A Unifying Framework for Probabilistic Belief Revision

Zhiqiang Zhuang[1], **James Delgrande**[2], **Abhaya Nayak**[3], **Abdul Sattar**[1]
[1] Institute for Integrated and Intelligent Systems, Griffith University, Australia
[2] School of Computing Science, Simon Fraser University, Canada
[3] Department of Computing, Macquarie University, Australia
z.zhuang@griffith.edu.au, jim@cs.sfu.ca, abhaya.nayak@mq.edu.au, a.sattar@griffith.edu.au

Abstract

In this paper we provide a general, unifying framework for probabilistic belief revision. We first introduce a probabilistic logic called p-logic that is capable of representing and reasoning with basic probabilistic information. With p-logic as the background logic, we define a revision function called p-revision that resembles partial meet revision in the AGM framework. We provide a representation theorem for p-revision which shows that it can be characterised by the set of basic AGM revision postulates. P-revision represents an "all purpose" method for revising probabilistic information that can be used for, but not limited to, the revision problems behind Bayesian conditionalisation, Jeffrey conditionalisation, and Lewis's imaging. Importantly, p-revision subsumes the above three approaches indicating that Bayesian conditionalisation, Jeffrey conditionalisation, and Lewis' imaging all obey the basic principles of AGM revision. As well our investigation sheds light on the corresponding operation of AGM expansion in the probabilistic setting.

1 Introduction

Since an agent acquires new information all the time, a key question is how such new information affects the agent's beliefs. This is the main subject of study in the area of *belief change* [Gärdenfors, 1988; Peppas, 2008]. The dominant approach in belief change is the *AGM framework* [Alchourrón *et al.*, 1985; Gärdenfors, 1988], which represents the agent's beliefs and input information as formulas of some background logic that subsumes classical logic.

In daily life, we most often deal with *uncertain* information, and quite often degrees of uncertainty are involved and play an essential role. For instance, we may decide to stay at home if the weather forecast says "the chance of a thunder storm is 90%" but we may seriously consider outdoor activities if the chance is 10%. More importantly, upon acquiring new information, the degree of uncertainty of our beliefs may change. For instance, the chance of an average American developing Type 1 diabetes by age 70 is 1%; however upon learning that the person has an immediate relative who has Type 1 diabetes, the chance rises to 10%–20%. Since the AGM framework does not take into account explicit measures of uncertainty, this kind of change is beyond its scope.[1] In this paper, we aim to develop a belief revision framework that deals with changes to degrees of uncertainty, among others.

To reach this aim, a crucial first step is to represent degrees of uncertainty properly. The most common representation of uncertainty in artificial intelligence is via probabilities, or more precisely *probability functions* which assign probabilities to propositional formulas. Also there are some established methods for incorporating new information into a probability function. The standard method is *Bayesian conditionalisation* or *conditionalisation* for short, which handles new information of the form "the probability of ϕ is 1," where ϕ represents some event. The best known generalisation of conditionalisation is *Jeffrey conditionalisation* [Jeffrey, 1965], which also handles new information with a degree of uncertainty such as "the probability of ϕ is 0.3." Neither conditionalisation nor Jeffrey conditionalisation is defined for the case where the probability of ϕ is initially zero. Lewis [1976] introduced a method called *imaging* that can handle such a "zero-probability" case. But imaging without further generalisation cannot handle new information with an associated degree of uncertainty.

We can identify at least two issues with the above methods. First, they assume that an agent's beliefs are represented by a single probability function, which means the agent has to know the exact probability for every event before any of the methods can be applied. In many cases, this is impossible, as very often we only have an estimation of the probability for a limited number of events. Second, each of these methods is limited in one way or another: either it handles only new information that is certain (i.e. imaging, conditionalisation) or only new information with non-zero initial probability (i.e. conditionalisation, Jeffrey conditionalisation).

We address both issues in our approach. The key is to deal with uncertainty through a probabilistic logic, which we call *p-logic*. Instead of probability functions, it is more natural and intuitive to work with individual probability statements such as "the probability of ϕ is 40%"; "the probability of ϕ is

[1] Arguably, the entrenchment based approach in the AGM framework [Gärdenfors and Makinson, 1988] captures some forms of uncertainty, but it neither represents explicit degrees of uncertainty nor changes to such degrees.

twice as likely as that of ψ"; or "the probability of ϕ is 20% more likely than that of ψ". P-logic, which originates from [Fagin *et al.*, 1990], is capable of representing and reasoning with such probability statements. As opposed to probability functions, p-logic can deal with incomplete information.

In our approach, we represent an agent's beliefs by a logically closed set of p-logic formulas called a *p-belief set*, which captures exactly the probabilistic information known to the agent. Then we define a variant of *partial meet revision* functions [Alchourrón *et al.*, 1985], called *p-revision*, for incorporating p-logic formulas into a p-belief set. P-revision is an "all purpose" revision method in that it can handle all forms of new information expressible by p-logic. P-revision is also well-behaved: we provide a representation theorem which shows that the class of p-revision functions can be characterised by the p-logic analogues of the basic AGM revision postulates.

