

A New Approach to Probabilistic Belief Change

Gavin Rens and Thomas Meyer

Centre for Artificial Intelligence Research,
University of KwaZulu-Natal, and CSIR Meraka, South Africa.
{grens, tmeyer}@csir.co.za

Abstract

One way for an agent to deal with uncertainty about its beliefs is to maintain a probability distribution over the worlds it believes are possible. A belief change operation may recommend some previously believed worlds to become impossible and some previously disbelieved worlds to become possible. This work investigates how to redistribute probabilities due to worlds being added to and removed from an agent’s belief-state. Two related approaches are proposed and analyzed.

1 Introduction

Suppose an agent maintains its beliefs in a knowledge-base (KB) of sentences. An agent’s *possible worlds* are simply the models (in terms of logic) of its current KB. In this work, we assume that an agent deals with uncertainty about its current situation by maintaining a probability distribution over possible worlds. When we say that a (possible) world ω has a probability p , we shall mean that the agent believes that there is a probability p that ω is the *actual* world. For instance, you may be in a situation where you care about exactly three things: whether you will receive a call on your cellphone in the next two minutes, whether the phone’s battery will last another two minutes on standby and whether you can reach a charging point within two minutes. There are eight ways for these three things to be true and false and there are various probability distributions one could attach to the eight worlds. We shall use this cellphone scenario as a running example throughout the paper.

We shall refer to the operations of belief revision, expansion, contraction and update collectively as *belief change* (Van Harmelen, Lifshitz, and Porter, 2008, e.g.). Due to some belief change operation on a KB, the KB may gain some models and it may lose some others. Then, in terms of possible worlds, some worlds may be added due to a belief change operation and some may be removed from the set of worlds thought possible (represented by the KB).¹

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We shall not restrict how the information contained in the KB is represented, as long as one can derive which worlds are possible and their assigned probabilities.

In this paper, we suggest a new unified approach to redistribute the probability mass of worlds in a KB resulting from any belief change operation.

Related Work “Lewis, in the context of providing semantics for conditionals, distinguished conditional probability from the probability of conditionals (Lewis, 1976). While employing Bayesian conditioning for the former, he devised imaging based on similarity or closeness between worlds to deal with the latter,” (Chhogyal et al., 2014). In particular, imaging solves the problem of update when the evidence Ψ is contradictory to current beliefs P (i.e., $P(\Psi) = 0$). “Lewis’ use of imaging based on closeness between possible worlds offers a way to overcome this limitation in the context of belief update (in a dynamic environment),” (Chhogyal et al., 2014).

Gärdenfors (1988) proposes a method which retains the basic idea behind imaging but that is also preservative² “However this approach is nonconstructive, and does not really provide a way for performing probabilistic belief revision,” (Chhogyal et al., 2014).

Chhogyal et al. (2014) explore the use of imaging as a means to construct probabilistic belief *revision*, while staying close to the spirit of Gärdenfors (1988)’s approach to probabilistic belief change. Specifically, they present explicit constructions of three candidate strategies to redistribute probabilities after a belief revision operation.

Although our approach makes use of the notion of similarity or closeness between worlds, we do not use imaging (with the assumption of a unique closest world for each world). Our approach is to look at *all* the currently believed worlds to inform what probabilities to assign to each world in the new set of believed worlds, after a belief change operation. Furthermore, whereas Gärdenfors (1988) and Chhogyal et al. (2014) focus on revision, our approach can be applied to belief revision *and* belief update.

Boutilier (1995) also presents a method for belief revision, however, in the setting where probabilistic epistemic states are represented by *Popper functions*.

²A probabilistic revision function $*$ is said to be preservative iff, for all probability functions P and for all propositions Ψ and Φ , if $P(\Psi) > 0$ and $P(\Phi) = 1$, then the probability of Φ is still 1 after $*$ is applied to P on Ψ . (Gärdenfors, 1988, p. 115)

Kern-Isberner (2008) applies the principle of minimum cross entropy to find the new distribution after belief revision or update. Furthermore, her work includes the notion of conditional beliefs.

Some notation and basic definitions are covered in Section 2. Section 3 lays the foundation of how to redistribute probabilities when worlds are either added or removed from the current set of possible worlds. We provide the details of our proposed framework for probabilistic belief change in Section 4, using the definitions developed in Section 3. We end with a discussion of our approach and pointers to future research. Due to limited space, all proofs are omitted.

