Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Belief Revision in a Probabilistic Setting

James P. Delgrande[1] , **Gerhard Lakemeyer**[2] , **Maurice Pagnucco**[3] , **Joshua Sack**[4]
[1]School of Computing Science, Simon Fraser University,Burnaby, B.C., V5A 1S6 Canada
[2]Dept. of Computer Science, RWTH Aachen University, D-52056 Aachen, Germany
[3]School of Comp. Sci. and Eng., The University of New South Wales, Sydney, NSW 2052, Australia
[4]Dept. of Math. and Statistics, California State University Long Beach, CA 90840, USA
jim@cs.sfu.ca, gerhard@kbsg.rwth-aachen.de, morri@cse.unsw.edu.au joshua.sack@csulb.edu,

Abstract

This work develops an approach to qualitative belief revision in a fully probabilistic setting. We begin with a logic where possible worlds are assigned probabilities. In this logic an agent may believe a formula is true even though the subjective probability of the formula is less than 1.0. Similarly, after revision by a formula ϕ, the agent will believe ϕ is true, even though the agent's subjective probability of ϕ may be less than 1.0. We establish a correspondence with the hallmark AGM postulates for belief revision. Moreover, we use Jeffrey Conditionalisation to establish a link with iterated belief change. To this end, we develop an approach that satisfies appropriately modified Darwiche-Pearl postulates (with clear justification). Thus, we provide a connection between quantitative probabilistic approaches on the one hand and the qualitative formulation of belief change, on the other. This work holds potential for the development of practical belief revision systems by applying a (qualitative) approach to belief change in probabilistic, uncertain domains.

1 Introduction

A primary focus of research in *Knowledge Representation and Reasoning* (KRR) is representing and reasoning with incomplete and uncertain information. Classical logic can be regarded as suitable for expressing incomplete information. Moreover, reasoning techniques, both deductive and nonmonotonic, have been developed for deriving further conclusions. In contrast to these *qualitative* approaches, probability theory provides a *quantitative* framework for dealing with uncertain information.

While probabilistic approaches have enjoyed considerable success,[1] classical probability may nonetheless be too fine-grained or demanding; and in many situations, determining exact numerical probabilities may be difficult or impossible either because 1) the information is missing; or, 2) it is simply not worth the effort to determine the exact probabilities. Furthermore, reasoning with exact probabilities can be complex and relies on certain (in)dependence assumptions. As a result, there has been interest in qualitative probabilistic reasoning where one may assert that some event is more probable than another without giving exact numerical probabilities (Gärdenfors 1975; Delgrande, Renne, and Sack 2019). There has also been work in providing a qualitative account of epistemic reasoning based on an underlying probabilistic semantics (Leitgeb 2013; Delgrande et al. 2022). In such accounts, an agent may believe that a formula ϕ is true (without qualification) even when the subjective probability of ϕ is less than 1.0. The semantics of these approaches is given in terms of *possible worlds*, where probabilities are assigned to worlds.[2] This approach to belief subsumes epistemic logic KD45 (also weak S5) (Chellas 1980) and, while based on a probabilistic semantics, leads to a qualitative notion of belief.

However, an agent's knowledge is generally not static but evolves as it learns new information. The area of *belief change* addresses how an agent's beliefs evolve in light of new information. The fundamental problem of belief change is: given a description of an agent (its *epistemic state*) and new information (a formula), what should the agent's new epistemic state be after it accepts the formula?[3] The central approach to belief change is the so-called *AGM approach* (Alchourrón, Gärdenfors, and Makinson 1985; Peppas 2008). *AGM revision postulates* are taken as characterising minimal conditions on an agent's beliefs that any revision operator should satisfy. Subsequent work, beginning with (Darwiche and Pearl 1997; Nayak et al. 1996), addressed conditions governing how an agent's epistemic state should evolve.

In this paper, we are interested in the relation between qualitative and quantitative approaches to belief change. In our approach, an agent's epistemic state is given by a probability assignment to possible worlds. This assignment implicitly defines a contingent set of beliefs, similar to the approaches in (Leitgeb 2013; Delgrande et al. 2022). Thus an agent may believe that a formula is true, even though the subjective probability of the formula is less than 1.0. Our goal is to investigate belief revision in this setting. The fundamental question we address is: *Under a probabilistic semantics, given an epistemic state and a formula ϕ, what should the agent's epistemic state look like after it comes to*

[1]It seems fair to say, more so than qualitative approaches.

[2]Informally, the probability attached to a world represents the agent's degree of belief that that world is the actual world.

[3]This is a simplification. There are other change operators, such as *contraction* and *merging*; and the original approaches to revision were concerned only with the agent's set of beliefs. However, this simplification is sufficient for our motivation.

accept ϕ? In addressing this question we show when a revision function satisfies the AGM postulates. We also show how the dynamics of the underlying epistemic state can be specified using Jeffrey Conditionalisation (Jeffrey 1990), and how this approach relates to iterated revision (Darwiche and Pearl 1997).

This paper makes several contributions to the boundary between quantitative and qualitative approaches to belief change. Foremost, we provide a comprehensive account of how an agent's contingent beliefs, given in terms of an underlying probabilistic semantics, can evolve following revision by a formula. We show that the operator we introduce is well-behaved; it adheres to well established principles of belief change. Moreover, we use Jeffrey Conditionalisation to define iterated belief change that satisfies appropriately modified Darwiche-Pearl postulates. Lastly, to date, research in belief change has been almost entirely theoretical. An important result of this work is that it provides the basis for a practical, realistic revision operator, where domain-specific information is expressed in terms of assertions with probabilities but where an agent's beliefs are contingent, as are the formulas for revision.

The next section provides the necessary background. §3 informally introduces our approach. §4 details our approach, first describing how a new set of contingent beliefs can be obtained following revision and, second, describing how the agent's epistemic state would evolve following revision. §5 provides conclusions and suggestions for future work.

2 Background

2.1 Formal Preliminaries

Let $\mathcal{L}$ be a propositional language over a finite set of atomic formulas $\mathcal{P}$ and connectives $\neg$ and $\wedge$. We use lower-case Roman letters $a, b, \ldots$, to denote atomic formulas. We use lower-case Greek letters ϕ, ψ, and χ, possibly with subscripts or superscripts, as metavariables for formulas in $\mathcal{L}$. The meaning of formulas $\neg\phi$ and $(\phi \wedge \psi)$ is defined as usual. Other connectives like $\vee, \supset, \equiv$ will be used freely with the usual abbreviations. The set of *possible worlds* W is defined as the set of all subsets of $\mathcal{P}$ (i.e. $2^{\mathcal{P}}$). Let $w \in W$ and let ϕ be a formula; then w *satisfies* ϕ, denoted $w \models \phi$, according to the following:

$$\begin{array}{lll} w \models p & \text{iff} & p \in w, \text{ where } p \in \mathcal{P} \\ w \models \neg\phi & \text{iff} & w \not\models \phi \\ w \models \phi \wedge \psi & \text{iff} & w \models \phi \text{ and } w \models \psi \end{array}$$

By $[\phi]$ we denote the set of possible worlds where ϕ is true (i.e. $[\phi] = \{w \in W \mid w \models \phi\}$). The worlds $[\phi]$ can be thought of as the proposition expressed by the formula ϕ. For $X \subseteq W$, $Cn(X) = \{\phi \mid w \models \phi \text{ for all } w \in X\}$ denotes the *belief set* entailed by X. We overload notation and for $\phi \in \mathcal{L}$ we also define $Cn(\phi) = \{\psi \mid \text{if } w \models \phi, \text{then } w \models \psi \text{ for all } w \in W\}$. Recall that a belief set is closed under logical deduction: $Cn(\phi) = Cn(Cn(\phi))$. See (Chellas 1980; Hughes and Cresswell 1996) for more on possible-world semantics.