An interesting outcomes of our approach is that we can show that each of the three methods, conditionalisation, Jeffrey conditionalisation, and imaging, is equivalent to some suitably-restricted class of p-revision functions. So, although these methods are devised in different contexts and were proposed long before the AGM framework, they all obey the basic principles of revision in the AGM framework. Moreover, our approach helps explicate the counterpart of AGM expansion in the probabilistic setting.

2 P-Logic

P-logic can be seen as a syntactically restricted version of the logic given in [Fagin *et al.*, 1990] for reasoning about probabilities.[2] Roughly speaking, a p-logic formula, called a *p-formula*, represents a constraint on the probabilities of some events.

Events are represented using formulas of classical propositional logic. Let $\mathcal{A}$ be a finite set of propositional atoms. The propositional language $\mathcal{L}$ is defined using the standard set of Boolean connectives, based on the atoms in $\mathcal{A}$ and the constants $\top$ and $\bot$. We write propositional atoms as lower case Roman letters $(a, b, c, \ldots)$ and propositional formulas as lower case Greek letters $(\phi, \psi, \ldots)$. Let Cn be the classical consequence operator. We denote a *propositional interpretation* or *possible world* as ω or μ possibly with a subscript, and sometimes as a bit vector such as $011\ldots$ which indicates that a is assigned false, b is assigned true, c is assigned true, and so on. For a set of propositional formulas S, we denote its set of models as $|S|$. For a propositional formula ϕ, $|\{\phi\}|$ is abbreviated to $|\phi|$. The set of all possible worlds is denoted Ω. If $\omega \in |\phi|$, we say ϕ is true in ω and ω is a ϕ-world.

With events represented by propositional formulas, we represent constraints on their probabilities with *atomic p-formulas* which take one of the following three forms:

1. $p(\phi) \bowtie t$
2. $p(\phi) \bowtie c \cdot p(\psi)$
3. $p(\phi) \bowtie p(\psi) + t$

[2]wherein all events are measurable.

where $\phi, \psi \in \mathcal{L}$, $\bowtie \in \{\leqslant, \geqslant, =\}$ and t, c are rational numbers such that $0 \leqslant t \leqslant 1$ and $0 < c$. A *p-formula* is a conjunction of atomic p-formulas (e.g., $(p(\phi) \geqslant 0.4) \wedge (p(\phi) = p(\psi) + 0.2)$). We write p-formulas as upper case Greek letters $(\Phi, \Psi, \ldots)$ and denote the set of all p-formulas as $\mathcal{L}_{\mathrm{P}}$. The three forms of atomic p-formulas are also referred as *Category 1*, *2*, and *3* p-formulas and constrain, respectively, the probability of a single event, the ratio between probabilities of two events, and the difference between probabilities of two events.[3] Our intention is that each category captures a specific type of commonly encountered constraints on probabilities of events.

Not all constraints are, however, covered here. For instance p-logic does not support inequalities involving more than two events (e.g., $p(\phi) + 3 \cdot p(\psi) \geqslant 2 \cdot p(\delta)$) or arbitrary Boolean combinations of inequalities (e.g., $(p(\phi) = 0.3) \vee (p(\phi) = 0.4)$). While restricted, p-formulas is in our opinion sufficient to capture most "commonsense" probability constraints. Also it will be clear that p-logic is more than enough for representing the revision process behind conditionalisation, Jeffrey conditionalisation, and imaging. Hence, the restricted syntax of p-logic demonstrates that not much formal machinery is needed from the probability standpoint to capture the three methods in a revision setting.

The basic semantic element for p-logic is a probability function. A probability function $P : \mathcal{L} \mapsto [0, 1]$ is a function that satisfies the Kolmogorov axioms:

- $0 \leqslant P(\phi) \leqslant 1$ for all $\phi \in \mathcal{L}$,
- $P(\top) = 1$, and
- $P(\phi \vee \psi) = P(\phi) + P(\psi)$ whenever $\neg\psi \in Cn(\phi)$.

We denote the set of all probability functions as $\mathcal{P}$. The letter P is reserved to denote a probability function throughout the paper. P *satisfies* an atomic p-formula

1. $p(\phi) \bowtie t$ iff $P(\phi) \bowtie t$;
2. $p(\phi) \bowtie\ c \cdot p(\psi)$ iff $P(\phi) \bowtie\ c \cdot P(\psi)$; and
3. $p(\phi) \bowtie p(\psi) + t$ iff $P(\phi) \bowtie P(\psi) + t$.

P *satisfies* a p-formula $\Phi \wedge \Psi$ iff it satisfies Φ and Ψ. P *satisfies* a set of p-formulas iff it satisfies all p-formulas in that set. Let X be a set of p-formulas and Φ a p-formula. If P satisfies X, then it is called a *p-model* of X. The set of p-models of X is denoted as $||X||$. We abbreviate $||\{\Phi\}||$ by $||\Phi||$. We say X (respectively Φ) is *consistent* iff $||X|| \neq \emptyset$ (respectively $||\Phi|| \neq \emptyset$); and X *entails* Φ under p-logic, denoted $X \models_{\mathrm{P}} \Phi$, iff $||X|| \subseteq ||\Phi||$. The *logical closure* of X, denoted $cl(X)$, is such that

$$cl(X) = \{\Phi \in \mathcal{L}_{\mathrm{P}} \mid X \models_{\mathrm{P}} \Phi\}.$$

A *p-belief set* B is a logically closed set of p-formulas, that is $B = cl(B)$. The letter B is reserved to denote a p-belief set throughout the paper.