2 Preliminaries

We'll work with propositional logic. Let $\mathcal{P} = \{a_1, a_2, \dots, a_n\}$ be the set of atomic propositional variables (*atoms*, for short). Formally, a *world* is a unique assignment of truth values to all the atoms in \mathcal{P} . There are thus 2^n conceivable worlds. An agent may consider some subset Ω of the conceivable worlds called the possible worlds.

In this work, we shall have nothing to say about exactly which worlds should be added and removed to accomplish a belief change operation. We assume that some well defined approaches are used. For instance, in the case of belief contraction and revision, the AGM approach (Alchourrón, Gärdenfors, and Makinson, 1985) may be used and for belief update, the approach of Katsuno and Mendelzon (1991) may be used.

We use a pseudo-distance measure between worlds, as defined by Lehmann, Magidor, and Schlechta (2001) and adopted by Chhogyal et al. (2014)

Definition 2.1 A pseudo-distance function $d : \Omega \times \Omega \rightarrow \mathbb{Z}$ satisfies the following four conditions: for all worlds $\omega, \omega', \omega'' \in \Omega$,

1. $d(\omega, \omega') \geq 0$ (*Non-negativity*)
2. $d(\omega, \omega) = 0$ (*Identity*)
3. $d(\omega, \omega') = d(\omega', \omega)$ (*Symmetry*)
4. $d(\omega, \omega') + d(\omega', \omega'') \geq d(\omega, \omega'')$ (*Triangular Inequality*)

The notions of distance between worlds and the worlds' probability should not be confounded; these are two different, but related notions: From the perspective of some world ω , an agent is assumed to consider some (usually other) world ω' in Ω as having a degree of similarity, and it is further assumed that every world in Ω can be compared to every other world in Ω by its degree of similarity from the perspective of ω . This is what the pseudo-distance function does. Any given pseudo-distance function is fixed for a given set of worlds.

A world's probability is also fixed given the current situation in which the agent finds itself. However, the likelihood of worlds depends on the agent's past actions and perceptions. And if information is received or if the situation changes, then the probability mass should be shifted around while considering the relationships between worlds. That is, probabilities are fixed for a particular situation, but may need to be amended due to new information or a new situation.

Suppose you believe that exactly two worlds are possible, ω_1 with probability 0.9 and ω_2 with probability 0.1. If you are told to remove only ω_1 due to some higher level reasoning, it might seem wrong, because ω_2 is the less likely world. In this work, we take the non-probabilistic aspect and not the stochastic aspect of an agent's beliefs to be primary. That is, deciding which worlds an agent should consider possible/believable is assumed to be based only on non-probabilistic considerations³. Only after a decision is made about which worlds (if any) to add and remove, is the probability distribution determined. We do not argue that logical and stochastic considerations *must* be separate; we separate the two aspects in this work only to simplify our analyses. Future work may include stochastic considerations when deciding which worlds to believe possible during a belief change operation.

In this work, as in the work of Lewis (1976) and Chhogyal et al. (2014), probabilistic belief change entails two steps: (1) determine which worlds to add and remove from the current belief-state, then (2) re-distribute the probability mass over the new set of possible worlds (if a change occurred).

A KB can be written as a single sentence k . Let $[k]$ be the set of models of k : $[k] = \{\omega \in \Omega \mid \omega \models k\}$. An expansion operation on k by information (sentence) Ψ (conventionally denoted $k + \Psi$) has the consequence (in semantic terms) of removing a (possibly empty) set of worlds R from $[k]$. It is assumed that $R \subseteq [k]$. A contraction operation on k by Ψ (conventionally denoted $k - \Psi$) has the consequence of adding a (possibly empty) set of worlds A to $[k]$. Here we assume that $A \cap [k] = \emptyset$. A revision operation on k by Ψ (conventionally denoted $k * \Psi$) has the consequence of adding (possibly empty) A to and removing (possibly empty) R from $[k]$. An update operation on k by Ψ (conventionally denoted $k \diamond \Psi$) also has the consequence of adding A to and removing R from $[k]$. Of course, in each case, the content of A may be different (similarly for R). In the case of revision and update, $[k * \Psi]$, respectively, $[k \diamond \Psi]$ could be accomplished by $([k] \cup A) \setminus R$ or $([k] \setminus R) \cup A$ or some other series of set theoretic operations resulting in $A \subseteq [k! \Psi]$ and $R \cap [k! \Psi] = \emptyset$, where $! \in \{*, \diamond\}$.

