapshot of an ideally rational belief system corresponds to a probability distribution; and (ii) rational belief change occurs by Bayesian conditionalization.

The most well known account of the classical theory of elementary probability is due to Kolmogorov in [5]. The theory consists of a set of constraints which any function P must satisfy in order to be a probability function.

KP1 P is defined on a σ-field of sets.
KP2 $0 \leq P(\alpha)$
KP3 $P(U) = 1$
KP4 If $\alpha \cap \beta = \varnothing$, then $P(\alpha \cup \beta) = P(\alpha) + P(\beta)$

Kolmogorov takes *a priori* probability $P(\alpha)$ as basic and defines conditional probability $P(\alpha, \beta)$ in terms of *a priori* probabilities. The expression "$P(\alpha, \beta)$" is read "the probability of α, given β" or "the probability of α, on condition that β." Conditional probability is defined as follows:

DF.KCP.1 If $P(\beta) \neq 0$ then $P(\alpha, \beta) = P(\alpha \cap \beta)/P(\beta)$.

Since it is usual to think of beliefs propositionally, it will be convenient if we think of probability distributions as defined over a language **L**. For our purposes we will assume **L** to be the language of classical propositional logic with the usual connectives: $\wedge$, $\vee$, $\supset$, and $\sim$. Valid inferences in the language are assumed to be governed by a standard classical consequence relation; we write "$\Gamma \vdash A$" for the claim that sentence A follows from the set of sentences Γ. Given that we are dealing with classical logic, we may presume the usual definition of maximal sets.

DF.MCON The set of sentences Γ is **maximal with respect to** sentence A iff (i) it is not the case that $\Gamma \vdash A$, and (ii) for all sentences B, if $B \notin \Gamma$, then $\Gamma \cup \{B\} \vdash A$. The set of sentences Γ is **maximal** iff there is a sentence A such that Γ is maximal with respect to A.

We take the sigma field of events of the usual Kolmogorov theory to be the sigma field of sets of maximal sets. Each proposition in the language **L** corresponds to the set of maximal sets which contain the proposition. The connection between the Kolmogorov constraints and the language **L** is then totally transparent. Thus we may take the probability distributions to be defined over sentences of the langauge **L**, and we can use conjunction, disjunction, and negation instead of intersection, union, and complement, respectively.

How should the belief system of the ideally rational agent be revised in light of new information C? According to the simple view mentioned above, rational belief change should take place in accord with DF.KCP.1. That is, if one's belief state at time t is given as probability distribution P, and at time t+1 one comes to accept evidence C, then one's revised belief state P[C] at t+1 should be given by:

BREV.1 $P[C](x) = P(x, C) = P(x \wedge C) / P(C)$

This proposal for updating beliefs is sometimes referred to as Bayesian revision or Bayesian conditionalization. For convenience, we have adopted the notation "P[C]", which carries an indication of the sentence C with respect to which the distribution P is conditionalized. We remind the reader that the distribution itself carries no such indication.

One of the most well known difficulties with Bayesian conditionalization concerns the situation in which the prior probability assigned by the agent to the newly accepted evidence was 0. When P(C) = 0, the proposed conditional probability distribution given by BREV.1 is mathematically undefined because of division by 0. (This vexing problem has been discussed widely in the literature, but this is not the place for a review; perhaps the earliest treatment is [2], but for more recent discussion see [1], and [14]). Partially in an attempt to avoid such difficulties, some authors (e.g. Gärdenfors [3]) extend the Kolmogorov functions to include the so called "trivial" function which assigns P(A) = 1 for all sentences A. It is then suggested that the Bayesian updating scheme be slightly amended as follows:

BREV.2 $P[C](x) = P(x, C) = P(x \wedge C) / P(C)$, if $P(C) \neq 0$
$= 1$, if $P(C) = 0$

One serious problem with this prescription is that it does not correspond to the way in which it seems rational to assign probabilities in some circumstances.