In many cases, it is more convenient to work with probabilities of possible worlds instead of propositional formulas. So we write $P(\omega_1, \ldots, \omega_n)$ or $P(|\phi|)$ as a "shorthand" for

[3]Note that the p-formula $p(\phi) \bowtie p(\psi)$ is both a Category 2 and a Category 3 p-formula.

$P(\phi)$ where $|\phi| = \{\omega_1, \ldots, \omega_n\}$. Similarly, $p(\omega_1, \ldots, \omega_n)$ or $p(|\phi|)$ is a shorthand for $p(\phi)$ where $|\phi| = \{\omega_1, \ldots, \omega_n\}$.

The following example helps to illustrate the specifics of p-logic.

Example 1. *Let $\mathcal{A} = \{a, b\}$. Then $\Omega = \{11, 00, 10, 01\}$. Consider the Category 1 p-formula $p(11) \geqslant 0.5$ (shorthand for: $p(a \wedge b) \geqslant 0.5$). Suppose the probability function P is such that $P(11) = 0.7$, and $P(10) = P(01) = P(00) = 0.1$. Then P satisfies $p(11) \geqslant 0.5$ since $P(11) \geqslant 0.5$. Note that $p(11) \geqslant 0.5 \models_{\mathrm{P}} p(11) \geqslant c$ for any $c < 0.5$, since any P that satisfies $p(11) \geqslant 0.5$ has $P(11) \geqslant 0.5$ which means, by basic inequality, $P(11) \geqslant c$ for any $c < 0.5$.*

Working with logics like p-logic offers several advantages over probability functions. While a probability function is a complex structure that encapsulates the probabilistic information for all events, with p-logic we have the flexibility to deal with each piece of information separately. Also with p-logic, one can deal with incomplete information. Hence the probability of some events can be given without specifying the probability of all events. As well a constraint can be placed on an event (such as the fact that its probability is less than 0.4) without giving the exact probability. At the same time, one can clearly specify a complete probability function using only atomic p-formulas of any category.

3 P-Revision Functions

In this section, we consider how to revise a p-belief set by a p-formula. To begin, we define *probabilistic expansion* in the same way as *AGM expansion*. The probabilistic expansion of B by Φ, denoted $B + \Phi$, is such that

$$B + \Phi = cl(B \cup \{\Phi\}).$$

Probabilistic expansion gives trivial results when Φ is inconsistent with B. To devise a non-trivial method, we follow the construction of partial meet revision in the AGM framework.[4] In revising B by Φ, we first determine the maximal subsets of B that are consistent with Φ; we call these subsets the *remainder sets* of B with respect to Φ.

Definition 1. *The set of* remainder sets *of B with respect to Φ, denoted $B \downarrow \Phi$, is such that $X \in B \downarrow \Phi$ iff*

1. *$X \subseteq B$,*
2. *$X \cup \{\Phi\}$ is consistent, and*
3. *If $X \subset X' \subseteq B$, then $X' \cup \{\Phi\}$ is inconsistent.*

We then select some of the "best" remainder sets. The intersection of these "best" remainder sets is expanded by Φ to form the revision outcome. The decision on which remainder sets to select is modelled by a selection function. A function γ is a *selection function* for B iff $\gamma(B \downarrow \Phi)$ is a non-empty subset of $B \downarrow \Phi$, unless $B \downarrow \Phi$ is empty, in which case $\gamma(B \downarrow \Phi) = \emptyset$. Then p-revision functions are defined as follows.

[4]Strictly speaking, in the AGM framework, partial meet revision is constructed indirectly from a (partial meet) contraction via the Levi identity. Here we employ an equivalent, direct, construction of partial meet revision.

Definition 2. *A function $* : 2^{\mathcal{L}_{\mathrm{P}}} \times \mathcal{L}_{\mathrm{P}} \mapsto 2^{\mathcal{L}_{\mathrm{P}}}$ is a p-revision function iff*

$$B * \Phi = \bigcap \gamma(B \downarrow \Phi) + \Phi$$

where γ is a selection function for B.

For properties of p-revision function, consider the following postulates:

$(P*1)$ $B * \Phi$ is a p-belief set.

$(P*2)$ $\Phi \in B * \Phi$.

$(P*3)$ $B * \Phi \subseteq B + \Phi$.

$(P*4)$ If $B + \Phi$ is consistent, then $B + \Phi \subseteq B * \Phi$.

$(P*5)$ $B * \Phi$ is inconsistent iff Φ is inconsistent.

$(P*6)$ If $\Phi \equiv \Psi$, then $B * \Phi = B * \Psi$.