We represent an agent's current beliefs as a *belief-state*

$$b = \{(\omega_1, p_1), (\omega_2, p_2), \dots, (\omega_n, p_n)\},$$

where p_i is the non-zero probability that ω_i is the actual world in which the agent is and every world ω_i appears at most once in b . We allow a belief-state to be an empty set, but if not empty, then $\sum_{(\omega, p) \in b} p = 1$. A knowledge-base k_b extracted from b is a sentence which has the worlds in b as models and no other models: k_b is extracted from b if and only if for all worlds $\omega \in \Omega$, $\omega \models k_b$ iff $(\omega, p) \in b$. The belief-state *after* a change (after removal or addition of worlds) is denoted $b^!$. The set of worlds which model a KB extracted from b will be denoted as K . We also say that K is the set of worlds in b . In other words, if k is the sentence representing the KB extracted from b , then

³Some pre-order on worlds may also be involved during the decisions, however, our approach is independent of that information.

$K = [k]$. Therefore, (i) $\{\omega \mid (\omega, p) \in b\} = K$, (ii) $\{\omega \mid (\omega, p) \in b^1\} = K \setminus R$, if worlds are only removed from K , (iii) $\{\omega \mid (\omega, p) \in b^1\} = K \cup A$, if worlds are only added to K and (iv) $\{\omega \mid (\omega, p) \in b^1\} = (K \cup A) \setminus R = (K \setminus R) \cup A$, if worlds are added to *and* removed from K .

3 Redistributing Probabilities

We make the following supposition. Two worlds are more similar the closer they are to each other. Hence, *in the absence of information to the contrary*, the closer two worlds are to each other, the more similar the likelihood should be that they represent the actual world. We use this insight to guide the design of the methods for both cases, when removing worlds from and when adding worlds to the current belief-state. We acknowledge that many probability distributions over worlds seem not to reflect this situation. For instance, a given probability distribution may well have three worlds ω_1, ω_2 and ω_3 and their respective probabilities p_1, p_2 and p_3 such that ω_2 is much closer to ω_1 than ω_3 is to ω_1 , yet the difference between p_1 and p_3 is much smaller than the difference between p_1 and p_2 . However, one can think of distance information as lying in the background of probability information; the less informative the probabilities become, the more the distances between worlds move to the foreground. In this work, both sources of information are used.

Whether adding or removing a world ω , our method requires a notion of (inverse-)distance-based weight associated with worlds in K with respect to ω .

In the next two subsections, the foundations of our general approach are laid. First, the method for removing worlds and distributing probability mass is detailed. Second, the method for adding worlds and assigning probabilities is detailed.

Removing Worlds

We consider the case when worlds are removed due to a belief change operation. The question is, ‘How should the probabilities of the worlds in R be distributed amongst the remaining worlds?’

Our reasoning is the following. The closer a world ω is to a world ω^\times which is going to be removed, the less likely/probable ω is: If the agent decides that ω^\times is impossible, then by the supposition that two worlds are more similar the closer they are to each other, ω should also tend towards impossibility. By moving more of ω^\times ’s probability p^\times to worlds far away, closer worlds become relatively less likely. Our approach is thus, for each world $\omega^\times \in R$ removed from K , assign part of p^\times to a remaining world in proportion to the removed world’s distance from the remaining world, for every remaining world. Let $S \subseteq \Omega$.

Definition 3.1 For $\omega^\times \in R$ and $\omega \in S$,

$$\delta^{rem}(\omega^\times, \omega, S) = \frac{d(\omega^\times, \omega)}{\sum_{\omega' \in S} d(\omega^\times, \omega')}.$$

That is, $\delta^{rem}(\omega^\times, \omega, S)$ is the normalized distance of ω^\times from world $\omega \in S$.

Let B be the set of all possible belief-states.

Definition 3.2 The probabilistic world removal operation $\langle prem \rangle : B \times 2^\Omega \rightarrow B$ is a function such that $b \langle prem \rangle R =$

$$\bigcup_{\omega \in K \setminus R} \{(\omega, p_{new}) \mid (\omega, p) \in b \text{ and } p_{new} = p + \sum_{\omega^\times \in R, (\omega^\times, p^\times) \in b} p^\times \delta^{rem}(\omega^\times, \omega, K \setminus R)\}.$$

Note that $p^\times \delta^{rem}(\omega^\times, \omega, K \setminus R)$ is added to p for every $\omega^\times \in R$, where p^\times is weighted by $\delta^{rem}(\dots)$. Note that Definition 3.2 implies that if $K \setminus R = \emptyset$ or if $b = \emptyset$, then $b \langle prem \rangle R = \emptyset$.

