

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/307577667>

On Revision of Partially Specified Convex Probabilistic Belief Bases

Conference Paper · September 2016

CITATIONS

4

READS

38

3 authors, including:



Gavin Rens
KU Leuven

31 PUBLICATIONS 67 CITATIONS

[SEE PROFILE](#)



Giovanni Casini
University of Luxembourg

30 PUBLICATIONS 218 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Agent Architectures [View project](#)



An Agent Logic for Reasoning under Uncertainty [View project](#)

On Revision of Partially Specified Convex Probabilistic Belief Bases¹

Gavin Rens² and Thomas Meyer³ and Giovanni Casini⁴

Abstract. We propose a method for an agent to revise its incomplete probabilistic beliefs when a new piece of propositional information is observed. In this work, an agent's beliefs are represented by a set of probabilistic formulae – a belief base. The method involves determining a *representative* set of 'boundary' probability distributions consistent with the current belief base, revising each of these probability distributions and then translating the revised information into a new belief base. We use a version of Lewis Imaging as the revision operation. The correctness of the approach is proved. An analysis of the approach is done against six rationality postulates. The expressivity of the belief bases under consideration are rather restricted, but has some applications. We also discuss methods of belief base revision employing the notion of optimum entropy, and point out some of the benefits and difficulties in those methods. Both the boundary distribution method and the optimum entropy methods are reasonable, yet yield different results.

1 INTRODUCTION

Suppose an agent represents its probabilistic knowledge with a set of statements; every statement says something about the probability of some features the agent is aware of. Ideally, the agent would want to have enough information to, at least, identify one probability distribution over all the situations (worlds) it deems possible. However, if the agent could not gather sufficient data or if it was not told or given sufficient information, it would not be able to pinpoint exactly one probability distribution. An agent with this sort of ignorance, can be thought of as having beliefs compatible with a *set of* distributions. Now, this agent might need to revise its beliefs when new (non-probabilistic) information is received, even though the agent's beliefs do not characterize a *particular* probability distribution over its current possible worlds.

Several researchers argue that using a single probability distribution requires the agent to make unrealistically precise uncertainty distinctions [9, 27, 29]. "One widely-used approach to dealing with this has been to consider sets of probability measures as a way of modeling uncertainty," [9]. However, simply applying standard probabilistic conditioning to each of the measures/distributions in the set individually and then combining the results is either not recommended

because it produces unsatisfactory results [9, 27] or because it is not computable due to the set being infinite. The framework presented in this paper proposes two ways to go from one 'probabilistically incomplete' belief base to another when new information is acquired.

Both belief revision methods presented, essentially follow this process: From the original belief base, determine a relatively small set of belief states / probability distributions which are 'compatible' with the belief base and in a sense, representative of the belief base. (We shall use the terms *belief state* and *probability distribution* interchangeably). Then revise only the belief states in this representative set. Finally, induce a new, revised belief base from the revised representative set.

We shall present two approaches to determine the representative set of belief states from the current belief base: (i) The approach we focus on involves finding belief states which, in a sense, are at the boundaries of the constraints implied by the belief base. These 'boundary belief states' can be thought of as drawing the outline of the convex space of beliefs. This outline is then revised to form a new outline shape, which can be translated into a new belief base. (ii) As a possible alternative approach, the representative set is a *single* belief state which can be imagined to be at the center of the outline of the first approach. This 'central' belief state is found by determining the one in the space of beliefs which is least biased or most entropic in terms of information theory [11, 4].

For approach (i) – where the canonical set is the set of boundary belief states – we shall prove that the revised canonical set characterizes the set of all belief states which would have resulted from revising all (including interior) belief states compatible with the original belief base.

The next section provides the relevant background theory and notation. Section 3 presents a generalized imaging method for revising probabilistic belief states. Then we describe the application of generalized imaging in our main contribution; revising boundary belief states instead of all belief states. The subsequent section explain two other approaches of revising our belief bases, based on optimum entropy. The first method finds a single representative belief state through maximum entropy inference and the second method revises boundary belief states using minimum cross-entropy inference. All three methods can be considered motivated methods, yet yield different results. Then, in Section 6, we shall list six traditional rationality postulates, and check how well the main approach fares against them. The related work is discussed in Section 7. We end with a section on future possible directions of research and some concluding remarks.