$(P*1)$ – $(P*6)$ are the p-logic versions of the six basic AGM revision postulates. We show that $(P*1)$ – $(P*6)$ fully characterise the class of p-revision functions.

Theorem 1. *A function $* : 2^{\mathcal{L}_{\mathrm{P}}} \times \mathcal{L}_{\mathrm{P}} \mapsto 2^{\mathcal{L}_{\mathrm{P}}}$ is a p-revision function iff it satisfies $(P*1)$ – $(P*6)$.*

For revising probabilistic information, p-revision has at least three advantages over the previously mentioned classical methods such as conditionalisation. First, an agent's beliefs are modelled as a p-belief set, which does not have to contain the probabilistic information for all events; therefore, p-revision does *not* require the agent to be "probabilistically omniscient" to start with. Second, new information is modelled as a p-formula; therefore, p-revision can accommodate more forms of new information than the classical methods. Thirdly, p-revision can handle the "zero-probability" cases that are undefined for conditionalisation and Jeffrey conditionalisation.

4 Equivalence with Conditionalisation, Jeffrey Conditionalisation, and Imaging

Due to the richness of p-logic, the problems behind many approaches to revising probabilistic information can be represented as sub-problems of those behind p-revision. In this section, we demonstrate that conditionalisation, Jeffrey conditionalisation, and imaging are each equivalent to specific p-revision functions.

4.1 Conditionalisation & Jeffrey Conditionalisation

Conditionalisation is one of the most common methods for revising probability functions. The conditionalisation of P on ϕ, denoted P^+_ϕ, is defined as

$$P^+_\phi(\psi) = \begin{cases} \dfrac{P(\phi \wedge \psi)}{P(\phi)} & \text{if } P(\phi) > 0 \\ \text{undefined} & \text{otherwise.} \end{cases}$$

The input ϕ is understood as "the probability of ϕ is 1." Jeffrey conditionalisation is the best known generalisation of conditionalisation. It deals with changes to a probability function induced by an input of the form "the probability of ϕ is c"

where $0 \leqslant c \leqslant 1$. For simplicity we write the input as $\phi = c$. The Jeffrey conditionalisation of a probability function P on $\phi = c$, denoted $P^J_{\phi=c}$, is defined as

$$P^J_{\phi=c}(\psi) = c \cdot P^+_\phi(\psi) + (1-c) \cdot P^+_{\neg\phi}(\psi)$$

in the general case, with special conditions attached to avoid division by zero.[5] It is easy to see that when $c = 1$ this reduces to conditionalisation.

Since we can represent a probability function P by a p-belief set B that has P as the only p-model, and the expression "the probability of ϕ is c" as the p-formula $p(\phi) = c$, the Jeffrey conditionalisation of P on $\phi = c$ corresponds to the problem of revising B by $p(\phi)=c$ where $||B|| = \{P\}$, which is a sub-problem for p-revision functions. We show that there are p-revision functions that yield an equivalent outcome to Jeffrey conditionalisation on this problem.

As its distinguishing feature, Jeffrey conditionalisation does not affect the probability ratios among formulas that imply ϕ and among those that imply $\neg\phi$.

Lemma 1. *If $P(\alpha) > 0$, $P(\beta) > 0$ and either both $\alpha \models \phi$ and $\beta \models \phi$ or both $\alpha \models \neg\phi$ and $\beta \models \neg\phi$, then*

$$\frac{P^J_{\phi=c}(\alpha)}{P^J_{\phi=c}(\beta)} = \frac{P(\alpha)}{P(\beta)}.$$

The key here is to formulate this ratio-preserving feature as a restriction to the selection functions for p-revision. To this end, we introduce the notion of *ratio-formulas*.

Definition 3. *A p-formula of Category 2, $p(\alpha) = c \cdot p(\beta)$, is a* ratio-formula *for ϕ iff $0 < c$ and either both $\alpha \models \phi$ and $\beta \models \phi$ or both $\alpha \models \neg\phi$ and $\beta \models \neg\phi$.*

A ratio-formula for ϕ is a Category 2 p-formula that describes the ratios between probabilities of propositional formulas that imply ϕ or those that imply $\neg\phi$. So if $\alpha \models \phi$ but $\beta \not\models \phi$, then $p(\alpha) = c \cdot p(\beta)$ is not a ratio-formula for ϕ. In revising B by $p(\phi) = c$ where $||B|| = \{P\}$, in order to keep the profile of probability ratios for ϕ and $\neg\phi$ untouched, we have to preserve all ratio-formulas for ϕ in B, which we denote as $R_\phi(B)$. We can show that $R_\phi(B)$ together with the new information $p(\phi) = c$, give us the Jeffrey conditionalisation of P on $\phi = c$.

Lemma 2. *If $||B|| = \{P\}$, then*

$$||R_\phi(B) \cup \{p(\phi) = c\}|| = \{P^J_{\phi=c}\}.$$

In relation to p-revision, there is a remainder set of B with respect to $p(\phi) = c$ that contains $R_\phi(B)$.