As a simple example, suppose I am considering the results of flipping a coin. From my beliefs about physics and my knowledge of the coin, I assign a probability value to various outcomes. Typically I only consider the two possibilities of the coin showing heads (H) or tails (T). But hecklers may suggest that we should also include the possibility of the coin coming to rest on edge (E). A perfectly rational agent may well assign a probability of 0.5 to each of the outcomes H and T and a probability of 0 to the possibility of E. The fact that when discussing actual coin tosses, real agents generally do not even mention E is strong evidence that such agents do not consider E a "live" possibility and assign it a probability of 0. But note that if the coin were to come to rest on edge, each of us would be perfectly able to revise our beliefs in a rational way. Given that the coin has come to rest on edge, we would assign a 0 probability to the claim that the coin is showing heads uppermost, and a 0 probability to the claim that the coin is showing tails uppermost. We may represent these beliefs as follows:

EX1.prior	P(H) = 0.5	P(T) = 0.5	P(E) = 0
EX1.revised	P'(H) = 0	P'(T) = 0	P'(E) = 1

The problem is that the belief state given by EX1.revised cannot be obtained from EX1.prior by using either BREV.1 or BREV.2.

One may regard H ∨ T as true, but not regard it as necessarily true; an agent might well be willing to consider a situation in which H ∨ T is false. However, the same agent might well regard the sentence H ∨ ~ H as necessarily true; that is, the agent might be unwilling to revise his/her belief in the truth of H ∨ ~ H no matter what additional information comes to light. Some sentences we think are true, but not necessarily true. But a Kolmogorov *a priori* probability distribution is simply an assignment of values to the sentences. Any two sentences that both take the value 1 (or 0) will be mathematically indistinguishable, given the Kolmogoroff constraints.

Asking me to revise my beliefs in such a way as to accept a logically false proposition might well push me to the position of being unable to reject anything; if I have to accept ~ (H ∨ ~ H) then my revised belief state will be the constant function 1. But asking me to revise my beliefs in such a way as to accept a sentence I previously believed was simply false, but not logically false, should not provoke a switch to the constant function 1.

Another serious problem with the simple view is that the Bayesian updating scheme provides for no way to change your mind. Once we have conditionalized a distribution on some information C, there is no way to re−conditionalize the distribution using the Bayesian scheme so that ~ C has a value other than 0. Once C has been accepted, there is no way to reverse that acceptance using BREV.1 or BREV.2.

As a simple example, let us return to the coin flipping experiment, but for now disregard the possibility of the coin coming up on edge. In this case, there will be only 4 distinct events. There is the tautological event, here designated by "t", the contradictory event, here designated by "f", and the two possibilities of heads "H" and tails "T". All other logical combinations reduce to one of these four. We want to consider two probability distributions. The first represents the state of an agent who does not know which face the coin is showing; the second represents the revised beliefs of the agent when he/she comes to accept that the coin is showing heads. Since the coin may be loaded in some way, we will not prejudice the issue by assigning probabilities of 0.5 to H and T in the first distribution; instead, we will use "n" for the probability of H and "1 − n" for the probability of T, where n is some number between 0 and 1.

	t	H	T	f
P	1	n	1−n	0
P[H]	1	1	0	0

It is clear that we cannot conditionalize P[H] on T or on f, since both have probability of 0. But conditionalizing P[H] on t or on H will leave P[H] unchanged. Thus once we have conditionalized on H, there is no way to change our minds and subsequently reject H.

Various attempts to get around these problems using distributions with infinitesimal values or using Jeffrey conditionalization have been proposed, but they can easily be shown to be unsuccessful.

II. Using Popper Instead of Kolmogorov

Simple reflection on real examples shows that our reasoning about probabilities is never *a priori*. We are always making some sort of background assumptions, such as that there are six faces on a die, the numbers on the faces are positive integers from 1 through 6, the faces do not change when the die is thrown, etc. As we have seen, at least part of the problem with giving a reasonable formal account of how an agent may rationally change his/her mind is associated with the adoption of the Kolmogorov account of probability theory which takes *a priori* distributions as basic.

If we use *a priori* probabilities to define conditional probabilities as in the Kolmogorov theory, then there is no way to distinguish contingently false propositions from necessarily false propositions. If P(C) = 0, then conditionalizing on C is just like conditionalizing on a logical falsehood; if P(C) = 0, then as evidence (or as background assumption) C is treated as no different from a flat contradiction. But recalling the coin example, we may well assign a 0 probability to the coin coming to rest on edge, and yet still distinguish a situation in which the coin does come to rest on edge from the bizarre situation in which we are asked to accept that some contradiction allegedly obtains. We surely should be able to conditionalize our beliefs on statements that we claim to know are contingently false. Even if we claim to know a sentence C is false, we may want to consider how the world might have been had C obtained; or we may recognize that we are not infallible, and even our most firmly held beliefs, e.g. that C is false, may subsequently prove to be mistaken.