2 PRELIMINARIES

We shall work with classical propositional logic. Let \mathcal{P} be the finite set of n atomic propositional variables (*atoms*, for short). Formally,

¹ This is a modified and extended version of a paper that was presented at the Sixteenth International Workshop on Non-Monotonic Reasoning, 22–24 April 2016, Cape Town, South Africa.

² Centre for AI Research, and University of KwaZulu-Natal, School of Mathematics, Statistics and Computer Science, and CSIR Meraka, South Africa, email: gavinrens@gmail.com

³ Centre for AI Research, and University of Cape Town, Department of Computer Science, and CSIR Meraka, South Africa, email: tmeyer@cs.uct.ac.za

⁴ University of Luxembourg, Computer Science and Communication Research Unit, Luxembourg, email: giovanni.casini@uni.lu

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 non-empty subset W of the conceivable worlds; W is called the possible worlds. Often, in the exposition of this paper, a world will be referred to by its truth vector. For instance, if the vocabulary is placed in order $\langle q, r \rangle$ and $w_3 \models \neg q \wedge r$, then w_3 may be referred to as 01.⁵ Let L be all propositional formulae which can be formed from \mathcal{P} and the logical connectives \wedge and \neg , with \top abbreviating tautology and \perp abbreviating contradiction.

Let β be a sentence in L . $[\beta]$ denotes the set of β -worlds, that is, the elements of W satisfying β . The worlds satisfying all sentences in a set of sentences K are denoted by $[K]$.

We define the probabilistic language $L^{prob} = \{(\alpha) \bowtie x \mid \alpha \in L, \bowtie \in \{\leq, =, \geq\}, x \in [0, 1]\}$. We propose a belief base (BB) to be a consistent (logically satisfiable) subset of L^{prob} . A BB specifies an agent's knowledge.

The basic semantic element of an agent's beliefs is a probability distribution or a *belief state*

$$b = \{(w_1, p_1), (w_2, p_2), \dots, (w_n, p_n)\},$$

where p_i is the probability that w_i is the actual world in which the agent is. $\sum_{(w,p) \in b} p = 1$. We may also use c to refer to a belief state. For parsimony, let $b = \langle p_1, \dots, p_n \rangle$ be the probabilities that belief state b assigns to w_1, \dots, w_n where $\langle w_1, w_2, w_3, w_4 \rangle = \langle 11, 10, 01, 00 \rangle$, and $\langle w_1, w_2, \dots, w_8 \rangle = \langle 111, 110, \dots, 000 \rangle$. Let Π be the set of all belief states over W .

$b(\alpha)$ abbreviates $\sum_{w \in W, w \models \alpha} b(w)$. b satisfies formula $(\alpha) \bowtie x$ (denoted $b \models (\alpha) \bowtie x$) iff $b(\alpha) \bowtie x$. If B is a set of formulae, then b satisfies B (denoted $b \models B$) iff $\forall \gamma \in B, b \models \gamma$. If B and B' are sets of formulae, then B entails B' (denoted $B \models B'$) iff for all $b \in \Pi$, $b \models B'$ whenever $b \models B$. If $B \models \{\gamma\}$ then we simply write $B \models \gamma$. B is logically equivalent to B' (denoted $B \equiv B'$) iff $B \models B'$ and $B' \models B$.

Instead of an agent's beliefs being represented by a single belief state, a BB B represents a *set* of belief-states: Let $\Pi^B := \{b \in \Pi \mid b \models B\}$. A BB B is *satisfiable* (*consistent*) iff $\Pi^B \neq \emptyset$. We can now also define entailment as $B \models B'$ iff $\Pi^B \subseteq \Pi^{B'}$.

The technique of *Lewis imaging* for the revision of belief states, requires a notion of distance between worlds to be defined. Various notions of distance are possible, however, a study of their influence on the imaging technique is beyond the scope of this paper. We use a pseudo-distance measure between worlds, as defined by Lehmann et al. [18] and adopted by Chhogyal et al. [3]. We add a 'faithfulness' condition, which we feel is lacking from the definition of Lehmann et al. [18]: without this condition, a pseudo-distance measure would allow all worlds to have zero distance between them. Boutilier [2] mentions this condition, and we use his terminology: "faithfulness".