Lemma 3. *Let $||B|| = \{P\}$. Then there exists some element X in $B \downarrow (p(\phi) = c)$ such that $R_\phi(B) \subseteq X$.*

Thus, according to Lemma 2, if the selection function for B picks a single remainder set that contains $R_\phi(B)$, then the p-revision function yields an equivalent outcome to Jeffrey conditionalisation.

[5] Here we follow [Halpern, 2005] and require that if $c > 0$, then $P(\phi) > 0$ and similarly if $1-c > 0$, then $P(\neg\phi) > 0$. Also $c \cdot P^+_\phi(\psi)$ is taken to be 0 if $c = 0$, even if $P(\phi) = 0$. Similarly $(1-c) \cdot P^+_{\neg\phi}(\psi)$ is taken to be 0 if $1-c = 0$, even if $P(\neg\phi) = 0$.

Theorem 2. *Let $||B|| = \{P\}$. Then there is a p-revision function $*$ such that*

$$||B * (p(\phi) = c)|| = \{P^J_{\phi=c}\}.$$

Since Jeffrey conditionalisation generalises conditionalisation, the equivalence result also applies to conditionalisation. The following example illustrates the equivalence between a p-revision function and Jeffrey conditionalisation.

Example 2. *Let $\Omega = \{11, 00, 10, 01\}$, $|\phi| = \{11, 00\}$, and P be a probability function such that $P(11) = P(10) = P(01) = 0.2$ and $P(00) = 0.4$. Then the Jeffrey conditionalisation of P on the input that "the probability of ϕ is 0.3," denoted $P^J_{\phi=0.3}$, is such that $P^J_{\phi=0.3}(10) = 0.35$, $P^J_{\phi=0.3}(01) = 0.35$, $P^J_{\phi=0.3}(11) = 0.1$, and $P^J_{\phi=0.3}(00) = 0.2$.*

The Jeffrey conditionalisation corresponds to the revision of the p-belief set B by the p-formula $p(\phi) = 0.3$ where $||B|| = \{P\}$. The set of ratio-formulas for ϕ in B, denoted $R_B(\phi)$, is such that $R_B(\phi) = \{p(10) = p(01), p(10, 01) = 2 \cdot p(01), p(10, 01) = 2 \cdot p(10), p(11) = 2 \cdot p(00), p(11, 00) = 3 \cdot p(11), p(11, 00) = 1.5 \cdot p(00)\}$. Note that $|\neg\phi| = \{10, 01\}$. Since $R_B(\phi)$ is consistent with $p(\phi) = 0.3$, there is a remainder set of B with respect to $p(\phi) = 0.3$ that contains $R_B(\phi)$.

Let $*$ be a p-revision function where the selection function γ for B is such that $\gamma(B \downarrow p(\phi) = 0.3) = \{X\}$ and $R_B(\phi) \subseteq X$. Then we have $B * (p(\phi) = 0.3) = cl(X \cup \{p(\phi) = 0.3\})$. Note that $p(11) = 2 \cdot p(00)$ together with $p(\phi) = 0.3$ entail $p(11) = 0.1$ and $p(00) = 0.2$. Also $p(\phi) = 0.3$ entails $p(\neg\phi) = 0.7$; and $p(10) = p(01)$ together with $p(\neg\phi) = 0.7$ entail $p(10) = 0.35$ and $p(01) = 0.35$. So the only probability function that satisfies $B * (p(\phi) = 0.3)$ is $P^J_{\phi=0.3}$.

4.2 Imaging

Imaging, introduced in [Lewis, 1976], is the starting point for many works on probabilistic belief revision [Ramachandran *et al.*, 2010; Chhogyal *et al.*, 2014; Rens *et al.*, 2016]. In contrast to Jeffrey conditionalisation, imaging gives non-trivial results for the "zero-probability" case. Lewis makes the assumption that for each possible world ω and each consistent propositional formula ϕ, there is a ϕ-world, denoted ω_ϕ, that is closest to ω among the ϕ-worlds. The *image* of a probability function P on ϕ is obtained by shifting the original probability of each world ω over to ω_ϕ. As in conditionalisation, the input ϕ is understood as "the probability of ϕ is 1."

While there is only one way to do Jeffrey conditionalisation, there are many ways to do imaging; and each is determined by how, for each consistent $\phi \in \mathcal{L}$ and each $\omega \in \Omega$, ω is assigned a closest ϕ-world. According to [Lewis, 1976], the only restriction for such an assignment is that if $\omega \in |\phi|$, then $\omega_\phi = \omega$, that is the closest ϕ-world to any ϕ-world is the ϕ-world itself. We model the assignment of closest world by a function $I : \Omega \times \mathcal{L} \mapsto \Omega$ such that $I(\omega, \phi)$ gives us the closest ϕ-world to ω and

1. $I(\omega, \phi) \in |\phi|$ and
2. $I(\omega, \phi) = \omega$ whenever $\omega \in |\phi|$.

We call such functions *closest world functions*.

Now the imaging process can be captured precisely by a *image function* defined as follows.