There is a simple way of modifying the basic Kolmogorov theory to make it conditional and to get around the problem of conditioning on sets of probability 0. I have taken the central idea for the solution from Popper [13]; I have previously formulated versions of these constraints in [8] and [11].

PP1 P is defined on ordered pairs from a σ-field of sets.
PP2 $0 \leq P(\alpha, \beta) \leq 1$
PP3 $P(U, \alpha) = 1$
PP4 If $\alpha \cap \beta = \varnothing$, then $P(\alpha \cup \beta, \gamma) = P(\alpha, \gamma) + P(\beta, \gamma)$, unless for all δ, $P(\delta, \gamma) = 1$.
PP5 $P(\alpha \cap \beta, \gamma) = P(\alpha, \gamma) \times P(\beta, \alpha \cap \gamma)$

Using the technique outlined in [11] it is easy to use the properly conditionalized Kolmogorov theory to extract the following constraints, based on the formal language.

NP1 $0 \leq P(A, \Gamma) \leq 1$
NP2 If $A \in \Gamma$ then $P(A, \Gamma) =1$.
NP3 $P(A \vee B, \Gamma) = P(A, \Gamma)+P(B, \Gamma)-P(A \wedge B, \Gamma)$
NP4 $P(A \wedge B, \Gamma) = P(A, \Gamma) \bullet P(B, \Gamma \cup \{A\})$
NP5 $P(\sim A, \Gamma) = 1-P(A, \Gamma)$ unless $P(D, \Gamma) = 1$ for all D.
NP6 $P(A \wedge B, \Gamma) = P(B \wedge A, \Gamma)$
NP7 $P(C, \Gamma \cup \{A \wedge B\}) = P(C, \Gamma \cup \{A, B\})$
NP8 $P(A \vee \sim A, \Gamma) = 1$
NP9 For all sets Γ and Δ, if $P(D, \Gamma) = 1$ for all D, then $P(D, \Gamma \cup \Delta) = 1$ for all D.

It should be noted that conditions NP1-9 do not depend on any notions from proof theory, and hence they are autonomous. Thus they are an acceptable basis for probabilistic semantics. Further, the constraints do not depend on any notions from classical formal semantics. Popper in [13] was perhaps first to define fully conditional probabilities directly on sentences of the formal language, although he did not derive them in the way just indicated. Our constraints differ slightly from those originally proposed by Popper; the most significant difference is that we allow the background information to be a set of formulas rather than a single formula. I refer to these constraints as neo-Popperian classical conditional probability theory, or more simply as Popper probability theory. Many authors (e.g. [3], [13], and [14]) have advocated Popper probability functions as better models for belief representation than Kolmogorov functions..

The notion of abnormality will play a significant role in the following material, so we will pause here to give a definition.

DF.PA.2 A set Γ of sentences is said to be **P-abnormal**, for conditional probability function P, just in case for all sentences D, $P(D, \Gamma) = 1$. A set that is not P-abnormal is said to be P-normal.

Trying to regard an *a priori* probability distribution as an instantaneous snapshot of the belief system of an ideally rational agent completely overlooks the fact that we hold many conditional beliefs. It may be possible to account for some conditional beliefs in an *a priori* distribution by treating them in the Bayesian fashion. For example, at this point in time, I cannot say for certain whether or not I will go out to dinner tomorrow evening; but, on the condition that I do go out, I believe that it is more probable that I will eat at Green Cuisine than at More Fat Food. This example does not involve probabilities with extreme values. However, from the coin toss example discussed previously, we know that starting with *a priori* distributions, we cannot account for conditional beliefs in which the background condition has initial probability 0, e.g. the flipped coin coming to rest on edge.

At this point it will be useful to clarify our terminology. When dealing with Popper functions, we will take "Bayesian" revision to mean the following:

BREV.3 $P[C](x, \Gamma) = P(x, \Gamma \cup \{C\})$

Unlike BREV.1, the revised distribution will always be defined, as there is no division by 0. And, unlike BREV.2, with BREV.3 and Popper distributions, we can make sense of belief revision even when the conditioning sentence has an initial probability of 0, as we shall see below.