Definition 1 A pseudo-distance function $d : W \times W \rightarrow \mathbb{Z}$ satisfies the following five conditions: for all worlds $w, w', w'' \in W$,

1. $d(w, w') \geq 0$ (Non-negativity)
2. $d(w, w) = 0$ (Identity)
3. $d(w, w') = d(w', w)$ (Symmetry)
4. $d(w, w') + d(w', w'') \geq d(w, w'')$ (Triangular Inequality)
5. if $w \neq w'$, then $d(w, w') > 0$ (Faithfulness)

Presently, the foundation theory, or paradigm, for studying belief change operations is commonly known as AGM theory [1, 7]. Typically, belief change (in a static world) can be categorized as expansion, revision or contraction, and is performed on a belief set, the set

of sentences K closed under logical consequence. Expansion (denoted $+$) is the logical consequences of $K \cup \{\alpha\}$, where α is new information and K is the current belief set. Contraction of α is the removal of some sentences until α cannot be inferred from K . Revision is when α is (possibly) inconsistent with K and K is (minimally) modified so that the new K remains consistent and entails α . In this view, when the new information is consistent with the original beliefs, expansion and revision are equivalent.

3 GENERALIZED IMAGING

It is not yet universally agreed what revision means in a probabilistic setting. One school of thought says that probabilistic expansion is equivalent to Bayesian conditioning. This is evidenced by Bayesian conditioning (BC) being defined only when $b(\alpha) \neq 0$, thus making BC expansion equivalent to BC revision. In other words, one could define expansion (restricted revision) to be

$$b \text{ BC } \alpha = \{(w, p) \mid w \in W, p = b(w \mid \alpha), b(\alpha) \neq 0\},$$

where $b(w \mid \alpha)$ can be defined as $b(\phi_w \wedge \alpha)/b(\alpha)$ and ϕ_w is a sentence identifying w (i.e., a complete theory for w).

To accommodate cases where $b(\alpha) = 0$, that is, where α contradicts the agent's current beliefs and its beliefs need to be revised in the stronger sense, we shall make use of *imaging*. Imaging was introduced by Lewis [20] as a means of revising a probability distribution, and has been discussed in other work too [7, 6, 3, 25]. Informally, Lewis's original solution for accommodating contradicting evidence α is to move the probability of each world to its closest, α -world. Lewis made the strong assumption that every world has a *unique* closest α -world. More general versions of imaging allows worlds to have *several*, equally proximate, closest worlds.

Gärdenfors [7] calls one generalization of Lewis's imaging, *general imaging*. Our method is also a generalization of Lewis's imaging. We thus refer to his as *Gärdenfors's general imaging* and to our method as *generalized imaging* to distinguish them. It should be noted that all three these imaging methods are general revision methods and can be used in place of Bayesian conditioning for expansion. "Thus imaging is a more general method of describing belief changes than conditionalization," [7, p. 112] in the sense that Bayesian conditioning cannot deal with contradicting evidence but imaging can.

Let $Min(\alpha, w, d)$ be the set of α -worlds closest to w with respect to pseudo-distance d . Formally,

$$Min(\alpha, w, d) := \{w' \in [\alpha] \mid \forall w'' \in [\alpha], d(w', w) \leq d(w'', w)\},$$

where $d(\cdot)$ is some pseudo-distance measure between worlds (e.g., Hamming or Dalal distance).

Example 1 Let the vocabulary be $\{q, r, s\}$. Let α be $(q \wedge r) \vee (q \wedge \neg r \wedge s)$. Suppose d is Hamming distance. Then

$$\begin{aligned} Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 111, d) &= \{111\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 110, d) &= \{110\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 101, d) &= \{101\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 100, d) &= \{110, 101\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 011, d) &= \{111\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 010, d) &= \{110\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 001, d) &= \{101\} \\ Min((q \wedge r) \vee (q \wedge \neg r \wedge s), 000, d) &= \{110, 101\} \end{aligned}$$

⁵ $w \models \alpha$ is read 'w is a model for/satisfies α '.