Definition 4. *A function $\circ : \mathcal{P} \times \mathcal{L} \mapsto \mathcal{P}$ is an image function iff*

$$P^{\circ}_{\phi}(\omega) = \begin{cases} \sum\limits_{\substack{\mu \in \Omega \\ \omega = I(\mu,\phi)}} P(\mu) & \text{if } \omega \in |\phi| \\ 0 & \text{otherwise} \end{cases}$$

where the probability function P°_{ϕ} is the outcome of $\circ$ on P and ϕ, and I is a closest world function. We say $\circ$ is determined by I.

From now on, we work with image functions rather than the general notion of imaging, as the former is precise on how the imaging is done.

Similar to the case of Jeffery conditionalisation, we can represent a probability function P by a p-belief set B that has P as the only p-model; and the expression "the probability of ϕ is 1" as the p-formula $p(\phi) = 1$. So for an image function $\circ$, the problem behind $P \circ \phi$ corresponds to the problem of revising B by $p(\phi) = 1$ where $||B|| = \{P\}$; this is again a sub-problem of p-revision functions. We next show the equivalence of some p-revision functions with the image function.

The imaging process can be characterised by a number of probability shifts, one shift for each world ω that "benefits" from the probability movement in the imaging process. We can picture such a shift for a beneficiary world ω as consisting of two steps: (a) identify all worlds, including ω, whose (relevantly) closest world is ω, and (2) shift their combined probability mass to ω. The key here is to formulate such probability shifts as restrictions on the selection functions underlying p-revision. To this end, we introduce the notion of a *shift-formula*.

Definition 5. *A p-formula of Category 1, $p(\alpha) = c$, is a* shift-formula *for ϕ iff $|\alpha| \cap |\phi|$ is a singleton.*

The intention here is that, given the desired outcome $P(\phi) = 1$, for each ϕ-world ω we identify a sentence α such that $|\alpha|$ contains exactly those worlds each of which has ω as its closest world among $|\phi|$. It is easily noted that, in such a case, $|\alpha| \cap |\phi| = \{\omega\}$ is a singleton. The resultant probability of ω (due to imaging) is the total probability it will receive from members of $|\alpha|$, and c is intended to capture that total. Since we have a shift-formula $p(\alpha) = c$ for every ϕ-world ω, we can capture every shift of probability behind $P \circ \phi$. We accordingly define the *shift-set* for ϕ (with respect to P and I) as the set of such relevant shift-formulas:

Definition 6. *The* shift-set *for ϕ with respect to P and I, denoted $S^I_P(\phi)$, is a set of shift formulas. A shift formula $p(\alpha) = c$ is in $S^I_P(\phi)$ iff there exists $\omega \in |\phi|$ such that*

1. *$|\alpha| = \{\mu \in \Omega \mid I(\mu, \phi) = \omega\}$, and*
2. *$c = P(\alpha)$.*

The following lemma shows that our intuition on capturing the shift of probability through a shift-set is correct. That is, P°_{ϕ}, where the image function $\circ$ is determined by I, is the only p-model of $S^I_P(\phi)$ and $p(\phi) = 1$.

Lemma 4. *Let $\circ$ be an image function that is determined by I. Then*

$$||S^I_P(\phi) \cup \{p(\phi) = 1\}|| = \{P^{\circ}_{\phi}\}.$$

In relation to p-revision, we have that if B is a p-belief set that has P as the only p-model, then there is a remainder set of B with respect to $p(\phi) = 1$ that contains $S^I_P(\phi)$.

Lemma 5. *Let $||B|| = \{P\}$. Then there exists element X in $B \downarrow (p(\phi) = 1)$ such that $S^I_P(\phi) \subseteq X$.*

It follows then from Lemma 4 that if the selection function for B picks a single remainder set that contains $S^I_P(\phi)$, then the p-revision function yields an outcome equivalent to the imaging-function $\circ$.

Theorem 3. *Let $\circ$ be an image function, and $||B|| = \{P\}$. Then there is a p-revision function $*$ such that*

$$||B * (p(\phi) = 1)|| = \{P^{\circ}_{\phi}\}.$$

The following example illustrates the equivalence between a p-revision function and an image function.

Example 3. *Let $\Omega = \{11, 00, 10, 01\}$ and $|\phi| = \{11, 00\}$. Let P be a probability function such that $P(11) = P(00) = 0$, $P(10) = 0.3$, and $P(01) = 0.7$. Let $\circ$ be an image function where the associated closest world function I is such that $I(10, \phi) = 11$ and $I(01, \phi) = 00$. Then we have $P^{\circ}_{\phi}(11) = 0.3$, $P^{\circ}_{\phi}(00) = 0.7$, $P^{\circ}_{\phi}(10) = 0$, and $P^{\circ}_{\phi}(01) = 0$.*

The problem behind $P \circ \phi$ corresponds to the the revision of the p-belief set B by the p-formula $p(\phi) = 1$ where $||B|| = \{P\}$. The shift-set for ϕ with respect to P and I, denoted $S^I_P(\phi)$, is such that $S^I_P(\phi) = \{p(01, 00) = 0.7, p(10, 11) = 0.3\}$. Since $S^I_P(\phi)$ is consistent with $p(\phi) = 1$, there is a remainder set of B with respect to $p(\phi) = 1$ that contains $S^I_P(\phi)$.