Fully conditional probability distributions are much better candidates than Kolmogorov functions as models for instantaneous snapshots of rational belief systems. As a

Table 1: Complete Popper Distribution for Coin Toss

$P(A, \Gamma)$		Γ							
		t	H	T	E	$H \vee T$	$H \vee E$	$T \vee E$	f
A	t	1	1	1	1	1	1	1	1
	H	0.5	1	0	0	0.5	1	0	1
	T	0.5	0	1	0	0.5	0	1	1
	E	0	0	0	1	0	0	0	1
	$H \vee T$	1	1	1	0	1	1	1	1
	$H \vee E$	0.5	1	0	1	0.5	1	0	1
	$T \vee E$	0.5	0	1	1	0.5	0	1	1
	f	0	0	0	0	0	0	0	1

Table 2:Coin Toss, Bayesian Conditioned on E

$P[E](A, \Gamma)$		Γ							
		t	H	T	E	$H \vee T$	$H \vee E$	$T \vee E$	f
A	t	1	1	1	1	1	1	1	1
	H	0	1	1	0	1	0	0	1
	T	0	1	1	0	1	0	0	1
	E	1	1	1	1	1	1	1	1
	$H \vee T$	0	1	1	0	1	0	0	1
	$H \vee E$	1	1	1	1	1	1	1	1
	$T \vee E$	1	1	1	1	1	1	1	1
	f	0	1	1	0	1	0	0	1

simple illustration of such a probability distribution, let us return to our coin toss example. Our language contains the three atomic propositions H, T, and E. We will use t for an arbitrary logical truth and f for an arbitrary logical falsehood. We presume that the coin must show one and only one of H, T, or E; e.g., "H" is equivalent to "$H \wedge \sim T \wedge \sim E$". Under these circumstances, there will be only three distinct maximal sets: (i) the set that contains H and all its logical consequences, (ii) the set that contains T and all its logical consequences, and (iii) the set that contains E and all its logical consequences.

We are dealing with a small finite language, so it will be possible to lay out a complete conditional probability distribution on a table. Since there are only three maximal sets, there will be only $2^3 = 8$ logically distinct propositions, each proposition being equivalent to a set of maximal sets and each set of maximal sets being equivalent to a proposition. If each proposition x corresponds to a set of maximal sets max(x), then each set of propositions $\{x_1, x_2, ..., x_n\}$ also corresponds to a set of maximal sets: $\{x_1, x_2, ...,x_n\}$ corresponds to $\max(x_1) \cap \max(x_2) \cap ... \cap \max(x_n)$. Thus each set of propositions is equivalent to some single proposition. So in our tabular representation of probability distributions for this simple case, we will represent each set of expressions by the single expression to which it is equivalent.

An example of a simple probability distribution is given in Table 1. The table corresponds to the case in which the agent believes that H and T are equally probable but believes the probability of E to be 0.

Carefully note, however, that Bayesian conditionalizing on E is not the same as conditionalizing on f, even though $P(E, \{t\}) = 0$. The agent knows that if the coin lands on edge, then neither heads nor tails will be uppermost. Table 2 gives the complete distribution obtained for conditionalizing in the Bayesian way on E; the resulting distribution is certainly NOT the trivial constant 1 function as would be the case if we were using Kolmogorov functions. Note that EX1.prior agrees with the column for $P(-, \{t\})$ on Table 1. And further, EX1.revised agrees with the column for $P[E](-, \{t\})$ on Table 2.

Table 3:Coin Toss, Bayesian Conditioned on H

P[H](A, Γ)		Γ							
		t	H	T	E	H ∨ T	H ∨ E	T ∨ E	f
A	t	1	1	1	1	1	1	1	1
	H	1	1	1	1	1	1	1	1
	T	0	0	1	1	0	0	1	1
	E	0	0	1	1	0	0	1	1
	H ∨ T	1	1	1	1	1	1	1	1
	H ∨ E	1	1	1	1	1	1	1	1
	T ∨ E	0	0	1	1	0	0	1	1
	f	0	0	1	1	0	0	1	1

In our example, an agent whose belief system is given by P, i.e. Table 1, holds that E is certainly false. However, the agent does not hold that E is necessarily false. With the Popper account of probability, it is easy to distinguish between propositions that are deemed to be simply true or false from those that are deemed to be necessarily true or necessarily false.

DEF.TF For conditional probability function P, we say that a proposition x is held to be P-true (or simply true) under assumption Γ if and only if P(x, Γ) = 1. A proposition x is held to be P−false (or simply false) under assumption Γ if ~ x is P-true. We say that x is P−true (or P-false) if P(x, {t}) = 1 (or P(x, {t}) = 0), for t a tautology.