We will deal with total preorders over the set of possible worlds. Recall that a total preorder $\preceq$ is a binary relation that is *reflexive* ($w \preceq w$), *transitive* ($w_1 \preceq w_2$ and $w_2 \preceq w_3$ implies $w_1 \preceq w_3$), and *connected* ($w_1 \preceq w_2$ or $w_2 \preceq w_1$). We define $w \prec w'$ as $w \preceq w'$ and $w' \not\preceq w$. For $A \subseteq W$, $\min(A, \preceq)$ denotes the $\preceq$-least members of A. Hence, for $\phi \in \mathcal{L}$ we have

$$\min([\phi], \preceq) = \{w \in W \mid w \models \phi \text{ and } \nexists w' \in [\phi] \text{ s.t. } w' \prec w\}.$$

A *probability assignment* P to possible worlds is a mapping $P : W \to (0, 1)$ such that $\sum_{w \in W} P(w) = 1$. This extends naturally to sets of possible worlds by, for $A \subseteq W$, $P(A) = \sum_{w \in A} P(w)$. The *conditional probability* of A given B (where $P(B) > 0$) can then be defined by $P(A|B) = \frac{P(A \cap B)}{P(B)}$. A and B are *independent* just when $P(A \mid B) = P(A)$ or equivalently when $P(A \cap B) = P(A) \cdot P(B)$. A and B are *conditionally independent given* C just when $P(A \mid B \cap C) = P(A \mid C)$ or equivalently when $P(A \cap B \mid C) = P(A \mid C) \cdot P(B \mid C)$. We will occasionally overload notation and say that formulas ϕ and ψ are (conditionally) independent just when $[\phi]$, $[\psi]$ are. See (Halpern 2003; Russell and Norvig 2010) for more on probability theory and its role in reasoning about uncertainty in AI.

A standard technique for adjusting probabilities is called *Jeffrey's Rule* (Jeffrey 1990), after its originator. For our purposes, consider where we have a probability assignment $P(\cdot)$ over possible worlds W, and where $A \subset W$, $0 < P(A) < 1$, is assigned a new probability $s = P'(A)$. Then the probability distribution for $B \subseteq W$ is updated as follows:

$$P'(B) = P'(A) \cdot P(B \mid A) + (1 - P'(A)) \cdot P(B \mid (W \setminus A)).$$

In terms of possible worlds, this simplifies to, for $w \in W$:

$$P'(w) = \begin{cases} s \cdot P(w)/P(A) & \text{if } w \in A \\ (1 - s) \cdot P(w)/(1 - P(A)) & \text{if } w \notin A \end{cases} \quad (1)$$

2.2 Belief Change

Belief change (Peppas 2008) studies how an agent's beliefs evolve in the presence of new information. The primary change operator is belief *revision*,[4] in which an agent comes to accept a formula ϕ as true, and the issue is how to consistently incorporate ϕ into its set of beliefs. The challenging situation is when ϕ is inconsistent with the agent's beliefs.

In the original AGM account of belief change (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988) an agent's belief state is given by a set of formulas closed under deduction, called a *belief set*. That is, a belief set is made up of those formulas the agent is committed to believing. A revision function maps a belief set and a formula into a new belief set in which that formula is believed. Such a function is not arbitrary, but instead satisfies certain constraints, which we get into below. A key point is that the revision function effectively encodes what the agent's new beliefs should be for each input formula. A significant limitation of the approach is that it cannot be iterated, or, more

[4] Other operators include belief *contraction* and *expansion* (Alchourrón, Gärdenfors, and Makinson 1985) and *merging* (Konieczny and Pino Pérez 2002).

accurately, that there are no meaningful constraints on a sequence of revision operations.

This was first addressed by (Darwiche and Pearl 1997). In the (Darwiche and Pearl 1997) approach an agent is characterised by its *epistemic state*. An epistemic state is composed not just of the agent's belief set, but it also contains information as to what a new epistemic state would look like following any revision. Thus in such approaches, revision is a function from an epistemic state and a formula to a new epistemic state. Under this perspective, standard AGM revision can be recast as a function from an epistemic state and a formula to a belief set. As a result, we have the following two broad approaches to belief revision.

AGM Revision Given an epistemic state and a formula, determine the agent's new belief set.

Iterated Revision Given an epistemic state and a formula, determine the agent's new epistemic state.

We next review these approaches.

AGM Revision The standard framework for belief change is the AGM approach (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988), named after its founders. There are two aspects to this approach. First, postulates that govern any rational belief revision function are provided. Second, formal constructions are specified for such functions. These two aspects are shown to be equivalent, in the sense that any function that satisfies the postulates can be specified via the construction, and vice versa.

The AGM Postulates for belief revision are as follows. Given a *belief set* K and a formula ϕ, the revision of K by ϕ, denoted $K * \phi$, is characterised by the following *rationality postulates*:

$(K*1)$ $K * \phi = Cn(K * \phi)$
$(K*2)$ $\phi \in K * \phi$
$(K*3)$ $K * \phi \subseteq Cn(K \cup \{\phi\})$
$(K*4)$ If $K \not\vdash \neg\phi$, then $Cn(K \cup \{\phi\}) \subseteq K * \phi$
$(K*5)$ $K * \phi = Cn(\bot)$ iff $\vdash \neg\phi$
$(K*6)$ If $\vdash \phi \equiv \psi$, then $K * \phi = K * \psi$
$(K*7)$ $K * (\phi \wedge \psi) \subseteq Cn(K * \phi \cup \{\psi\})$
$(K*8)$ If $\neg\psi \notin K * \phi$, then $Cn(K * \phi \cup \{\psi\}) \subseteq K * (\phi \wedge \psi)$

The result of a revision operation as expressed by postulate (K*1) is a belief set, that is to say, a set of formulas.

One of the principal constructions for AGM belief revision is Grove's (1988) *systems of spheres*. Katsuno and Mendelzon (KM) (1991) introduced a simplification of the Grove approach, in that they deal with a finitary language, whereas Grove's approach allows an infinitary language. The KM approach has arguably become the most referred-to construction, and we adopt it here in our framework.

Towards a *semantic construction* of revision operators, following Katsuno and Mendelzon (1991), we introduce the notion of a *faithful assignment* over possible worlds.[5]

Definition 1 (Faithful Assignment) *Let K be a belief set. A* faithful assignment *over possible worlds W is a total pre-order $\preceq_K$ satisfying the following properties:*

1. *If $w, w' \in [K]$, then $w' \preceq_K w$; and,*
2. *If $w \in [K]$ and $w' \notin [K]$, then $w \prec_K w'$.*

Katsuno and Mendelzon (1991) provide a result showing the equivalence between operators satisfying the postulates and faithful assignments.

Theorem 2 (KM 1991) *A revision operator $*$ satisfies conditions (K*1)–(K*8) if and only if there exists a faithful assignment that maps each belief set K to a total pre-order $\preceq_K$ such that $[K * \phi] = \min([\phi], \preceq_K)$.*

Informally, an agent's epistemic state is given by a total preorder over worlds (the faithful assignment) where the K worlds (i.e., $[K]$) are minimal in the ordering (i.e., are most preferred), and revision of this total preorder by a formula ϕ is characterised by the minimal ϕ worlds in the preorder.