Definition 2 (GI) Then generalized imaging (denoted GI) is defined as

$$b \text{ GI } \alpha := \{(w, p) \mid w \in W, p = 0 \text{ if } w \notin [\alpha], \\ \text{else } p = \sum_{\substack{w' \in W \\ w \in \text{Min}(\alpha, w', d)}} b(w') / |\text{Min}(\alpha, w', d)|\}.$$

In words, $b \text{ GI } \alpha$ is the new belief state produced by taking the generalized image of b with respect to α . Notice how the probability mass of non- α -worlds is shifted to their closest α -worlds. If a non- α -world w^\times with probability p has n closest α -worlds (equally distant), then each of these closest α -worlds gets p/n mass from w^\times .

We define $b_\alpha^\circ := b \circ \alpha$ so that we can write $b_\alpha^\circ(w)$, where \circ is a revision operator.

Example 2 Continuing on Example 1: Let $b = \langle 0, 0.1, 0, 0.2, 0, 0.3, 0, 0.4 \rangle$.

$(q \wedge r) \vee (q \wedge \neg r \wedge s)$ is abbreviated as α .

$$b_\alpha^{\text{GI}}(111) = \sum_{\substack{w' \in W \\ 111 \in \text{Min}(\alpha, w', d)}} b(w') / |\text{Min}(\alpha, w', d)| \\ = b(111) / |\text{Min}(\alpha, 111, d)| + b(011) / |\text{Min}(\alpha, 011, d)| \\ = 0/1 + 0/1 = 0.$$

$$b_\alpha^{\text{GI}}(110) = \sum_{\substack{w' \in W \\ 110 \in \text{Min}(\alpha, w', d)}} b(w') / |\text{Min}(\alpha, w', d)| \\ = b(110) / |\text{Min}(\alpha, 110, d)| + b(100) / |\text{Min}(\alpha, 100, d)| \\ + b(010) / |\text{Min}(\alpha, 010, d)| + b(000) / |\text{Min}(\alpha, 000, d)| \\ = 0.1/1 + 0.2/2 + 0.3/1 + 0.4/2 = 0.7.$$

$$b_\alpha^{\text{GI}}(101) = \sum_{\substack{w' \in W \\ 101 \in \text{Min}(\alpha, w', d)}} b(w') / |\text{Min}(\alpha, w', d)| \\ = b(101) / |\text{Min}(\alpha, 101, d)| + b(100) / |\text{Min}(\alpha, 100, d)| \\ + b(001) / |\text{Min}(\alpha, 001, d)| + b(000) / |\text{Min}(\alpha, 000, d)| \\ = 0/1 + 0.2/2 + 0/1 + 0.4/2 = 0.3.$$

And $b_\alpha^{\text{GI}}(100) = b_\alpha^{\text{GI}}(011) = b_\alpha^{\text{GI}}(010) = b_\alpha^{\text{GI}}(001) = b_\alpha^{\text{GI}}(000) = 0$. \square

4 REVISION VIA GI AND BOUNDARY BELIEF STATES

The most obvious way to revise a given belief base (BB) B is to revise every individual belief state in Π^B and then induce a new BB from the set of revised belief states. Formally, given observation α , first determine a new belief state b^α for every $b \in \Pi^B$ via the defined revision operation:

$$\Pi^{B^\alpha} = \{b^\alpha \in \Pi \mid b^\alpha = b \text{ GI } \alpha, b \in \Pi^B\}.$$

If there is more than only a single belief state in Π^B , then Π^B contains an infinite number of belief states. Then how can one compute Π^{B^α} ? And how would one subsequently determine B^α from Π^{B^α} ?

In the rest of this section we shall present a finite method of determining Π^{B^α} . What makes this method possible is the insight that Π^B can be represented by a finite set of ‘boundary’ belief states – those belief states which, in a sense, represent the limits or the convex hull of Π^B . We shall prove that the set of revised boundary belief states defines Π^{B^α} . Inducing B^α from Π^{B^α} is then relatively easy, as will be seen.