Proposition 3.1 The probabilistic world removal operation $b \langle prem \rangle R$ results in a probability distribution over $K \setminus R$.

If b is a uniform probability distribution, then, in general, $b \langle prem \rangle R$ is not. For instance, $\{(\omega_{111}, 0.25), (\omega_{110}, 0.25), (\omega_{100}, 0.25), (\omega_{000}, 0.25)\} \langle prem \rangle \{\omega_{111}\} = \{(\omega_{110}, 0.29), (\omega_{100}, 0.33), (\omega_{000}, 0.38)\}$.

Suppose the set of atoms is $\{\text{call, battery, charge}\}$, where *call* stands for the proposition ‘Will receive call in next two minutes’, *battery* stands for the proposition ‘The battery will last two minutes on standby’ and *charge* stands for the proposition ‘Can reach a charging point within two minutes’. We abbreviate worlds by their truth vectors, where, for instance, 111 indicates that *call, battery* and *charge* are true, 110 indicates that *call* and *battery* are true and *charge* is false, and so on. We may also write $\omega_{111}, \omega_{110}$, and so on. For the purpose of illustration, we take the distance measure to be the Hamming distance $d^H(\cdot)$ between truth vectors. So, for instance, $d^H(111, 101) = 1$, $d^H(011, 101) = 2$ and $d^H(111, 000) = d^H(010, 101) = 3$. (Hereafter $d^H(\cdot)$ is denoted simply as $d(\cdot)$.) Keep in mind though that any distance measure conforming to the definition of the pseudo-distance function (Def. 2.1) can be used.

Suppose that the current belief-state is $b = \{(\omega_{111}, 0.25), (\omega_{110}, 0.25), (\omega_{100}, 0.25), (\omega_{000}, 0.25)\}$ ($K = \{\omega_{111}, \omega_{110}, \omega_{100}, \omega_{000}\}$). In natural language, the agent believes the following situations are equally likely. A call will come in in the next two minutes, except if both the battery will not last another two minutes and a charger can be reached before then, and the situation in which none of the propositions is true. And suppose that, due to new information and some higher level reasoning, the agent wants to remove 100 and 000 from its belief-state. This situation can be represented pictorially as in Figure 1.

World 111 gets an extra $0.25\delta(100, 111, K) + 0.25\delta(000, 111, K) = 0.25(2/3) + 0.25(3/5) = 0.317$ and 110 gets an extra $0.25\delta(100, 110, K) + 0.25\delta(000, 110, K) = 0.25(1/3) + 0.25(2/5) = 0.183$. The new belief-state is thus $b^1 = \{(\omega_{111}, 0.567), (\omega_{110}, 0.433)\}$.



Adding Worlds

We now consider the case when worlds are added due to a belief change operation. The question is, ‘How should probabilities be assigned to the worlds in A and how should the probabilities of the current worlds change?’

There are several ways to assign the A -world probabilities and change the current worlds’ probabilities. In this paper, we can only introduce some basic ideas and investigate one or two intuitive and reasonable approaches. Two operators

⁴We usually write function signature $\langle x \rangle(b, S)$ as an operator in infix notation: $b \langle x \rangle S$, where x and S are just place-holders here.




Figure 1: Four worlds in the current belief-state. Crossed out worlds are to be removed.




Figure 2: Two worlds (empty circles) to be added to two worlds in the current belief-state.

will be defined, a simpler $\langle padd_1 \rangle$ and a more sophisticated $\langle padd_2 \rangle$.

Recall that when removing a world ω^\times , remaining worlds closer to ω^\times become relatively less likely than remaining worlds farther from ω^\times . When adding worlds, a similar but opposite movement of probability mass should occur: When adding a world ω^+ , current worlds closer to ω^+ should become relatively more likely than current worlds farther from ω^+ . When removing worlds, the probabilities of the worlds to remove are known. However, the worlds to add do not come with probabilities, we can thus not immediately apply the approach of redistributing probabilities as is done in the previous subsection.

One could estimate probabilities for the A -worlds by using a method based on our supposition that two worlds are more similar the closer they are to each other.

We shall use the notion of the *inverse* distance of $\omega \in K$ from an added world. There are several ways to define the inverse of a set of values, each way having slightly different results. It is beyond the scope of this paper to investigate the effect of each instantiation of the definition of *inverse*. We choose $1/d(\omega^+, \omega)$ because of its simplicity and because it integrates well with the rest of our proposed methods: Let $S \subseteq \Omega$.