Iterated Revision The AGM approach says nothing about how the total preorder $\preceq_K$ changes following a revision. This was first addressed by (Darwiche and Pearl 1994) where revision takes an *epistemic state* $\mathcal{K}$ and formula ϕ to a new epistemic state $\mathcal{K}'$.[6] The idea is that an epistemic state includes information regarding how this state would change following a revision. For (Darwiche and Pearl 1997) an epistemic state $\mathcal{K}$ *induces* a total preorder on worlds and, following revision by ϕ, in the new epistemic state $\mathcal{K}'$ the minimal worlds in the new induced preorder are exactly $\min([\phi], \preceq_{\mathcal{K}})$. They provide four postulates for iteration and four properties on how an epistemic state changes, along with a representation result linking these approaches. We defer listing these postulates and properties until we come to iteration in our approach but, informally, in revising by a formula ϕ, the ϕ worlds are moved "down" in the total preorder relative to the $\neg\phi$ worlds so that, first, ϕ worlds (resp. $\neg\phi$ worlds) retain their relative positions in the ranking, and second, in the resulting total preorder, all minimum worlds are ϕ worlds. Related to this, Boutilier (Boutilier 1996) developed an approach in which the $\min([\phi], \preceq_{\mathcal{K}})$ worlds become the least worlds in the new preorder, but the rest of the preorder is unchanged.

While the focus in such work is on total preorders over possible worlds, this need not be the case. For example, an approach based on distances between worlds is investigated in (Lehmann, Magidor, and Schlechta 2001). Indeed, for our approach described herein, an epistemic state is given by a probability assignment to possible worlds, which in turn *induces* a total preorder over worlds.

2.3 Belief in a Probabilistic Setting

Our interests lie in defining belief revision in a probabilistic setting, that is in a possible-world model in which there is a probability assignment P to the worlds in W. Before we can do this, we require a notion of *belief* in a probabilistic setting. A common intuition is that a formula ϕ is believed by an agent just when its subjective probability of ϕ is greater than some threshold r. A difficulty with this intuition is that

[5] This is given in terms of the models of a belief set to anticipate our approach to follow.

[6] Subsequent work in the area includes (Darwiche and Pearl 1997; Nayak et al. 1996; Boutilier 1996; Jin and Thielscher 2007; Delgrande and Jin 2012).

one may not have the principle of *Conjunction* of beliefs: it allows belief in formulas ϕ_1 and ϕ_2 but not necessarily in $\phi_1 \wedge \phi_2$. One can pursue this line of reasoning;[7] however it requires that any resulting approach would not subsume the commonly-accepted logic of belief KD45 (equivalently, weak S5) (Chellas 1980).

Leitgeb (2013) addresses this issue in an algebraic approach, in which belief is defined in a probabilistic setting and in which *Conjunction* is shown to hold. The crucial notion in this approach is what he calls *P-stability*.

Definition 3 (P-stable (Leitgeb 2013)) *Let P be a probability distribution over W,*[8] *and $0.5 \leq r < 1.0$ a belief threshold. Then $X \subseteq W$ is P-stable^r iff for every $Y \subseteq W$, such that $X \cap Y \neq \emptyset$ and $P(Y) > 0$, we have $P(X|Y) > r$.*

Intuitively, for a P-stable^r set X, the worlds in X are "much" more probable than the worlds in $W \setminus X$. How much more? Leitgeb gives the equivalent definition: X is P-stable^r iff:

$$\text{for every } w \in X, \text{ we have } P(w) > \frac{r}{1-r} \cdot P(W \setminus X). \quad (2)$$

If we take $r = 0.5$, this reduces to: X is P-stable^r iff:

$$\text{for every } w \in X \text{ we have } P(w) > P(W \setminus X). \quad (3)$$

It is shown that if X_1 and X_2 are P-stable^r and contain no zero-probability worlds, then either $X_1 \subseteq X_2$ or $X_2 \subseteq X_1$. Hence the set of such P-stable^r sets can be linearly ordered by set inclusion.

If an agent's beliefs are characterised by a P-stable^r set, then the agent may believe a proposition to be true even though that proposition has (subjective) probability less than 1.0. Moreover, the involvement of a P-stable^r set ensures the underlying logic subsumes the logic of belief KD45 (or: weak S5) and so the Conjunction principle is satisfied: If an agent believes ϕ_1, ϕ_2, then it will also believe $\phi_1 \wedge \phi_2$. Consider the following example.

Example 4 *We have that arthritis (a) is causally related to having a sore back (b) and swollen fingers (f). Also, b and f are conditionally independent given a. We have the following probabilities: $P(a) = 0.1$, $P(b \mid a) = 0.75$, $P(b \mid \neg a) = 0.1$, $P(f \mid a) = 0.6$, $P(f \mid \neg a) = 0.05$. This gives rise to the following full probability distribution.*

	a		$\neg a$	
	b	$\neg b$	b	$\neg b$
f	0.045	0.015	0.0045	0.0405
$\neg f$	0.030	0.010	0.0855	0.7695

Take $r = 0.5$. Let $S_1 = \{\neg a \neg b \neg f\}$, $S_2 = \{abf, ab\neg f, \neg ab\neg f, \neg a\neg bf\}$, $S_3 = \{a\neg bf, a\neg b\neg f\}$, $S_4 = \{\neg abf\}$. The P-stable sets are given by S_1, $S_1 \cup S_2$, $S_1 \cup S_2 \cup S_3$, $S_1 \cup S_2 \cup S_3 \cup S_4$. Consequently in the minimal P-stable set, S_1, it is believed that none of arthritis, sore back or swollen fingers is the case.

[7] As Kyburg (1961) does in examining the *Lottery Paradox*.

[8] Leitgeb uses the more general notions of σ-algebras and probability measures, which we do not need as W is finite.

Overall, Leitgeb presents an algebraic approach to belief and conditional belief, in which belief is defined in terms of a sufficiently-high probability. This approach has a parallel with the AGM approach to revision wherein conditioning is analogous to revision. Leitgeb shows this in a representation result: Let an agent's belief set be characterised by the least P-stable^r set X and let a revision by proposition B be given by $A \cap B$ where A is the least P-stable^r set such that $A \cap B \neq 0$ and for which $P(A|B) > r$. This can be precisely captured via a set of postulates that characterise the notion of belief along with others analogous to the AGM postulates.

The notion of P-stability arises independently in (Delgrande et al. 2022) in which a logic of qualitative probability is presented. This logic contains two modal operators: $B\phi$, expresses that the agent believes that ϕ is true, and $\phi \preceq \psi$, expresses that the agent believes that ψ is at least as likely as ϕ. The logic contains the axiom:

$$(B\phi \wedge (\phi \preceq \psi)) \supset B\psi.$$

It is shown that this axiom is captured semantically by the condition of P-stability.

2.4 Belief Change and Probability

Here we briefly review work related to probability and belief revision. Early research in belief revision and probability was generally concerned with revising a probability function. In these approaches, an agent's belief set K is determined by those formulas with probability 1.0, called the *top* of the probability function. For a revision $K * \phi$, the probability function is revised by ϕ, and the belief set corresponding to $K * \phi$ is given by the top of the resulting probability function. One difficulty in revising probability functions is the *non-uniqueness problem*; there are many different functions that have K as their top. Lindström and Rabinowicz (1989) consider various ways of dealing with this problem. Boutilier (1995) considers the same general framework but focuses on issues of iterated belief revision. Gärdenfors (1988, Ch. 5) also considers an extension of the AGM approach to the revision of probability functions.

In a different vein, (Delgrande 2012) addresses how probability update via Jeffrey's Rule can be reconciled with AGM revision, or alternatively, how Jeffrey's Rule can be considered an instance (or proto-instance) of AGM-style revision. Unsurprisingly, few of the AGM postulates are satisfied via changing beliefs in this way.

We have already reviewed Leitgeb's (2013) approach. To some extent this work was anticipated by Snow's (1994) account of *atomic bound systems*, in which there is a total order on atoms such that the probability of an atom is greater than the sum of probabilities of atoms below that atom. Similarly, Benferhat et al. (1999) do the same thing with respect to interpretations with their *big-stepped probabilities*.