Let W^{perm} be every permutation on the ordering of worlds in W . For instance, if $W = \{w_1, w_2, w_3, w_4\}$, then $W^{\text{perm}} = \{\langle w_1, w_2, w_3, w_4 \rangle, \langle w_1, w_2, w_4, w_3 \rangle, \langle w_1, w_3, w_2, w_4 \rangle, \dots, \langle w_4, w_3, w_2, w_1 \rangle\}$. Given an ordering $W^\# \in W^{\text{perm}}$, let $W^\#(i)$ be the i -th element of $W^\#$; for instance, $\langle w_4, w_3, w_2, w_1 \rangle(2) = w_3$. Suppose we are given a BB B . We now define a function which, given a permutation of worlds, returns a belief state where worlds earlier in the ordering are assigned maximal probabilities according to the boundary values enforced by B .

Definition 3 *MaxASAP*($B, W^\#$) is the $b \in \Pi^B$ such that for $i = 1, \dots, |W|$, $\forall b' \in \Pi^B$, if $b' \neq b$, then $\sum_{j=1}^i b(W^\#(j)) \geq \sum_{k=1}^i b'(W^\#(k))$.

Example 3 Suppose the vocabulary is $\{q, r\}$ and $B_1 = \{(q) \geq 0.6\}$. Then, for instance, *MaxASAP*($B_1, \langle 01, 00, 11, 10 \rangle$) = $\{(01, 0.4), (00, 0), (11, 0.6), (10, 0)\} = \{(11, 0.6), (10, 0), (01, 0.4), (00, 0)\}$. \square

Definition 4 We define the boundary belief states of BB B as the set $\Pi_{\text{bnd}}^B := \{b \in \Pi^B \mid W^\# \in W^{\text{perm}}, b = \text{MaxASAP}(B, W^\#)\}$.

Example 4 Suppose the vocabulary is $\{q, r\}$ and $B_1 = \{(q) \geq 0.6\}$. Then

$$\Pi_{\text{bnd}}^{B_1} = \{\{(11, 1.0), (10, 0.0), (01, 0.0), (00, 0.0)\}, \\ \{(11, 0.0), (10, 1.0), (01, 0.0), (00, 0.0)\}, \\ \{(11, 0.6), (10, 0.0), (01, 0.4), (00, 0.0)\}, \\ \{(11, 0.6), (10, 0.0), (01, 0.0), (00, 0.4)\}, \\ \{(11, 0.0), (10, 0.6), (01, 0.4), (00, 0.0)\}, \\ \{(11, 0.0), (10, 0.6), (01, 0.0), (00, 0.4)\}\}.$$

\square

Next, the revision operation is applied to every belief state in Π_{bnd}^B . Let $(\Pi_{\text{bnd}}^B)_\alpha^{\text{GI}} := \{b' \in \Pi \mid b' = b_\alpha^{\text{GI}}, b \in \Pi_{\text{bnd}}^B\}$.

Example 5 Suppose the vocabulary is $\{q, r\}$ and $B_1 = \{(q) \geq 0.6\}$. Let α be $(q \wedge \neg r) \vee (\neg q \wedge r)$. Then

$$(\Pi_{\text{bnd}}^{B_1})_\alpha^{\text{GI}} = \{\{(11, 0.0), (10, 0.5), (01, 0.5), (00, 0.0)\}, \\ \{(11, 0.0), (10, 1.0), (01, 0.0), (00, 0.0)\}, \\ \{(11, 0.0), (10, 0.3), (01, 0.7), (00, 0.0)\}, \\ \{(11, 0.0), (10, 0.6), (01, 0.4), (00, 0.0)\}, \\ \{(11, 0.0), (10, 0.8), (01, 0.2), (00, 0.0)\}\}.$$

(Two revision operations produce $\{(11, 0), (10, 0.5), (01, 0.5), (00, 0)\}$.) \square

To induce the new BB B_{bnd}^α from $(\Pi_{\text{bnd}}^B)_\alpha^{\text{GI}}$, the following procedure is executed. For every possible world, the procedure adds a sentence enforcing the upper (resp., lower) probability limit of the world, with respect to all the revised boundary belief states. Trivial limits are excepted.

For every $w \in W$, $(\phi_w) \leq \bar{y} \in B^\alpha$, where $\bar{y} = \max_{b \in (\Pi_{bnd}^\alpha)^{GI}} b(w)$, except when $\bar{y} = 1$, and $(\phi_w) \geq \underline{y} \in B^\alpha$, where $\underline{y} = \min_{b \in (\Pi_{bnd}^\alpha)^{GI}} b(w)$, except when $\underline{y} = 0$.