DEF.NTF For conditional probability function P, we say that a proposition x is held to be necessarily P-true (or just necessarily true) under assumption Γ if and only if P(x, Γ ∪ Δ) = 1 for all sets of sentences Δ. A proposition x is held to be necessarily P-false (or just necessarily false) under assumption Γ if and only if ~ x is necessarily P-true. We say that x is necessarily P−true if x is necessarily P−true under assumption {t}, for t a tautology; we say that x is necessarily P−false if ~ x is necessarily P−true.

As examples, note that in our coin toss example, the sentence H ∨ T would be deemed to be true, under assumption {t}, since P(H ∨ T, {t}) = 1. But H ∨ T would not be deemed to be necessarily true, under assumption {t}, since P(H ∨ T, {t} ∪ {E}) = 0.

This distinction between simple truth and necessary truth in probabilistic terms was elaborated in [6] and there used as the basis for a probabilistic semantics for the modal logic S5. The distinction between simply true propositions and necessarily true propositions cannot be made if we begin with the Kolmogorov account and define conditional probabilities in the usual way. The relationship between necessity and abnormality is elementary.

Theorem NP.2 A proposition C is necessarily P−false, given Γ, if and only if Γ ∪ {C} is P−abnormal.

Popper distributions are much better than Kolmogorov distributions for the representation of instantaneous states of belief of ideally rational agents. Using Popper functions, we can represent Bayesian conditionalization on sentences with initial probability 0 in a much more reasonable way than if restricted to Kolmogorov functions. In fact, Gärdenfors has formulated a set of postulates for rational belief revision in probabilistic terms, and he claims (without proof) that his postulates are "... essentially equivalent to Popper's axiomatization of conditional probability functions" ([3], p 123).

As encouraging as these results seem to be, there is still a difficulty with using Bayesian conditionalization as a representation of belief change. The problem is that even with Popper functions, Bayesian conditionalization still does not provide a mechanism for an agent to change his/her mind. Let us return to the coin toss example. Suppose an agent begins with the state of belief given by P as specified by Table 1, in which the agent believes H and T are equally probable but E is false. Suppose subsequently that the agent comes to accept H. Then if we adopt the Bayesian model, the agent's revised belief state should be represented by P conditionalized on H, which is given in Table 3.

Now, suppose the agent subsequently decides that it was a mistake to accept H and that really ~ H should be accepted. There is no sentence the agent could use to conditionalize the distribution in Table 3 to represent such a change of mind using the Bayesian scheme.

III. A Better Model for Belief Revision

Unlike Kolmogorov distributions, Popper distributions allow us to distinguish between simple truths and necessary truths. The ability to make that distinction makes Popper functions much better models of rational belief than Kolmogorov functions. However, using Bayesian conditionalization as a model for belief revision once more

Table 4: Coin Toss, M−conditioned on H

P<H>(A, Γ)		Γ							
		t	H	T	E	H ∨ T	H ∨ E	T ∨ E	f
A	t	1	1	1	1	1	1	1	1
	H	1	1	0	0	1	1	0	1
	T	0	0	1	0	0	0	1	1
	E	0	0	0	1	0	0	0	1
	H ∨ T	1	1	1	0	1	1	1	1
	H ∨ E	1	1	0	1	1	1	0	1
	T ∨ E	0	0	1	1	0	0	1	1
	f	0	0	0	0	0	0	0	1

confuses simple truths and necessary truths. The primary problem with Bayesian conditionalization as a model for belief revision is that propositions accepted using the Bayesian scheme become necessary true propositions in the conditionalized distribution.

We will now introduce a better scheme for belief revision; we will call the new technique variously “M−revision”, “M−updating”, or “M−conditionalization”. We will use the notation P<A> for the result of M−revising distribution P by accepting sentence A. We will write “P<A,B>” instead of “P<A><B>”. In general, the notational index $\langle A_1, A_2, ..., A_n\rangle$ is to be thought of as an ordered n−tupple, the first revision being A_1 and the most recent revision being A_n. If P is a Popper probability distribution, then we define P<A> as follows.