3 The Approach: Intuitions

The previous section introduced an epistemic logic in which possible worlds have associated probabilities. In this logic an agent may believe a proposition to be true even though that proposition has probability less than 1.0, and the logic

itself subsumes KD45 (or: weak S5). Our goal is to provide a theory of (iterated) belief change in this logical setting.

In our approach, an epistemic state is a tuple $\mathcal{K} = \langle W, P, r\rangle$. Derived from $\mathcal{K}$ is a set of worlds $X \subseteq W$ characterising the agent's contingent beliefs that we call a *pw-belief set*. This is just the minimum P-stable set in W (see Definition 8 below). A revision by a formula ϕ, denoted $\mathcal{K} * \phi$, yields a new epistemic state $\mathcal{K}' = \langle W, P', r\rangle$ in which, for the resulting pw-belief set X', ϕ is believed; i.e. $X' \subseteq [\phi]$. This can be called *epistemic state revision*.

We first show that the AGM approach is compatible with this framework. We show that an epistemic state $\mathcal{K} = \langle W, P, r\rangle$ induces a total preorder over worlds, $\preceq_{\mathcal{K}} = \langle K_0, K_1, \ldots\rangle$ by taking K_0 to be the pw-belief set in W, K_1 to be the pw-belief set wrt $W \setminus K_0$, and so on. Then, for proposition ϕ we can define $\min([\phi], \preceq_{\mathcal{K}})$ as usual. If we define the agent's belief set following revision by ϕ to be given by $\min([\phi], \preceq_{\mathcal{K}})$, it can be shown that the AGM postulates are satisfied. We also show the converse, that any operator satisfying the AGM postulates can be captured by such an epistemic state revision.

While this shows that the basic AGM approach can be expressed, it doesn't say anything regarding the dynamics of an epistemic state. That is, for an epistemic state $\mathcal{K}$ and proposition ϕ, there is the issue of how $\mathcal{K}$ itself should be revised so that ϕ is believed in the new epistemic state $\mathcal{K}'$. This in turn means that, if the agent doesn't believe that ϕ holds but revises its beliefs by ϕ, the subjective probability of ϕ will increase to the extent that ϕ is believed to be true. In our setting, we apply Jeffrey's Rule so that ϕ is accepted as being true, in that ϕ is true at all worlds in the resulting pw-belief set.

Unless ϕ is unsatisfiable, there is some probability s that can be given to ϕ so that ϕ is subsequently accepted. Since we do not allow possible worlds to be assigned a probability of 0, s cannot be assigned probability 1.0. Therefore, $s \in (0.0, 1.0)$, i.e. we can choose $s = 1.0 - \epsilon$ for suitably small value of ϵ (or less). There will also be some minimum value of s such that, in the resulting epistemic state, ϕ is believed.

A guiding intuition in belief revision is the *principle of minimal change*, which states that an agent should change its beliefs as little as possible in order to accept new information and so, in applying Jeffrey's Rule, one should choose the minimum value of s such that ϕ is believed and that properties of a pw-belief set hold. To this end, we provide conditions on s such that Jeffrey's Rule yields an epistemic state in which ϕ is believed and the AGM postulates are satisfied. However, we also show that, on occasion, the AGM postulates (K*3) and (K*7) may not be satisfied; however the remaining postulates are always satisfied in the approach. We also show that a slightly weaker version of the Darwiche-Pearl postulates holds.

4 Belief Change

In this section we present our approach. We begin by defining our notion of an epistemic state and providing some simple results on epistemic states. After this we consider (uniterated) AGM-style revision, essentially recasting Leitgeb's approach in a propositional setting. The major part of the paper is given over to an account of epistemic state revision, in which we examine the Darwiche-Pearl approach.

4.1 Epistemic States

We begin by defining a *(probabilistic) epistemic state*.

Definition 5 (Epistemic State) *Let $\mathcal{P}$ be a finite set of atomic formulas. An* epistemic state *is a tuple $\mathcal{K} = \langle W, P\rangle$, where:*

- $W = 2^{\mathcal{P}}$ *is the set of possible worlds;*
- $P : W \to (0, 1)$ *is a probability assignment (i.e.* $\sum_{w \in W} P(w) = 1$*).*

$\mathcal{ES}$ is the set of all epistemic states.

For simplicity, and in contrast with Leitgeb (2013), we make two simplifying assumptions:

1. The belief threshold is fixed with $r = 0.5$. This yields the strongest underlying approach that satisfies *Conjunction*.
2. Possible worlds are indeed possible, and so have non-zero probability.

Given an epistemic state $\mathcal{K} = \langle W, P\rangle$ and $\emptyset \neq W' \subseteq W$, we define the notion of a *P-stable subset X with respect to W'*.

Definition 6 (P-stable) *Let $\mathcal{K} = \langle W, P\rangle$ be an epistemic state and let $\emptyset \neq W' \subseteq W$. Then $X \subseteq W'$ is P-stable with respect to W' iff*

$\forall w \in X$, *we have* $P(w) > P(W' \setminus X)$.

If X is P-stable wrt $W' = W$, we will simply say that X is P-stable.

We introduce some results and elaborations on epistemic states that will be useful later.[9]

Proposition 7 *Let $\mathcal{K} = \langle W, P\rangle$ be an epistemic state and $\emptyset \neq W' \subseteq W$. Let X be P-stable with respect to W'. Then:*

1. *If $w \in X$ and $w' \in W' \setminus X$, then $P(w) > P(w')$;*
2. *If $w, w' \in W'$ and $P(w) = P(w')$, then $w \in X$ iff $w' \in X$;*
3. *If X_1 and X_2 are P-stable wrt W', then $X_1 \subseteq X_2$ or $X_2 \subseteq X_1$;*
4. *W' is the largest P-stable set with respect to W'.*
5. *If X is P-stable wrt W' and $X \not\subseteq X'$ then $X \setminus X'$ is P-stable wrt $W' \setminus X'$.*
6. *If X is P-stable wrt W', $X' \cap W' = \emptyset$ and $\forall w \in X'$ $P(w) > P(W' \setminus X)$, then $X \cup X'$ is P-stable wrt $W' \cup X'$.*

Proof:

1. Since $w \in X$ we have $P(w) > P(W' \setminus X)$. Since $x' \in W' \setminus X$ we have $P(W' \setminus X) \geq P(w')$. Hence $P(w) > P(w')$.
2. If $w, w' \in W'$ and $P(w) = P(w')$, then from the preceding item it cannot be the case that $w \in X$ and $w' \notin X$ nor can it be that $w \notin X$ and $w' \in X$. Hence either $w, w' \in X$ or $w, w' \notin X$, which is to say $w \in X$ iff $w' \in X$.

[9] The first four items appear in (Leitgeb 2013) but it is convenient to also list them here.

3. Assume to the contrary that X_1 and X_2 are P-stable wrt W' but $X_1 \not\subseteq X_2$ and $X_2 \not\subseteq X_1$. Hence $\exists w_1 \in X_1$ where $w_1 \notin X_2$ and $\exists w_2 \in X_2$ where $w_2 \notin X_1$.

   Since $w_1 \in X_1$ and $w_2 \notin X_1$, by Part 1 of this Proposition we obtain that $P(w_1) > P(w_2)$. Similarly, since $w_2 \in X_2$ and $w_1 \notin X_2$ by Part 1 of this Proposition we obtain that $P(w_2) > P(w_1)$, contradiction.

   Consequently, if X_1 and X_2 are P-stable, it must be that $X_1 \subseteq X_2$ or $X_2 \subseteq X_1$.