Definition 3.3 For $\omega^+ \in A$ and $\omega \in S$, $\delta^{add}(\omega^+, \omega, S) = \frac{(1/d(\omega^+, \omega))}{\sum_{\omega' \in S} (1/d(\omega^+, \omega'))}$.

That is, $\delta^{add}(\omega^+, \omega, S)$ is the normalized inverse-distance of ω^+ from world $\omega \in S$.

Definition 3.4 Let

$$\sigma = 1 + \sum_{\omega^+ \in A} \sum_{(\omega, p) \in b} p \delta^{add}(\omega^+, \omega, K),$$

a normalizing factor. Operation $\langle padd_1 \rangle : B \times 2^\Omega \rightarrow B$ is a function such that if $b \neq \emptyset$, then

$$b \langle padd_1 \rangle A = \bigcup_{\omega \in K} \{(\omega, p/\sigma) \mid (\omega, p) \in b\} \cup \bigcup_{\omega^+ \in A} \{(\omega^+, p^+/\sigma) \mid p^+ = \sum_{(\omega, p) \in b} p \delta^{add}(\omega^+, \omega, K)\},$$

else $\emptyset \langle padd_1 \rangle A = \{(\omega^+, 1/|A|) \mid \omega^+ \in A\}$.

Note that ω^+ gets the proportion $\delta^{add}(\omega^+, \omega, K)$ of p for every $\omega \in K$ (in the case $b \neq \emptyset$).

Instead of estimating the probabilities for the A -worlds by using $\langle padd_1 \rangle$, another approach is to *initialize* the probabilities of the A -worlds to reflect the average probability of the worlds already believed, and then adjust added worlds according to their distances from the worlds already believed. In other words, set the probability of every world in A to $1/|b|$ ($= 1/|K|$), or if $b = \emptyset$, then to $1/|A|$. The problem is that the probability mass of the add-set cannot simply be shifted to the current worlds—as is done when removing worlds—because the A -worlds must remain and they still need their probabilities. One could recover from this problem by removing the add-set in a way analogous to that of $\langle prem \rangle$, and then determine probabilities for the A -worlds by the application of $\langle padd_1 \rangle$. A removal process $\langle prem' \rangle$ adapted for the new addition operator $\langle padd_2 \rangle$ is defined next.

Definition 3.5 Let J be extracted from b' . Operation $\langle prem' \rangle : B \times 2^\Omega \rightarrow B$ is a function such that

$$b' \langle prem' \rangle A = \bigcup_{\omega \in J \setminus A} \{(\omega, p_{new}) \mid (\omega, p) \in b'\}$$

and $p_{new} = p + \sum_{\omega^+ \in A, (\omega^+, p^+) \in b'} p^+ \delta^{add}(\omega^+, \omega, J \setminus A)$.

Note that $p^+ \delta^{add}(\omega^+, \omega, J \setminus A)$ is added to p for every $\omega^+ \in A$, where p^+ is weighted by $\delta^{add}(\dots)$ (not $\delta^{rem}(\dots)$) as in Def. 3.2).

Next we give the formal definition of $\langle padd_2 \rangle$. Algorithm 1 gives the addition operation procedurally, as an aid in clarifying the definition.

Definition 3.6 Let $\sigma = 1 + |A|/|K|$, a normalizing factor. The probabilistic world addition operation $\langle padd_2 \rangle : B \times 2^\Omega \rightarrow B$ is a function such that $b \langle padd_2 \rangle A = (\{(\omega, p/\sigma) \mid (\omega, p) \in b\} \cup \{(\omega^+, 1/\sigma|b|) \mid \omega^+ \in A\}) \langle prem' \rangle A \langle padd_1 \rangle A$.

Division by σ in Definition 3.6 simulates normalization seen at line 3 in the algorithm.

Algorithm 1: $\langle padd_2 \rangle$

Input: b : belief-state, A : set of worlds to add to b
Output: new belief-state b' incl. worlds in A and their probabilities

- 1 **foreach** $\omega^+ \in A$ **do**
- 2 $b \leftarrow b \cup \{(\omega^+, 1/|b|)\}$
- 3 **Normalize**(b)
- 4 $b^- \leftarrow b \langle prem' \rangle A$
- 5 $b' \leftarrow b^- \langle padd_1 \rangle A$
- 6 **return** b'

Proposition 3.2 The probabilistic world addition operation $b \langle padd_2 \rangle A$ results in a probability distribution over $K \cup A$.