MREV.P $P\langle A\rangle(x, \Gamma) = P(x, \Gamma \cup \{A\})$, if $\Gamma \cup \{A\}$ is P−normal
$= P(x, \Gamma)$, if $\Gamma \cup \{A\}$ is P−abnormal

When we write “P<A>”, the index <A> is simply a notational convenience, in the same way that we have used [A] as a notational convenience to represent Bayesian conditionalization on A. Just as with Bayesian or Jeffrey revision, a function that results from M−revision carries no explicit indication of its origins.

When revising a conditional probability distribution, we must determine new values for P(x, Γ) for each sentence x and each set of background beliefs Γ. The intuitive idea behind M−revision is quite simple. When we revise a conditional probability distribution by accepting A, we simply add A to our background beliefs if that can be done in a non−absurd way. If adding A to the background beliefs produces a probabilistically absurd set, then we just do not add A but preserve the background set unaltered.

As a simple example, we return once more to our coin toss experiment. Recall that Table 1 represents the original function P for the coin toss. In Table 4 we have given P<H>, i.e. the original distribution M−conditioned on H. It will be useful to indicate how some representative entries in Table 4 were determined. First, consider the assumption set {t}. Since {t} ∪ {H} is P−normal, the values for P<H>(x, {t}) are just the the values P(x, {t} ∪ {H}), for all x. On the other hand, the set {T} ∪ {H} is P−abnormal, so the values for P<H>(x, {T}) are the values P(x, {T}), for all x.

Table 4 represents the state of belief of an agent who has come to accept that heads is uppermost on the coin. In Table 4, it is easy to see that H is P<H>−true. However, H is not a necessary P<H>−truth; simply note that neither P<H>(H, {T}) nor P<H>(H, {E}) has the value 1. Recall that Table 3 represents P[H], which is the result of conditioning on H in the Bayesian way. There it is apparent that H is a necessary P[H]−truth. So M−conditionalization and Bayesian conditionalization are rather different. In general, M−conditionalization makes the accepted sentence true without making it necessarily true, whereas Bayesian condtionalization makes the accepted sentence necessarily true.

The first thing that we need to establish about M−revision is that an M-revised probability distribution is still an appropriate probability distribution.

Theorem M.1 If P satisfies NP1−9, then P<A> satisfies NP1−9.

The next thing we wish to establish about M−revision is that it introduces no spurious abnormal sets or necessary truths.

Theorem M.2 Γ is P<A>−abnormal if and only if Γ is P−abnormal.

Theorem M.3 B is necessarily P<A>−true, given Γ, if and only if B is necessarily P−true, given Γ.

The primary reason for introducing a new scheme for conditionalization was that neither the Bayesian scheme nor

Table 5: Coin Toss, M−conditioned on H, then M−conditioned on ~ H

P<H,~H>(A, Γ)		Γ							
		t	H	T	E	H ∨ T	H ∨ E	T ∨ E	f
A	t	1	1	1	1	1	1	1	1
	H	0	1	0	0	0	0	0	1
	T	1	0	1	0	1	0	1	1
	E	0	0	0	1	0	1	0	1
	H ∨ T	1	1	1	0	1	0	1	1
	H ∨ E	0	1	0	1	0	1	0	1
	T ∨ E	1	0	1	1	1	1	1	1
	f	0	0	0	0	0	0	0	1

Table 6: Coin Toss, M−conditioned on ~ H

P<~H>(A, Γ)		Γ							
		t	H	T	E	H ∨ T	H ∨ E	T ∨ E	f
A	t	1	1	1	1	1	1	1	1
	H	0	1	0	0	0	0	0	1
	T	1	0	1	0	1	0	1	1
	E	0	0	0	1	0	1	0	1
	H ∨ T	1	1	1	0	1	0	1	1
	H ∨ E	0	1	0	1	0	1	0	1
	T ∨ E	1	0	1	1	1	1	1	1
	f	0	0	0	0	0	0	0	1

the Jeffrey scheme constituted a reasonable model for how an agent could accept A at one time and subsequently change his/her mind at a later time and accept ~ A. We now wish to show that M−revision handles such cases well.

Let us begin with an example. Recall that Table 4 represents P<H>, i.e. the distribution P, but M−conditioned on H. In Table 5 we give the values for P<H,~H>, which is the distribution P<H>, but M−conditioned on ~ H. For Bayesian conditionalization, the distribution P[H,~H] will always be the trivial constant 1 function. However, it is obvious that Table 5 is not the trivial constant 1 function. Note that P<H,~ H>(~ H, {t}) = 1; that is, ~ H is true in P<H,~ H>, as desired.