4. Follows trivially from the definition of P-stability (Definition 6).

5. Since X is P-stable wrt W', we have $\forall w \in X$ that $P(w) > P(W' \setminus X)$. As well, $P(W' \setminus X) \geq P((W' \setminus X') \setminus X)$. Thus, $\forall w \in X \setminus X'$, we have $P(w) > P((W' \setminus X') \setminus X)$. i.e. $X \setminus X'$ is P-stable wrt $W' \setminus X'$.

6. Since X is P-stable wrt W', we have, $\forall w \in X$, that $P(w) > P(W' \setminus X)$. Also, by assumption, $\forall w \in X'$, $P(w) > P(W' \setminus X)$. Therefore, $\forall w \in X \cup X'$, $P(w) > P(W' \setminus X)$. (A)

   However, given that we have $X' \cap W' = \emptyset$ and $X \subseteq W'$, this gives that $P(W' \setminus X) = P((W' \cup X') \setminus (X \cup X'))$.

   Thus from (A) we obtain: $\forall w \in X \cup X'$, $P(w) > P((W' \cup X') \setminus (X \cup X'))$, i.e. that $X \cup X'$ is P-stable wrt $W' \cup X'$. $\square$

We introduce the notion of a *possible-world belief set* (*pw-belief set*) *with respect to* $W' \subseteq W$, which is the smallest set of possible worlds that may reasonably be believed in W'.

Definition 8 (pw-belief set) *Given an epistemic state $\mathcal{K}$ and subset $W' \subseteq W$, define the* possible-world belief set (pw-belief set) wrt W' *to be the smallest subset $X \subseteq W'$ such that X is P-stable wrt W'.*

If X is a pw-belief set wrt $W' = W$ we simply say that X is a pw-belief set. While not obvious, it can be shown that a pw-belief set is unique for any given epistemic state.

Proposition 9 *Given an epistemic state $\mathcal{K}$ and $\emptyset \neq W' \subseteq W$, there is a unique pw-belief set wrt W'.*

It can be noted from Proposition 7 that the P-stable sets in an epistemic state $\mathcal{K}$ can be arranged by set containment. For P-stable sets $X_0 \subset \cdots \subset X_n$ we define the sequence $\mathcal{X}_\mathcal{K} = \langle X_0, \ldots, X_n \rangle$.

4.2 AGM-Style Revision

An epistemic state $\mathcal{K}$ naturally induces a total preorder $\preceq_\mathcal{K} = \langle W_0, \ldots, W_n \rangle$ over possible worlds. The agent's beliefs are most naturally given by the minimal P-stable set, W_0, in W. If the worlds in W_0 were determined to not be the case, then the agent's beliefs in $W \setminus W_0$ would be given by the minimal P-stable set, W_1, with respect to $W \setminus W_0$. Repeating this process gives us a total preorder over W. Formally, we have:

Definition 10 (Total preorder induced on W) *Let $\mathcal{K} = \langle W, P \rangle$. Define $\preceq_\mathcal{K} = \langle W_0, \ldots, W_n \rangle$ as follows:*

1. *W_0 is the pw-belief set wrt W.*
2. *For $i > 0$ and where $W \setminus \bigcup_{j=0}^{i-1} W_j \neq \emptyset$:*
   W_i is the pw-belief set wrt $W \setminus \bigcup_{j=0}^{i-1} W_j$.

Note that $\preceq_\mathcal{K}$ defines a total preorder if we take $w_i \preceq_\mathcal{K} w_j$ iff $w_i \in W_i$, $w_j \in W_j$, and $i \leq j$. In case $i < j$ we write $w_i \prec_\mathcal{K} w_j$. When we come to iterated belief revision, it will be important to note the difference between $w_1 \preceq_\mathcal{K} w_2$ and $P(w_1) \geq P(w_2)$; the latter implies the former but the converse may not hold.

For epistemic state $\mathcal{K}$, we can relate the total preorders $\preceq_\mathcal{K}$ and $\mathcal{X}_\mathcal{K}$ as follows:

Proposition 11 *Let $\mathcal{K} = \langle W, P \rangle$. For*
$\preceq_\mathcal{K} = \langle W_0, \ldots, W_n \rangle$ *and* $\mathcal{X}_\mathcal{K} = \langle X_0, \ldots, X_m \rangle$
we have that $W_0 = X_0$, that $n = m$, and $\forall i, 0 < i \leq n$:

1. $X_i = \bigcup_{j=0}^{i} W_j$.
2. $W_i = X_i \setminus X_{i-1}$

Using Definition 10, basic AGM-style revision can be characterised via Theorem 2 to provide the following result.

Theorem 12 *Let $\mathcal{K} = \langle W, P \rangle$ be an epistemic state and let $\phi \in \mathcal{L}$. If we define $\mathcal{K} * \phi = Cn(\min([\phi], \preceq_\mathcal{K}))$ then $*$ satisfies the AGM postulates $(K * 1)$–$(K * 8)$.*

Leitgeb(2013) provides an analogous result. However his result is purely algebraic, expressed in terms of possible worlds, whereas here the results are in terms of an associated propositional language. We also obtain the other half of a representation result:

Theorem 13 *Let $* : 2^\mathcal{L} \times \mathcal{L} \mapsto 2^\mathcal{L}$ satisfy the AGM postulates. There is an epistemic state $\mathcal{K}$ with associated pw-belief set W_0 such that for $K = Cn(W_0)$ and for every $\phi \in \mathcal{L}$, $K * \phi = Cn(\min([\phi], \preceq_\mathcal{K}))$.*

4.3 Epistemic State Revision

A limitation of the AGM approach is that it says nothing about how an epistemic state changes following revision. In this subsection we characterise revision of an epistemic state $\mathcal{K}$ by formula ϕ, giving a new epistemic state $\mathcal{K}'$ in which ϕ is believed. That is, if an agent comes to believe that ϕ is true, then its subjective probabilities are modified, so that in the resulting state $\mathcal{K}'$, the agent's pw-belief set corresponds to $\min([\phi], \preceq_\mathcal{K})$. We modify probabilities using Jeffrey's rule, as it is a prominent generalization of probabilistic conditioning that allows the proposition[10] one updates by to result in a probability less than 1. We repeat Jeffrey's rule here for convenience. Where proposition A is to be assigned probability s:

$$P'(w) = \begin{cases} s \cdot P(w)/P(A) & \text{if } w \in A \\ (1-s) \cdot P(w)/(1-P(A)) & \text{if } w \notin A \end{cases}$$

In revising by ϕ, it clearly only makes sense to increase the probability of ϕ worlds. The most obvious choice of A is that $A = [\phi]$, that is, the agent's subjective probability for all $[\phi]$ worlds increases. This yields a probabilistic approach to (Darwiche and Pearl 1997), which we consider

[10] Proposition can mean a formula φ or it can mean a set of worlds A. Here we mean a set of worlds A, and discuss below the relationship between φ and A.

now. (We consider a probabilistic approach to (Boutilier 1996) in future work.) However, regardless of the choice of A, the task is to determine a value for s, where applying Jeffrey's Rule yields new epistemic state $\mathcal{K}'$ wherein pw-belief set W_0' (i.e., the minimal P-stable set in $\mathcal{K}'$) corresponds to $\min([\phi], \preceq_\mathcal{K})$. We can think of this as constituting two constraints: first that $\min([\phi], \preceq_\mathcal{K})$ is P-stable in $\mathcal{K}'$ and second that $\min([\phi], \preceq_\mathcal{K})$ is the minimal P-stable set in $\mathcal{K}'$. We have the following constraints on P', and hence on s:

C1: For all $w \in \min([\phi], \preceq_\mathcal{K})$, we have $P'(w) > P'(W \setminus \min([\phi], \preceq_\mathcal{K}))$; and,

C2: For every set X' where $\emptyset \neq X' \subset \min([\phi], \preceq_\mathcal{K})$, there is $w' \in X'$ such that $P'(w') \leq P'(W \setminus X')$.