Suppose that the current belief-state is $b = \{(\omega_{111}, 0.5), (\omega_{110}, 0.5)\}$ ($K = \{\omega_{111}, \omega_{110}\}$). In natural language, given that the agent thinks s/he will receive a call in the next two minutes and that the battery will last another two minutes, it is equally as likely that s/he will be able to reach a charger within two minutes as not. And suppose that, due to new information and some higher level reasoning, the agent wants to add 100 and 000 to its belief-state. This situation can be represented pictorially as in Figure 2.

We work through the example in terms of Algorithm 1. Lines 2 and 3: b becomes $\{(\omega_{111}, 0.25), (\omega_{110}, 0.25), (\omega_{100}, 0.25), (\omega_{000}, 0.25)\}$. Line 4: b^- is set to $b\langle prem \rangle A$ which is $\{(\omega_{111}, 0.43), (\omega_{110}, 0.56)\}$. Line 5: b^+ is set to $b^-\langle padd_1 \rangle A$ which we work out in some detail (cf. Def. 3.4): Initially, before normalizing, $p_{100}^+ = p_{111}\delta(100, 111, K) + p_{110}\delta(100, 110, K) = 0.43(0.333) + 0.56(0.667) = 0.522$ and $p_{000}^+ = p_{111}\delta(000, 111, K) + p_{110}\delta(000, 110, K) = 0.43(0.4) + 0.56(0.6) = 0.513$. The final, normalized belief-state is $b^+ = \{(\omega_{111}, 0.213), (\omega_{110}, 0.278), (\omega_{100}, 0.257), (\omega_{000}, 0.252)\}$.

The assigned probabilities change as follows when the input belief-state has a different distribution: $b = \{(\omega_{111}, 0.9), (\omega_{110}, 0.1)\}$. Then $b^+ = \{(\omega_{111}, 0.328), (\omega_{110}, 0.190), (\omega_{100}, 0.236), (\omega_{000}, 0.246)\}$. Note that even though 000 is on average farther away from the current worlds than 100, 000 is assigned a greater probability than 100. This is because worlds farther away from the current worlds tend to get probabilities closer to the *average* probability of the current worlds. In the previous example, however, 100 is assigned the greater probability.

4 The Belief-Change Framework

A change in belief-state may require the addition *and* removal of worlds. A and R are determined appropriately in a pre-processing step for the particular initial belief-state and sentence with which to extend, contract, revise or update.

In this section, we shall propose two strategies for general probabilistic belief change. Then we mention some philosophical and technical problems that can occur.

Recall that K is the set of worlds in the current belief-state, A is the (possibly empty) set to be added and R the (possibly empty) set to be removed. Recall that $A \cap K = \emptyset$ and $R \subseteq K$ (with the consequence that $A \cap R = \emptyset$). One more assumption is made: $(K \setminus R) \cup A = (K \cup A) \setminus R \neq \emptyset$, that is, after adding and removing worlds, the resulting belief-state will contain at least one world.

In the following, when we do not distinguish between $\langle padd_1 \rangle$ and $\langle padd_2 \rangle$, then we write $\langle padd \rangle$.

Two obvious strategies for performing probabilistic belief change on b to determine b^+ , when given an add-set A and a remove-set R are

- (1) First determine $b' = b\langle padd \rangle A$, then determine $b^+ = b'\langle prem \rangle R$.
- (2) First determine $b' = b\langle prem \rangle R$, then determine $b^+ = b'\langle padd \rangle A$.

We define two belief change operators called *accepting*

(denoted $\langle pacc \rangle$) based on strategy (1) and *rejecting* (denoted $\langle prej \rangle$) based on strategy (2).⁵

Definition 4.1 The *probabilistically accepting belief change operation* $\langle pacc \rangle$ is defined as $b^+ = b\langle pacc \rangle A, R$ if and only if $b^+ = (b\langle padd \rangle A)\langle prem \rangle R$.

Definition 4.2 The *probabilistically rejecting belief change operation* $\langle prej \rangle$ is defined as $b^+ = b\langle prej \rangle A, R$ if and only if $b^+ = (b\langle prem \rangle R)\langle padd \rangle A$.

With $\langle pacc \rangle$, when the A -worlds are added, the probabilities of the R -worlds (not yet removed) are used to determine what probabilities to assign to the A -worlds. This is arguably counter-intuitive, because R -worlds are about to be removed; the R -worlds should, in effect, have zero probability. Their probabilities should thus not influence the A -worlds. The use of $\langle pacc \rangle$ is also affected in this way when R -worlds are removed; part of the probability mass of R -worlds is moved to A -worlds. This is arguably counter-intuitive, because A -worlds do not have a historic relationship with R -worlds. If one takes the stance that the assignment of probabilities to A -worlds is *allowed* to be influenced by the probabilities of R -worlds and that the distribution of R -world probabilities is allowed to influence the A -worlds, then $\langle pacc \rangle$ is acceptable.