The intention is that the procedure specifies B^α to represent the upper and lower probability envelopes of the set of revised boundary belief states. And thus, by Theorem 1, B^α defines the entire revised belief state space.

Example 6 Continuing Example 5, using the translation procedure just above, we see that $B_{1bnd}^\alpha = \{(\phi_{11}) \leq 0, (\phi_{10}) \geq 0.3, (\phi_{01}) \leq 0.7, (\phi_{00}) \leq 0.0\}$.

Note that if we let $B' = \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1, (q \wedge \neg r) \geq 0.3\}$, then $\Pi^{B'} = \Pi^{B_{1bnd}^\alpha}$. \square

Example 7 Suppose the vocabulary is $\{q, r\}$ and $B_2 = \{(\neg q \wedge \neg r) = 0.1\}$. Let α be $\neg q$. Then

$$\begin{aligned} \Pi_{bnd}^{B_2} &= \{(11, 0.9), (10, 0), (01, 0), (00, 0.1)\}, \\ &\quad \{(11, 0), (10, 0.9), (01, 0), (00, 0.1)\}, \\ &\quad \{(11, 0), (10, 0), (01, 0.9), (00, 0.1)\}, \end{aligned}$$

$$\begin{aligned} (\Pi_{bnd}^{B_2})_\alpha^{GI} &= \{(11, 0), (10, 0), (01, 0.9), (00, 0.1)\}, \\ &\quad \{(11, 0), (10, 0), (01, 0), (00, 1)\} \text{ and} \end{aligned}$$

$$B_{2bnd}^\alpha = \{(\phi_{11}) \leq 0, (\phi_{10}) \leq 0, (\phi_{01}) \leq 0.9, (\phi_{00}) \geq 0.1\}.$$

Note that if we let $B' = \{(\neg q) = 1, (\neg q \wedge r) \leq 0.9\}$, then $\Pi^{B'} = \Pi^{B_{2bnd}^\alpha}$. \square

Note that every world in W can be associated with the size of $Min(\alpha, w, d)$ for some α and d . Denote this size as $\#(w)$. Let $W^{Min(\alpha, d)}$ be a partition of W such that every block (equivalence class) blk of the partition is defined as follows. $blk = \{w_1, \dots, w_k\}$ iff $\#(w_1) = \dots = \#(w_k)$. Let $[w]$ denote block blk iff $w \in blk$. Finally, let w^i indicate that $i = \#(w)$, in other words, $[w^i]$ is the block containing all worlds such that $i = |Min(\alpha, w, d)|$. Let $m := \max_{w \in W} \{|Min(\alpha, w, d)|\}$.

Observation 1 Let $\delta_1, \delta_2, \dots, \delta_m$ be positive integers such that $i < j$ iff $\delta_i < \delta_j$. Let $\nu_1, \nu_2, \dots, \nu_m$ be values in $[0, 1]$ such that $\sum_{k=1}^m \nu_k = 1$. Associate with every ν_i a maximum value it is allowed to take: $most(\nu_i)$. For every ν_i , we define the assignment value

$$av(\nu_i) := \begin{cases} most(\nu_i) & \text{if } \sum_{k=1}^i \nu_k \leq 1 \\ 1 - \sum_{k=1}^{i-1} \nu_k & \text{otherwise} \end{cases}$$

Determine first $av(\nu_1)$, then $av(\nu_2)$ and so on. Then

$$\frac{av(\nu_1)}{\delta_1} + \dots + \frac{av(\nu_m)}{\delta_m} > \frac{\nu'_1}{\delta_1} + \dots + \frac{\nu'_m}{\delta_m}$$

whenever $\nu'_i \neq av(\nu_i)$ for some i . \square

For instance, let $\delta_1 = 1, \delta_2 = 2, \delta_3 = 3, \delta_4 = 4$. Let $most(\nu_1) = 0.5, most(\nu_2) = 0.3, most(\nu_3) = 0.2, most(\nu_4) = 0.3$. Then $av(\nu_1) = 0.5, av(\nu_2) = 0.3, av(\nu_3) = 0.2, av(\nu_4) = 0$ and

$$\frac{0.5}{1} + \frac{0.3}{2} + \frac{0.2}{3} + \frac{0}{4} = 0.716.$$

But

$$\frac{0.49}{1} + \frac{0.3}{2} + \frac{0.2}{3} + \frac{0.01}{4} = 0.709.$$

And

$$\frac{0.5}{1} + \frac{0.29}{2} + \frac{0.2}{3} + \frac{0.01}{4} = 0.714.$$

Lemma 1 essentially says that the belief state in Π^B which causes a revised belief state to have a maximal value at world w (w.r.t. all belief states in Π^B), will be in Π_{bnd}^B .