For C1, the last part is the same as $P'(w) > P'([\neg\phi]) + P'([\phi] \setminus \min([\phi], \preceq_\mathcal{K}))$. Using Jeffrey's Rule with $A = [\phi]$ having probability strictly between 0 and 1, this is the same as

$$\frac{s \cdot P(w)}{P([\phi])} > \frac{(1-s) \cdot P([\neg\phi])}{1 - P([\phi])} + \frac{s \cdot P([\phi] \setminus \min([\phi], \preceq_\mathcal{K}))}{P([\phi])}.$$

By arithmetic manipulation we get

$$s > \frac{P([\phi])}{P(w) + P(\min([\phi], \preceq_\mathcal{K}))}. \tag{4}$$

The maximum of the right-hand side occurs when w has minimum probability; call any such w to be $\min([\phi], \preceq_\mathcal{K})$-*minimal*.

Proposition 14 *For any $w \in \min([\phi], \preceq_\mathcal{K})$, there exists $s < 1$ such that* (4) *is satisfied. If w is $\min([\phi], \preceq_\mathcal{K})$-minimal, then $P([\phi])/(P(w) + P(\min([\phi], \preceq_\mathcal{K})))$ is the greatest lower bound of values of s satisfying* (4)*.*

The constraint C2 expresses that, if we try to shrink $\min([\phi], \preceq_\mathcal{K})$ to a smaller set X', then there is some member of X' that violates the purported P-stability of X'. We only need to consider those sets X' whose elements have probability P' greater than those in $\min([\phi], \preceq_\mathcal{K}) \setminus X'$ since, if there were a $w' \in X'$ where $P'(w') \leq P'(w'')$ for some $w'' \in \min([\phi], \preceq_\mathcal{K}) \setminus X'$, then w' will witness the failure of X' being P-stable.

To satisfy C2, for every such X' (where every element of X' has probability greater than every element of $\min([\phi], \preceq_\mathcal{K}) \setminus X'$), there must exist a $w' \in X'$ such that:

$$P'(w') \leq P'([\phi] \setminus X') + P'([\neg\phi]).$$

In terms of Jeffrey's Rule (1) with $A = [\phi]$, we can rewrite this as:

$$\frac{s \cdot P(w')}{P([\phi])} \leq \frac{s \cdot P([\phi] \setminus X')}{P([\phi])} + \frac{(1-s) \cdot P([\neg\phi])}{1 - P([\phi])}.$$

After simplification, we obtain:

$$s \leq \frac{P([\phi])}{P(w') + P(X')}. \tag{5}$$

(Compare with Equation 4.) For every such X', there must exist a $w' \in X'$ such that s is no greater than the right-hand side. In other words, s must be less than the smallest (among the X') of the largest (among the $w' \in X'$) possible right-hand side. For each such X', we can select $w' \in X'$ that has least probability among those elements of X'.

Consider the constraints C1 and C2, and the restrictions that they place on probability s (Equations (4) and (5)). The following example illustrates cases where these restrictions are and are not met.

Example 15 *Consider an epistemic state with $W = \{w_1, w_2, w_3, w_4\}$, $[\phi] = \{w_1, w_2, w_4\}$, as given in Table 1.*

The first three columns give what we need regarding the epistemic state. It can be verified that $W_0 = \{w_1, w_2, w_3\}$, $W_1 = \{w_4\}$, and that $\min([\phi], \preceq_\mathcal{K}) = W_0 \cap [\phi] = \{w_1, w_2\}$. Solving (4), we obtain that $s > \frac{0.75}{0.25+0.65} = \frac{0.75}{0.90} = 0.8\dot{3}$. If we take a value slightly less than this, $s = 0.82$ (violating (4)), we get the results in the fourth column. (In all cases we retain accuracy to three decimal places.) Unsurprisingly, in the new epistemic state we get $W_0 = \{w_1, w_2, w_3\}$. Choosing $s = 0.85 > \frac{0.75}{0.90}$ yields $W_0 = \{w_1, w_2\}$, as desired.

The second constraint C2 requires that, for every non-empty subset $X' \subset \min([\phi], \preceq_\mathcal{K})$, we have some $w' \in X'$ wherein $P'(w') \leq P'(W \setminus X')$. Solving (5) for our example, we obtain that $s \leq \frac{0.75}{0.4+0.4} = \frac{0.75}{0.80} = 0.9375$. If we take a value slightly less than this, $s = 0.92$, we obtain $\min([\phi], \preceq_\mathcal{K}) = \{w_1, w_2\}$, as desired. If we take a value slightly larger than this, $s = 0.95$ (violating (5)), we obtain $W_0 = \{w_1\}$ in the new epistemic state.

These considerations raise the question of whether there is always a probability s that satisfies both (4) and (5). Clearly there will be a problem if the right-hand side of (4) is not less than the right-hand side of (5), that is if

$$\frac{P([\phi])}{P(w') + P(X')} \leq \frac{P([\phi])}{P(w) + P(\min([\phi], \preceq_\mathcal{K}))}.$$

Simplifying, this will be the case exactly when

$$P(\min([\phi], \preceq_\mathcal{K}) \setminus X') \leq P(w') - P(w). \tag{6}$$

Proposition 16 *The following are equivalent:*

1. *There is an s that jointly satisfies C1 and C2.*
2. *For any $X' \subseteq \min([\phi], \preceq_\mathcal{K})$ and $w' \in X'$ and any $w \in W$ that is $\min([\phi], \preceq_\mathcal{K})$-minimal, $P(w') - P(w) < P(\min([\phi], \preceq_\mathcal{K}) \setminus X')$.*

Here is a simple example of the problem.

Example 17

		P
w_1	ϕ	0.5
w_2	ϕ	0.25
w_3	$\neg\phi$	0.25

Here there is one P-stable set, $\{w_1, w_2, w_3\}$ and so $w_i \preceq_\mathcal{K} w_j$ for each i, j. As well, $\min([\phi], \preceq_\mathcal{K}) = \{w_1, w_2\}$ and, for the lower bound in C1, $w = w_2$. Thus

$$s > \frac{0.75}{0.25 + 0.75} = 0.75$$

For C2, for the upper bound, choose $X' = \{w_1\}$ and $w' = w_1$. Thus

$$s \leq \frac{0.75}{0.5 + 0.5} = 0.75.$$

		P	$s = 0.82$	$s = 0.85$	$s = 0.92$	$s = 0.95$
w_1	ϕ	0.40	0.437	0.453	0.491	0.507
w_2	ϕ	0.25	0.273	0.283	0.307	0.317
w_3	$\neg\phi$	0.25	0.180	0.150	0.080	0.050
w_4	ϕ	0.10	0.109	0.113	0.123	0.127
W_0		w_1, w_2, w_3	w_1, w_2, w_3	w_1, w_2	w_1, w_2	w_1

Table 1: Example 15.