$\langle prej \rangle$ does not have these problems, however, the case where $K \setminus R = \emptyset$ when using $\langle pacc \rangle$ may be technically undesirable / philosophically debatable.

We now look at examples involving two cases: when $K \setminus R \neq \emptyset$ and when $K \setminus R = \emptyset$. Our two belief change operations are applied to both cases. ($\langle pacc_2 \rangle$ is used for $\langle pacc \rangle$.)

Suppose that the current belief-state is $b = \{(\omega_{111}, 0.16), (\omega_{110}, 0.16), (\omega_{101}, 0.16), (\omega_{100}, 0.16), (\omega_{011}, 0.16), (\omega_{010}, 0.16)\}$ and that $A = \{\omega_{001}, \omega_{000}\}$ and $R = \{\omega_{011}, \omega_{010}\}$. Note that $K \setminus R \neq \emptyset$.

Using $\langle pacc \rangle$: By Definition 4.1, $b' = b\langle padd_2 \rangle \{\omega_{011}, \omega_{010}\} = \{(\omega_{111}, 0.114), (\omega_{110}, 0.114), (\omega_{101}, 0.130), (\omega_{100}, 0.130), (\omega_{011}, 0.130), (\omega_{010}, 0.130), (\omega_{001}, 0.126), (\omega_{000}, 0.126)\}$ and $b^+ = b'\langle prem \rangle \{\omega_{011}, \omega_{010}\} = \{(\omega_{111}, 0.149), (\omega_{110}, 0.149), (\omega_{101}, 0.189), (\omega_{100}, 0.189), (\omega_{001}, 0.162), (\omega_{000}, 0.162)\}$.

Using $\langle prej \rangle$: By Definition 4.2, $b' = b\langle prem \rangle \{\omega_{011}, \omega_{010}\} = \{(\omega_{111}, 0.23), (\omega_{110}, 0.23), (\omega_{101}, 0.27), (\omega_{100}, 0.27)\}$ and $b^+ = b'\langle padd_2 \rangle \{\omega_{001}, \omega_{000}\} = \{(\omega_{111}, 0.14), (\omega_{110}, 0.14), (\omega_{101}, 0.19), (\omega_{100}, 0.19), (\omega_{001}, 0.17), (\omega_{000}, 0.17)\}$.

Now suppose that the current belief-state is $b = \{(\omega_{111}, 0.2), (\omega_{110}, 0.3), (\omega_{100}, 0.5)\}$ and that $A = \{\omega_{001}, \omega_{000}\}$ and $R = \{\omega_{111}, \omega_{011}, \omega_{010}\}$. Note that $K \setminus R = \emptyset$.

Using $\langle pacc \rangle$: $b' = b\langle padd_2 \rangle \{\omega_{001}, \omega_{000}\} = \{(\omega_{111}, 0.134), (\omega_{110}, 0.165), (\omega_{100}, 0.281), (\omega_{001}, 0.197), (\omega_{000}, 0.223)\}$ and $b^+ = b'\langle prem \rangle \{\omega_{111}, \omega_{011}, \omega_{010}\} = \{(\omega_{001}, 0.537), (\omega_{000}, 0.463)\}$.

Using $\langle prej \rangle$: $b^+ = \{(\omega_{001}, 0.5), (\omega_{000}, 0.5)\}$.

⁵The names of the operators are only suggestive of the flavor of their definitions. The names can also be used as mnemonics: the accepting operator first adds worlds; the rejecting operator first removes worlds.

Note that the two operations produce different results.

The Levi identity states that revision of a sentence k by Ψ is equivalent to first contraction of k by $\neg\Psi$, then expansion of the result by Ψ (Gärdenfors, 1988). It is always the case that the first step (contraction) results in a (possibly new) sentence k' such that $[k] \subseteq [k']$. In other worlds, the first step may add worlds to $[k]$, but never remove worlds. The second step (expansion) results in a (possibly new) sentence k'' such that $[k''] \subseteq [k']$. In other worlds, the second step may remove worlds from $[k']$, but never adds worlds. If belief revision were to be performed, we assume that the add-set A and the remove-set R can be determined in some pre-processing step. Operation $\langle pace \rangle$ thus mirrors belief revision translated as the Levi identity. Belief expansion, contraction and update can also be simulated with this operation. If expansion is the operation to perform, A will be empty. If contraction is the operation, R will be empty. With belief update, there may be worlds to add and worlds to remove, like with revision. Operation $\langle prej \rangle$ does not mirror belief revision translated as the Levi identity. Nonetheless, $\langle prej \rangle$ satisfies the stance that probabilities of the add-set and of the remove-set should be independent, while $\langle pace \rangle$ does not.