Let $*$ be a p-revision function where the selection function γ for B is such that $\gamma(B \downarrow p(\phi) = 1) = \{X\}$ and $S^I_P(\phi) \subseteq X$. Then we have $B * (p(\phi) = 1) = cl(X \cup \{p(\phi) = 1\})$. Note that since $|\phi| = \{00, 11\}$, we have $p(01, 00) = 0.7$ together with $p(\phi) = 1$ entail $p(00) = 0.7$; $p(10, 11) = 0.3$ together with $p(\phi) = 1$ entail $p(11) = 0.3$; and $p(\phi) = 1$ entails $p(10) = 0$ and $p(01) = 0$. So the only probability function that satisfies $B * (p(\phi) = 1)$ is P°_{ϕ}.

5 A Note on Probabilistic Expansion

It is generally accepted, starting with [Gärdenfors, 1988], that the probabilistic analogue of AGM expansion is conditionalisation. In both the cases, exposure to evidence that does not contravene current knowledge leads to a new consistent belief state which incorporates both the old and the new knowledge; however, if the evidence "contradicts" old knowledge, the agent lands in the "epistemic hell". In the case of AGM expansion the epistemic hell is the inconsistent set $Cn(\perp)$, and in case of conditionalisation it is an undefined belief state. Neither of these approaches employs any mechanism to resolve inconsistency. Note in particular that, in the probability setting, a piece of evidence with non-zero prior does not contravene current knowledge since there is at least one world with a non-zero probability mass that satisfies the evidence.

However, when a probabilistic belief state P is represented by a p-logic belief set B, the picture we get of conditionalisation alters drastically. Consider the conditionalisation of P by ϕ when $P(\phi) = 0.3$. Here the conditionalisation outcome is properly defined, and hence would be considered a case of probabilistic expansion. In the context of p-logic, however, this conditionalisation would correspond to incorporation of the evidential sentence $p(\phi) = 1$ into B where, presumably, B has $p(\phi) = 0.3$ as a member. Since no probability function can satisfy both $p(\phi) = 0.3$ and $p(\phi) = 1$, if we were to employ "expansion", we would expand to inconsistency. However conditionalisation demands otherwise, and what corresponds to conditionalisation in the context of p-logic is instead a proper p-revision function that involves resolving inconsistency. Hence we contend that it is our probabilistic expansion, namely a set union operation follow by logical closure, and not conditionalisation, that is the appropriate formulation of AGM expansion in a probabilistic setting.

We note that Voorbraak [1999] also claims that conditionalisation is different from AGM expansion, and his main argument is that AGM expansion aims to reduce ignorance whereas conditionalisation aims to reduce uncertainty.

6 Related Work

In contrast to p-revision, most approaches to probabilistic belief revision deal with changes to a single probability function induced by input of the form "the probability of ϕ is 1".

Apart from the classical methods mentioned previously, Gärdenfors [1988] also provides a revision method for probability functions that adheres to the AGM framework. The main idea is that, in revising P by ϕ, first pick a probability function Q such that $Q(\phi) > 0$ then take the conditionalisation of Q on ϕ (i.e., Q^+_ϕ) as the revision outcome. In addition to revision, Gärdenfors also discusses the corresponding operations of AGM contraction and expansion for probability functions. The approach by Chhogyal *et al.* [2014] is based on imaging. The authors propose various ways of assigning the closest worlds that are intuitively appealing, and evaluate their methods against some AGM-style revision postulates. Also based on the idea of imaging, the corresponding operation of AGM contraction is studied in [Ramachandran *et al.*, 2010] for probability functions.

An approach that is closer to ours is by Rens *et al.* [2016] which works with a syntactically-restricted version of p-logic. To be precise, their language is restricted to a subset of Category 1 p-formulas. Since a set of such p-formulas may have more than one satisfying probability function, an agent's beliefs are modelled as a set of probability functions. They focus on a generalisation of imaging called *general imaging* [Gärdenfors, 1988]. While imaging assumes that a possible world ω has a single most-similar ϕ-world (viz. ω_ϕ), general imaging allows multiple most-similar worlds. In revising $\{P_1, \ldots, P_n\}$ by ϕ, the approach of Rens *et al.* [2016] is to obtain the (general) image of P_i on ϕ for $1 \leqslant i \leqslant n$, then take the set of imaging outcomes as the revision outcome. Frequently there is an infinite set of satisfying probability functions. Their main contribution is the identification of a finite set of boundary probability functions among the infinite set, such that revision of the finite set gives an identical outcome to revision of the infinite set. Note that although their logic allows one to express new information with a degree of uncertainty, their revision method, like imaging, handles only new information of the form "the probability of ϕ is 1". Also concerning changes to a set of probability functions, Grove and Halpern [1998] articulate some desirable properties and evaluate several revision methods over the properties. The methods they considered however can not handle the "zero-probability" case.