Together, $0.75 < s \leq 0.75$*, which cannot be satisfied.*

We turn next to the relation of this approach with accounts of belief revision. Before doing so, we consider the special case of revision by an inconsistent formula. In the AGM approach, revision by such a formula yields the inconsistent belief set, given by postulate (K*5). In a practical setting this is problematic since Jeffrey's Rule is inapplicable, and there is no elegant way to exit an incoherent epistemic state. Consequently, we suggest that perhaps the most natural outcome of revision by an inconsistent formula is to do nothing, that is, to leave the epistemic state unchanged, ignoring the inconsistent ϕ. This requires modification to two of the AGM postulates:

$(K * 2')$ If $\nvdash \neg\phi$, then $\phi \in K * \phi$
$(K * 5')$ If $\vdash \neg\phi$, then $K * \phi = K$

We obtain the following result.[11]

Theorem 18 *Define* $* : \mathcal{ES} \times \mathcal{L} \mapsto \mathcal{ES}$ *as follows. Let* $\mathcal{K} = \langle W, P\rangle$ *be an epistemic state,* $0.5 < s < 1.0$*, and* ϕ *a formula. Define* $\mathcal{K} * \phi = \langle W, P'\rangle$ *where:*

1. *if* $\vdash \neg\phi$*, then* $P' = P$*;*
2. *otherwise*

$$P'(w) = \begin{cases} P(w) \cdot s/P([\phi]) & \text{if } w \in [\phi] \\ P(w) \cdot (1-s)/(1-P([\phi])) & \text{if } w \notin [\phi] \end{cases}$$

Then:

1. $\mathcal{K} * \phi$ *satisfies (K*1), (K*2'), (K*4), (K*5'), (K*6), (K*8) if* s *satisfies (4).*
2. $\mathcal{K} * \phi$ *satisfies (K*3), (K*7) if* s *satisfies (4) and (5).*

Contrast this result with Theorem 12 wherein revision in the latter yields a set of formulas. Here revision yields a new epistemic state. As the first part shows, we may not always be able to satisfy all AGM postulates. In this case, informally, in revising by ϕ, the relative probability of ϕ worlds vs. $\neg\phi$ worlds may increase to the extent that the agent comes to believe that some worlds in $\min([\phi], \preceq_{\mathcal{K}})$ are not plausible candidates for the real world. This in turn means that, in a probabilistic setting, the notion of minimal change captured by (K*3) and (K*7) may be overly stringent. On the other hand, we have shown that (4) is always satisfiable, and one can always define a revision operator satisfying the postulates in Part 1 of Theorem 18. We note also that the other direction of Theorem 18 is given by Theorem 13.

We also obtain the following regarding the relation between probabilistic independence and revision:

Proposition 19 *1. If* ϕ *and* ψ *are independent in the epistemic state* $\mathcal{K}$*, then* ϕ *and* ψ *are independent in the epistemic state* $\mathcal{K} * \phi$*.*

2. *If* ϕ *and* ψ *are conditionally independent given* χ *in* $\mathcal{K}$*, then* ϕ *and* ψ *are conditionally independent given* χ *in* $\mathcal{K} * \phi$*.*
3. *If* ϕ *and* ψ *are conditionally independent given* χ *in* $\mathcal{K}$*, then* ϕ *and* ψ *are independent in* $\mathcal{K} * \chi$*.*

Consider a continuation of Example 4.

Example 20 *Consider revising by b. We have* $P([b]) = 0.165$ *and* $P(\min([b], \preceq_{\mathcal{K}})) = 0.1605$*. From (4) we have* $s > 0.165/(0.03 + 0.1605) \approx 0.866$*. For (5), we find the maximum denominator ranging over the following two* (X', w') *pairs:* $(\{\neg ab\neg f\}, \neg ab\neg f)$ *and* $(\{abf, \neg ab\neg f\}, abf)$*. Max occurs with pair* $(\{abf, \neg ab\neg f\}, abf)$*, and hence from (5) we have* $s \leq 0.165/(0.045 + 0.1305) \approx 0.94$*. Choosing* $s = 0.9$ *we obtain the following probabilities (correct to 4 decimal places).*

	a		$\neg a$	
	b	$\neg b$	b	$\neg b$
f	0.2455	0.0018	0.0245	0.0049
$\neg f$	0.1636	0.0012	0.4664	0.0923

Via Definition 10 the total preorder over sets of worlds $S_1 = \{abf, ab\neg f, \neg ab\neg f\}$*,* $S_2 = \{\neg a\neg b\neg f\}$*,* $S_3 = \{\neg abf\}$*,* $S_4 = \{\neg a\neg bf\}$*,* $S_5 = \{a\neg bf\}$*,* $S_6 = \{a\neg b\neg f\}$*.* P'*-stable sets are then given by* $\bigcup_{1 \leq j \leq i} S_j$ *for* $i = 1, \ldots, 6$*.*

Revising by b it is believed the agent has a sore back, and either has arthritis or else doesn't have swollen fingers.

Note (as per Proposition 19) that in the new set of probabilities $P'(b \mid af) = P'(b \mid a)$*. That is, in detail,* $P'(b \mid af) = P'(baf)/P'(af) = 0.2455/(0.2455 + 0.0018) = 0.9927$ *and* $P'(b \mid a) = P'(ba)/P'(a) = (0.2455 + 0.1636)/(0.2455 + 0.1636 + 0.0018 + 0.0012) = 0.9927$*.*

4.4 Revising an Epistemic State: Iterated Revision

We turn next to revision in the Darwiche and Pearl (1997) approach for iterated revision. They give the following postulates (with the convention described in Footnote 11).

$(DP1)$ If $\psi \vdash \phi$, then $(\mathcal{K} * \phi) * \psi \equiv \mathcal{K} * \psi$
$(DP2)$ If $\psi \vdash \neg\phi$, then $(\mathcal{K} * \phi) * \psi \equiv \mathcal{K} * \psi$
$(DP3)$ If $\phi \in \mathcal{K} * \psi$, then $\phi \in (\mathcal{K} * \phi) * \psi$
$(DP4)$ If $\neg\phi \notin \mathcal{K} * \psi$, then $\neg\phi \notin (\mathcal{K} * \phi) * \psi$

[11] An expression like $\mathcal{K} * \phi$ is an epistemic state. We continue the convention in (Darwiche and Pearl 1997) that, if an epistemic state occurs where a belief set is expected, the belief set corresponding to that epistemic state is to be understood. More accurate but more cumbersome would be to use something like $Bel(\mathcal{K} * \phi)$.

They also provide the following properties on worlds.

(DPS1) If $w_1 \models \phi$ and $w_2 \models \phi$,
then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K}*\phi} w_2$

(DPS2) If $w_1 \models \neg\phi$ and $w_2 \models \neg\phi$,
then $w_1 \preceq_{\mathcal{K}} w_2$ iff $w_1 \preceq_{\mathcal{K}*\phi} w_2$

(DPS3) If $w_1 \models \phi$ and $w_2 \models \neg\phi$,
then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \prec_{\mathcal{K}*\phi} w_2$

(DPS4) If $w_1 \models \phi$ and $w_2 \models \neg\phi$,
then $w_1 \preceq_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K}*\phi} w_2$

They then provide a result showing the equivalence of their postulates and conditions:

Theorem 21 (DP 1997) *Suppose a revision operator satisfies the AGM postulates. The operator satisfies postulates* $(DP1)-(DP4)$ *iff the operator and its corresponding faithful assignment satisfy* $(DPS1)-(DPS4)$ *respectively.*

Suppose that we have a revision operator that satisfies (DPS1)–(DPS4). In revising by a formula ϕ, informally the relative positions among ϕ worlds are unchanged, as are the relative positions among $\neg\phi$ worlds. Moreover (some or all of the) ϕ worlds are "moved" so that, in the resulting total preorder, the minimum worlds are $\min([\phi], \preceq_{\mathcal{K}})$.