Whereas the Levi identity defines revision in terms of contraction and expansion, the Harper identity defines contraction in terms of revision. Let Φ be a belief set—the set of conjuncts of k . The Harper identity states that the contraction of k by Ψ is equivalent to the intersection of the sentences in Φ with the result of the revision of Φ by $\neg\Psi$ (Gärdenfors, 1988). Semantically, this is adding and removing worlds due to revision by $\neg\Psi$, resulting in $[k']$, then taking $[k'] \cup [k]$ as the final result. In the special case where $\neg\Psi$ is consistent with k , you would remove $[\Psi]$ from $[k]$ and then add $[\Psi]$ again (when you add the models of k ; so your result is exactly $[k]$ again). Operation $\langle prej \rangle$ has this sequence. The question in this case is, ‘Is the probability distribution obtained the same as the one started off with in our framework?’ It turns out that, in general, $b\langle prej \rangle[\Psi] \neq b$. An example is when $b = \{(\omega_{011}, 0.3), (\omega_{010}, 0.7)\}$ and $[\Psi]$ is $\{\omega_{010}\}$. Then $b\langle prej \rangle[\Psi] = \{(\omega_{011}, 0.31), (\omega_{010}, 0.69)\}$. Strictly speaking, the operation as used in the example above is undefined. Recall that we assume that a pre-processing step will determine the add-set and remove-set; moreover, we demand that $A \cap R = \emptyset$, which is not the case in the special case under discussion. In our framework, one would rather determine before hand that $A = R$ and thus not process b . In other worlds, we would perform the trivial process of removing and adding nothing, so that $b = b^!$.

5 Conclusion

Our research into this problem uncovered that whether or not one takes a stance that the probabilities of worlds to be added and the probabilities of worlds to be removed should interact, makes a difference to the approach to be taken and the resulting probability distribution of the new belief-state. It is arguable whether the probabilities of the add-set and remove-set should be independent or whether they need not be.

Interpreting why $\langle pace \rangle$ and $\langle prej \rangle$ produce different results is left for the future. We would also like to assess which of the two operations is in some sense the ‘better’ operation, or which operation is better suited to which kinds of situations.

We would also like to prove or disprove the following conjecture. When performing the $\langle pace \rangle$ or $\langle prej \rangle$ operation, the probability mass of the add-set is non-decreasing with an increase in the probability mass of the remove-set.

An in-depth comparison of the presented framework with the related work is still necessary. This would also including an analysis of the presented approach in terms of the postulates of belief change with respect to the axioms of probability (Gärdenfors, 1988, Chap. 5, e.g.), and in terms of the theory of Bayesian conditioning and Bayesian update.

Update in the framework presented here is dealt with with exactly the same processes as for belief revision, contraction and expansion. This is beneficial from the standpoint of economy and simplicity. It also allows one to think of all the belief change operations as part of one family, directly comparable. However, by placing a probability distribution over possible worlds, a belief-state may change only because of a change in probability distribution over the *same* set of worlds, due to some new information or change in situation. Our framework does not consider this possibility. More research in this direction is required.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *J. of Symbolic Logic* 50(2):510–530.
- Boutilier, C. 1995. On the revision of probabilistic belief states. *Notre Dame J. of Formal Logic* 36(1):158–183.
- Chhogyal, K.; Nayak, A.; Schwitter, R.; and Sattar, A. 2014. Probabilistic belief revision via imaging. In Pham, D., and Park, S., eds., *Proc. of PRICAI 2014*, volume 8862 of LNCS, 694–707. Springer-Verlag.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.
- Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In *Proc. of KR-91*, 387–394.
- Kern-Isberner, G. 2008. Linking iterated belief change operations to nonmonotonic reasoning. In *Proc. of KR-08*, 166–176. AAAI Press.
- Lehmann, D.; Magidor, M.; and Schlechta, K. 2001. Distance semantics for belief revision. *J. of Symbolic Logic* 66(1):295–317.
- Lewis, D. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 85(3):297–315.
- Van Harmelen, F.; Lifshitz, V.; and Porter, B. 2008. *Handbook of Knowledge Representation*. Elsevier Science.