Lemma 1 For all $w \in W$, $\arg \max_{b_X \in \Pi^B} \sum_{w' \in W} b_X(w') / |Min(\alpha, w', d)|$ is in Π_{bnd}^B .

Proof:

Note that $\sum_{w' \in W} b(w') / |Min(\alpha, w', d)|$ can be written in the form

$$\frac{\sum_{w' \in [w^1]} b(w')}{1} + \dots + \frac{\sum_{w' \in [w^m]} b(w')}{m}.$$

Observe that there must be a $W^\# \in W^{perm}$ such that $W^\# = \langle w_1^1, \dots, w_{n1}^1, \dots, w_1^m, \dots, w_{nm}^m \rangle$. Then by the definition of the set of boundary belief states (Def. 4), $MaxASAP(B, W^\#)$ will assign maximal probability mass to $[w^1] = \{w_1^1, \dots, w_{n1}^1\}$, then to $[w^2] = \{w_1^2, \dots, w_{n2}^2\}$ and so on.

That is, by Observation 1, for some $b_X \in \Pi_{bnd}^B$, $b_X(w) = \max_{b_X \in \Pi^B} \sum_{w' \in W} b_X(w') / |Min(\alpha, w', d)|$ for all $w \in W$. Therefore, $\arg \max_{b_X \in \Pi^B} \sum_{w' \in W} b_X(w') / |Min(\alpha, w', d)|$ is in Π_{bnd}^B . \square

Let

$$\bar{x}^w := \max_{b \in \Pi_{bnd}^B} b(w) \quad \bar{X}^w := \max_{b \in \Pi^B} b(w)$$

$$\bar{y}^w := \max_{b \in (\Pi_{bnd}^B)_\alpha^{GI}} b(w) \quad \bar{Y}^w := \max_{b \in (\Pi^B)_\alpha^{GI}} b(w)$$

$$\underline{x}^w := \min_{b \in \Pi_{bnd}^B} b(w) \quad \underline{X}^w := \min_{b \in \Pi^B} b(w)$$

$$\underline{y}^w := \min_{b \in (\Pi_{bnd}^B)_\alpha^{GI}} b(w) \quad \underline{Y}^w := \min_{b \in (\Pi^B)_\alpha^{GI}} b(w)$$

Lemma 2 states that for every world, the upper/lower probability of the world with respect to Π_{bnd}^B is equal to the upper/lower probability of the world with respect to Π^B . The proof requires Observation 1 and Lemma 1.

Lemma 2 For all $w \in W$, $\bar{y}^w = \bar{Y}^w$ and $\underline{y}^w = \underline{Y}^w$.

Proof:

Note that if $w \notin [\alpha]$, then $\bar{y}^w = \bar{Y}^w = 0$ and $\underline{y}^w = \underline{Y}^w = 0$.

We now consider the cases where $w \in [\alpha]$.

$$\bar{y}^w = \bar{Y}^w$$

iff

$$\max_{b \in (\Pi_{bnd}^B)} b(w) = \max_{b \in (\Pi^B)} b(w)$$

iff

$$\begin{aligned} &\max_{b_X \in \Pi_{bnd}^B} \sum_{w' \in W} b_X(w') / |Min(\alpha, w', d)| \\ &= \max_{b_X \in \Pi^B} \sum_{w' \in W} b_X(w') / |Min(\alpha, w', d)| \end{aligned}$$

if $\bar{b}_x(w) = \bar{b}_X(w)$, where

$$\bar{b}_x(w) := \max_{b_x \in \Pi_{bnd}^B} \sum_{\substack{w' \in W \\ w \in \text{Min}(\alpha, w', d)}} b_x(w') / |\text{Min}(\alpha, w', d)|$$

and

$$\bar{b}_X(w) := \max_{b_X \in \Pi^B} \sum_{\substack{w' \in W \\ w \in \text{Min}(\alpha, w', d)}} b_X(w') / |\text{Min}(\alpha, w', d)|.$$