In [Boutilier, 1995] an agent's beliefs are captured by a Popper function [Popper, 2002] that takes the notion of conditional probability as the primitive. In this setting, Boutilier [1995] discusses issues of iterated belief revision. In [Delgrande, 2012], an agent's belief is modelled as a tuple that consists of a probability function and a confidence level c. Then methods for incorporating new and uncertain information with a confidence level exceeds c are introduced. Lindström and Rabinowicz [1989] discussed ways of dealing with the *non-uniqueness problem*, that is a belief set can be associated with many different probability functions. Kern-Isberner [2001] brings conditionals and entropy measures into consideration for probabilistic belief revision. Bona *et al.* [2016] investigate ways of consolidating probabilistic information through belief base contraction [Hansson, 1999]. Noticeably, their approach is based on a syntactically-restricted version of p-logic (i.e., subset of Category 1 p-formulas). Chhogyal *et al.* [2015] give a concrete construction of Gärdenfors' contraction method through argumentation.

7 Conclusion

In this paper, we have proposed p-logic, which is capable of representing and reasoning with some commonly encountered probability assertions. This allows us to deal with uncertainty in a more natural and familiar way. This also allows us to refer to each item of probabilistic information separately, which makes it easy to represent the change of information when our beliefs represented as p-formulas are revised.

With p-logic as the basis for dealing with uncertainty, we proposed p-revision, which is an "all purpose" revision method that complies with all the basic AGM revision postulates. We show that p-revision subsumes conditionalisation, Jeffrey conditionalisation, and imaging, which are the classical methods for revising probability functions. Significantly, the result implies that although these classical methods were introduced much earlier than the AGM framework, they all obey the basic principles of AGM revision.

We note that Category 3 p-formulas are redundant in establishing the correspondence of p-revision with the classical methods. Our results show that Category 2 p-formulas are sufficient to capture the revision process of Jeffrey conditionalisation and Category 1 p-formulas alone are sufficient for imaging. Whether there are meaningful revision methods whose representation requires Category 3 p-formulas is an interesting question that we will explore in our future work.

Acknowledgements

We thank Gavin Rens for useful comments on an earlier draft of this paper. This research has been supported by the Australian Research Council (ARC), Discovery Project: DP150104133, and the Natural Sciences and Engineering Research Council of Canada.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Bona *et al.*, 2016] Glauber De Bona, Marcelo Finger, Márcio Moretto Ribeiro, Yuri David Santos, and Renata Wassermann. Consolidating probabilistic knowledge bases via belief contraction. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR-2016)*, pages 125–134, 2016.

[Boutilier, 1995] Craig Boutilier. On the revision of probabilistic belief states. *Notre Dame Journal of Formal Logic*, 36, 1995.

[Chhogyal *et al.*, 2014] Kinzang Chhogyal, Abhaya C. Nayak, Rolf Schwitter, and Abdul Sattar. Probabilistic belief revision via imaging. In *Proceedings of the 13th Pacific Rim International Conference on Artificial Intelligence (PRICAI-2014)*, pages 694–707, 2014.

[Chhogyal *et al.*, 2015] Kinzang Chhogyal, Abhaya C. Nayak, Zhiqiang Zhuang, and Abdul Sattar. Probabilistic belief contraction using argumentation. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2015)*, pages 2854–2860, 2015.

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

[Fagin *et al.*, 1990] R. Fagin, Joseph Y. Halpern, and Nimrod Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 87(1):78–128, 1990.

[Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd conference on Theoretical Aspects of Reasoning about Knowledge (TARK-1988)*, pages 83–95, 1988.

[Gärdenfors, 1988] Peter Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. MIT Press, 1988.

[Grove and Halpern, 1998] Adam J. Grove and Joseph Y. Halpern. Updating sets of probabilities. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-1998)*, pages 173–182, 1998.

[Halpern, 2005] Joseph Y. Halpern. *Reasoning about uncertainty*. MIT Press, 2005.

[Hansson, 1999] Sven Ove Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer, 1999.

[Jeffrey, 1965] Richard Jeffrey. *The Logic of Decision*. New York: McGraw-Hill, 1965.

[Kern-Isberner, 2001] Gabriele Kern-Isberner. Revising and updating probabilistic beliefs. In Mary-Anne Williams and Hans Rott, editors, *Frontiers in Belief Revision*, pages 393–408. Springer Netherlands, 2001.

[Lewis, 1976] David Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85:279–315, 1976.

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

[Peppas, 2008] Pavlos Peppas. Belief revision. In *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, 2008.

[Popper, 2002] Karl Raimund Popper. *The Logic of Scientific Discovery*. Routledge, 2002. 1st English Edition:1959.

[Ramachandran *et al.*, 2010] Raghav Ramachandran, Abhaya C. Nayak, and Mehmet A. Orgun. Belief erasure using partial imaging. In *Proceedings of the 23rd Australasian Joint Conference (AI-2010)*, pages 52–61, 2010.

[Rens *et al.*, 2016] Gavin Rens, Thomas Andreas Meyer, and Giovanni Casini. On revision of partially specified convex probabilistic belief bases. In *Proceedings of 22nd European Conference on Artificial Intelligence (ECAI-2016)*, pages 921–929, 2016.

[Voorbraak, 1999] Frans Voorbraak. Probabilistic belief change: Expansion, conditioning and constraining. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI-1999)*, pages 655–662, 1999.