Of course one might well wonder how the function P<H, ~ H> relates to other distributions. Once again, we consider an example. For the sake of comparison, Table 6 represents the distribution P<~ H>. It may come as a surprise to see that Table 5 and Table 6 are exactly the same! So P<H,~ H> represents the result of an agent first accepting H, but then subsequently changing his/her mind and accepting ~ H instead.

This happy identity between P<H,~ H> and P<~ H> is not just a fortuitous accident of our example, as the following theorem demonstrates. For any sentence A, M−revision on A followed by M−revision on ~ A is just the same as M−revision on ~ A. The change−of−mind situations so troublesome for Bayesian and Jeffrey revision are easily handled by M−revision.

Theorem M.4 For all sentences A and all Popper distributions P, P<A,~ A> = P<~ A>.

Gärdenfors and others have suggested that we need two types of belief change functions: one function for belief expansion, and one function for belief contraction. A change of mind of the sort we have been considering is then done by first retracting a previously held belief, and then secondly augmenting the retracted belief set with the new belief. The contraction function seems to be the most problematic. In general there will be a multitude of different ways to “simplify” a set of sentences so that the set is compatible with some specified sentence, but as yet nothing proposed seems intuitively compelling. See [3] for further details.

Although it may seem appealing at first sight, on deeper reflection the contraction−expansion approach can be seen to be fundamentally flawed. Suppose we want to revise some prior conditional distribution to reflect the acceptance of A, and suppose we are considering background assumption set Γ. If A is incompatible with Γ, then it seems we should alter Γ in some way, say to Γ(A) so that we can add A without incompatibility. That is, we should first apply some contraction function to remove from Γ all those sentences that are incompatible with A, producing a logically weaker set Γ(A); then we should take as new values the original distribution conditionalized on $\Gamma(A) \cup \{A\}$. But adopting such a stategy would force A to be a necessary truth in the revised probability function. In other words, the contraction−expansion strategy once more seems to blur the distinction between accepting A as true and making A necessarily true.

On the face of it, M−revision obviously does not follow the contraction−expansion strategy and yet it seems to accomplish both jobs. It is worth while trying to understand how this double function is achieved. Intuitively, when we M−conditionalize on sentence A, we consider each background set Γ. If A is compatible with Γ, we use as new values those of the old distribution conditional on $\Gamma \cup \{A\}$; if A is incompatible with Γ, then we abandon A and take as new values those of the old distribution conditional on just Γ. It may seem at first blush that such a strategy is mistaken; instead of abandoning incompatible information in the background assumptions, as is dictated by a contraction scheme, M−revision counsels the abandonment of the new information when incompatibilities arise. The view of M−revision is that if adding a new belief to a set of background assumptions is probabilistically absurd, then one simply should not add that belief to that set of background assumptions. If the new belief is not itself abnormal, then it will be compatible with {t}, so we can always make a normal new belief true in the revised distribution. On the other hand, there is no reason to refuse to countenance other sets of background beliefs that are incompatible with the new information. After all, we are only trying to make the new belief true, not necessarily true. In this way, crucial aspects of the original distribution are preserved in the new distribution so that if a change of mind occurs at a later time, this crucial information can be recovered. Just as Popper functions allow us to muse in a conditional way about events we firmly believe will not occur, so does M−revision allow us to muse in a conditional way about situations incompatible with newly adopted beliefs. If we subsequently change our mind about a newly adopted belief, these conditional musings can be used to establish the newly revised distribution in a reasonable way.

As previously remarked, Bayesian conditionalization of distribution P on sentence A always results in a distribution P[A] in which A is a necessary truth, even if A is P−abnormal. If A is P−abnormal, then P[A] is the trivial constant 1 function, and every sentence becomes a necessary truth. But with M−revision, P and P<A> have exactly the same necessary truths and the same absurd propositions. So if A is P−abnormal, then it will also be P−abnormal in P<A>. If A is P−abnormal, then $\Gamma \cup \{A\}$ will also be P−abnormal for all Γ. So M−revision on a P−abnormal set results in no change to P at all. But if A is P−normal, then A will always be be a simple truth of P<A>; but unless A is necessarily P−true, it will not be necessarily P<A>−true.

Theorem M.5 If A is P−abnormal, then P<A> = P.