It can be noted that (DPS1)–(DPS4) hold trivially if $w_1 \preceq_{\mathcal{K}} w_2$ is interpreted as $P(w_1) \leq P(w_2)$. For example, corresponding to (DPS1), it is obvious that if $w_1 \models \phi$ and $w_2 \models \phi$, then $P(w_1) \leq P(w_2)$ iff $P'(w_1) \leq P'(w_2)$ where P' is the probability assignment after applying Jeffrey's Rule with respect to $A = [\phi]$. However, this does not carry over to (DPS1)–(DPS4) expressed in terms of the total preorder $\preceq_{\mathcal{K}}$. In particular, the first two conditions need to be weakened to hold in our approach so that $\preceq_{\mathcal{K}}$ is replaced by $\prec_{\mathcal{K}}$ and the "iff" is replaced by an "if-then".[12]

Consider an addition to Example 15, where the probabilities in the third column of Table 1 are reduced by multiplying by a factor of $(1-\epsilon)$. We also add four worlds to the epistemic state, where these worlds have their probability scaled by ϵ. See Table 2, where X' is the minimum set of worlds among $\{w_5, w_6, w_7, w_8\}$ in the total preorder. The factor ϵ is chosen to be sufficiently small so that, in the total preorder, the worlds $w_1 - w_4$ are essentially independent of $w_5 - w_8$ with regards to a revision; here a couple of orders of magnitude would do. Thus choosing $s = 0.9$ will give us a revision operator that satisfies the AGM postulates. The extension to the ordering is then:

$$W_0 = \{w_1, w_2, w_3\}, W_1 = \{w_4\},$$
$$W_2 = \{w_5, w_6, w_7\}, W_3 = \{w_8\}.$$

After applying Jeffrey's Rule with $s = 0.8$, we obtain

$$W'_0 = \{w_1, w_2\}, W'_1 = \{w_4\}, W'_2 = \{w_3\},$$
$$W'_3 = \{w_5\}, W'_4 = \{w_7\}, W'_5 = \{w_8\}, W'_6 = \{w_6\}.$$

[12]This, in fact, may be an advantage. It's worth noting that DPS2 has been criticised as being overly strong. For example, if one revised by a formula such as $\phi = p_1 \wedge \cdots \wedge p_{1000}$ and then discovered that a fact was incorrect, $\psi = \neg p_1$, DPS2 requires that the revision by ϕ followed by ψ be the same as just revising by ψ. While this may be desirable in some cases, it seems overly strong to make such behaviour a requirement.

The set $\{w_5, w_6, w_7\}$ is P-stable wrt $\{w_5, w_6, w_7, w_8\}$. Hence, via Proposition 7.6, we have $w_7 \preceq_{\mathcal{K}} w_5$, even though $P(w_5) \neq P(w_7)$. Following revision using $s = 0.8$, we obtain the results of the last column of Table 2. Thus $w_7 \not\preceq_{\mathcal{K}'} w_5$ in the resulting epistemic state, violating DPS1.

		P	$s = 0.9$
w_5	ϕ	0.45ϵ	0.54ϵ
w_6	$\neg\phi$	0.25ϵ	0.10ϵ
w_7	ϕ	0.20ϵ	0.24ϵ
w_8	ϕ	0.10ϵ	0.12ϵ
X'		w_5, w_6, w_7	w_5

Table 2: DPS1 Counterexample.

Consider the following modified postulates:

(DPW1) If $w_1 \models \phi$ and $w_2 \models \phi$,
then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K}*\phi} w_2$

(DPW2) If $w_1 \models \neg\phi$ and $w_2 \models \neg\phi$,
then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K}*\phi} w_2$

(DPW3) If $w_1 \models \phi$ and $w_2 \models \neg\phi$,
then $w_1 \prec_{\mathcal{K}} w_2$ implies $w_1 \preceq_{\mathcal{K}*\phi} w_2$

These three postulates are simply "strict" versions of postulates (DP1) — (DP3). This leads to the following result.

Theorem 22 *Let $\mathcal{K}$ be an epistemic state and $\phi \in \mathcal{L}$ a formula. Then $\mathcal{K} * \phi$ satisfies (DPW1) – (DPW3).*

In other words, our account of belief change based on a probabilistic semantics utilising Jeffrey conditionalisation to update probabilities, also provides an account of iterated belief change that is consistent with the Darwiche-Pearl approach, in fact, a slightly stricter approach.

5 Conclusions and Future Work

This paper contributes to research in qualitative and quantitative probability. It explores a mechanism for qualitative belief change where the underlying epistemic state has a probabilistic semantics. Specifically, we use Jeffrey Conditionalisation to adjust probabilities assigned to possible worlds. In this way, we establish a correspondence with the hallmark AGM postulates for belief change and show under what conditions these postulates can be fully captured. We also establish a link with the Darwiche-Pearl approach to iterated belief change and provide formal results that put our work in the context of their approach. To our knowledge, this is the first work on qualitative probability providing an account of iterated belief change.

Several avenues for future work present themselves. There are various approaches to iterated belief change in the literature and we have addressed just the central one. This work can also be extended to give an account of iterated belief change for other types of belief change operators. Perhaps of most interest, there are extensive probabilistic knowledge bases (whether expressed as Bayesian networks, or whatever), and we suggest that our approach has significant potential for developing practical belief revision systems in probabilistic, uncertain domains.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2):510–530.

Benferhat, S.; Dubois, D.; and Prade, H. 1999. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation* 9(6):873–895.

Boutilier, C. 1995. On the revision of probabilistic belief states. *Notre Dame Journal of Formal Logic* 36(1):158–183.

Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic* 25:263–305.

Chellas, B. 1980. *Modal Logic*. Cambridge: Cambridge University Press.

Darwiche, A., and Pearl, J. 1994. On the logic of iterated belief revision. In Fagin, R., ed., *Proceedings of the Fifth Conference on the Theoretical Aspects of Reasoning about Knowledge*, 5–23.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.

Delgrande, J., and Jin, Y. 2012. Parallel belief revision: Revising by sets of formulas. *Artificial Intelligence* 176(1):2223–2245.

Delgrande, J. P.; Sack, J.; Lakemeyer, G.; and Pagnucco, M. 2022. Epistemic logic of likelihood and belief. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2599–2605.

Delgrande, J. P.; Renne, B.; and Sack, J. 2019. The logic of qualitative probability. *Artificial Intelligence* 275:457–486.

Delgrande, J. P. 2012. Revising beliefs on the basis of evidence. *International Journal of Approximate Reasoning* 53(3):396–412.

Gärdenfors, P. 1988. *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. Bradford/MIT, Cambridge, MA.

Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17(2):157–170.

Gärdenfors, P. 1975. Qualitative probability as an intensional logic. *Journal of Philosophical Logic* 4(2):171–185.

Halpern, J. Y. 2003. *Reasoning about Uncertainty*. MIT Press.

Hughes, G., and Cresswell, M. 1996. *A New Introduction to Modal Logic*. London and New York: Routledge.

Jeffrey, R. 1990. *The logic of decision*. University of Chicago Press, 2nd edition.

Jin, Y., and Thielscher, M. 2007. Iterated belief revision, revised. *Artificial Intelligence* 171:1–18.

Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.

Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12(5):773–808.

Kyburg, Jr., H. 1961. *Probability and the Logic of Rational Belief*. Wesleyan University Press.

Lehmann, D.; Magidor, M.; and Schlechta, K. 2001. Distance semantics for belief revision. *Journal of Symbolic Logic* 66(1):295–317.

Leitgeb, H. 2013. Reducing belief simpliciter to degrees of belief. *Ann. Pure Appl. Logic* 164(12):1338–1389.

Lindström, S., and Rabinowicz, W. 1989. On probabilistic representation of non-probabilistic belief revision. *Journal of Philosophical Logic* 11(1):69–101.

Nayak, A. C.; Foo, N. Y.; Pagnucco, M.; and Sattar, A. 1996. Changing conditional belief unconditionally. In *TARK*, 119–135.

Peppas, P. 2008. Belief revision. In van Harmelen, F.; Lifschitz, V.; and Porter, B., eds., *Handbook of Knowledge Representation*. San Diego, USA: Elsevier Science. 317–359.

Russell, S. J., and Norvig, P. 2010. *Artificial Intelligence: A Modern Approach*. Pearson Education, 3rd edition.

Snow, P. 1994. The emergence of ordered belief from initial ignorance. In *Proceedings of the National Conference on Artificial Intelligence*, 281–286.