Theorem M.6 If A is P−normal, then P<A>(A, {t}) = 1

For Bayesian conditionalization, the order in which revisions are made is irrelevant. But given Theorem M.4, it is obvious that at least in some cases of M−revision, the order in which revisions are made is a relevant consideration. In particular, first accepting A and subsequently accepting ~ A will generally result in a very different distribution from that obtained by first accepting ~ A and subsequently accepting A. However, when there is no conflict among background assumptions and newly accepted information, order does not matter.

Theorem M.7 If $\Gamma \cup \{A, B\}$ is P−normal, then for all sentences x

$$P\langle A,B\rangle(x, \Gamma) = P\langle B,A\rangle(x, \Gamma) = P\langle A \wedge B\rangle(x, \Gamma).$$

Space restrictions prevent going into more detail. But, it should be clear from this brief account that M−revision with Popper functions is without doubt a better model for rational belief than *a priori* distributions and Bayesian conditionalization. As we have seen, M−revision easily accounts for the change−of−mind examples in which an agent first accepts some sentence A and then subsequently rejects A in favor of ~ A. Further, M-revision has the great advantage of tremendous theoretical simplicity compared with alternatives such as [1] or [3]. M-revision is a fully probabilistic model of belief revision, requiring no complex algebraic constructions of “core beliefs”, no “contraction” or “expansion” schemes or similar machinery.

However, there is another sort of problem to which we should turn our attention. Suppose an agent comes to accept A, but then at some later point decides this acceptance was a mistake. The agent may not be in a position to accept ~ A, but rather would just like to return to the original state of belief before A was accepted. It seems that M−revision by itself is not able to handle this situation. The solution to this problem requires combining M−revision with a conditionalized version of Jeffrey revision (see [4]), but space limitations prevent elaboration here.

In current research we are investigating M−revision schemes appropriate for comparative probability relations of the sort discussed in [7], [10], and [11]. As discussed in [9] and [12], we advocate that belief revision is closely linked to non–monotonic reasoning, and we are exploring the connections between such reasoning and M–revision.

References

[1] Arló-Costa, H. 2001. Iterative Probability Kinematics. *Journal of Philosophical Logic* 30: 479-524.

[2] De Finetti, B. 1936. Les probabilités nulles. *Bulletins des Sceinces Mathématiques (premiere partie*) 60: 275-288.

[3] Gärdenfors, P. 1988. *Knowledge in Flux.* Cambridge, Mass: MIT Press.

[4] Jeffrey, R. C. 1965. *The Logic of Decision*. New York, N.Y.: McGraw–Hill.

[5] Kolmogorov, A. 1950. *Foundations of the Theory of Probability*, New York, N.Y.: Chelsea Publishing Company.

[6] Morgan, C.G. 1982. Simple Probabilistic Semantics for Propositional K, T, B, S4, and S5. *Journal of Philosophical Logic* 11: 443-458.

[7] Morgan, C.G. 1984. Weak conditional comparative probability as a formal semantic theory. *Zeitschr. f. math. Logik und Grundlagen d. Math.* 30: 199–212.

[8] Morgan, C.G. 1991. Logic, probability theory, and artificial intelligence – Part I: the probabilistic foundations of logic. *Computational Intelligence* 7: 94–109.

[9] Morgan, C.G. 2000. The Nature of Nonmonotonic Reasoning. *Minds and Machines* 10: 321–360.

[10] Morgan, C.G. 2002 Les probabilités comparatives comme fondements de la logique. *Sémantique et théories de la vérité: l'héritage carnapien*. Lepage F. and F. Fivenc (eds.) Paris/Montréal: Vrin/Bellarmin. 177-202 .

[11] Morgan, C.G. 2003. Two Values, Three Values, Many Values, No Values. *Beyond Two: Theory and Applications of Multiple Valued Logic.* Melvin Fitting and Ewa Orlowska (eds). Heidelberg: Physica-Verlag. 349-374.

[12] Morgan, C.G. 2004 Higher Order Probability Theory and Defeasible Reasoning. *Truth and Probability*. Bryson Brown and Francois Lepage (eds.) Forthcoming.

[13] Popper, K.R. (1959) *The Logic of Scientific Discovery*, Basic Books, Inc., New York.

[14] van Fraassen, B.C. 1995. Fine-grained Opinion, Probability, and the Logic of Full Belief. *Journal of Philosophical Logic* 24: pp. 349-377.