Note that

$$\sum_{\substack{w' \in W \\ w \in \text{Min}(\alpha, w', d)}} b(w') / |\text{Min}(\alpha, w', d)|$$

can be written in the form

$$\frac{\sum_{\substack{w' \in [w^1] \\ w \in \text{Min}(\alpha, w', d)}} b(w')}{1} + \dots + \frac{\sum_{\substack{w' \in [w^m] \\ w \in \text{Min}(\alpha, w', d)}} b(w')}{m}.$$

Then by Observation 1, $\bar{b}_X(w)$ is in Π_{bnd}^B . And also by Lemma 1, the belief state in Π_{bnd}^B identified by $\bar{b}_X(w)$ must be the one which maximizes

$$\sum_{\substack{w' \in W \\ w \in \text{Min}(\alpha, w', d)}} b_x(w') / |\text{Min}(\alpha, w', d)|,$$

where $b_x \in \Pi_{bnd}^B$. That is, $\bar{b}_x = \bar{b}_X$.

With a symmetrical argument, it can be shown that $\underline{y}^w = \underline{Y}^w$. \square

Informally, the following theorem says that the BB determined through the method of revising boundary belief states captures exactly the same beliefs and ignorance as the belief states in Π^B which have been revised. This correspondence relies on the fact that the upper and lower probability envelopes of Π^B can be induced from Π_{bnd}^B , which is what Lemma 2 states.

Theorem 1 Let $(\Pi_{\alpha}^B)^{\text{Gl}} := \{b_{\alpha}^{\text{Gl}} \in \Pi \mid b \in \Pi^B\}$. Let B_{bnd}^{α} be the BB induced from $(\Pi_{bnd}^B)^{\text{Gl}}_{\alpha}$. Then $\Pi_{bnd}^B = (\Pi^B)^{\text{Gl}}_{\alpha}$.

Proof:

We show that $\forall b' \in \Pi, b' \in \Pi_{bnd}^B \iff b' \in (\Pi^B)^{\text{Gl}}_{\alpha}$.

(\Rightarrow) $b' \in \Pi_{bnd}^B$ implies $\forall w \in W, \underline{y}^w \leq b'(w) \leq \bar{y}^w$ (by definition of B_{bnd}^{α}). Lemma 2 states that for all $w \in W, \bar{y}^w = \bar{Y}^w$ and $\underline{y}^w = \underline{Y}^w$. Hence, $\forall w \in W, \underline{Y}^w \leq b'(w) \leq \bar{Y}^w$. Therefore, $b'(w) \in (\Pi^B)^{\text{Gl}}_{\alpha}$.

(\Leftarrow) $b'(w) \in (\Pi^B)^{\text{Gl}}_{\alpha}$ implies $\forall w \in W, \underline{Y}^w \leq b'(w) \leq \bar{Y}^w$. Hence, by Lemma 2, $\forall w \in W, \underline{y}^w \leq b'(w) \leq \bar{y}^w$. Therefore, by definition of B_{bnd}^{α} , $b' \in \Pi_{bnd}^B$. \square

5 REVISION VIA OPTIMUM ENTROPY INFERENCE

Another approach to the revision of a belief base (BB) is to determine a representative of Π^B (call it b_{rep}), change the representative belief state via the defined revision operation and then induce a new BB from the revised representative belief state. Selecting a representative probability distribution from a family of such functions is not new [8, 22, e.g.]. More formally, given observation α , first determine $b_{rep} \in \Pi^B$, then compute its revision b_{rep}^{α} , and finally induce B^{α} from b_{rep}^{α} .

We shall represent Π^B (and thus B) by the single ‘least biased’ belief state, that is, the belief state in Π^B with *highest entropy*:

Definition 5 (Shannon Entropy)

$$H(b) := - \sum_{w \in W} b(w) \ln b(w),$$

where b is a belief state.

Definition 6 (Maximum Entropy) Traditionally, given some set of distributions Π , the most entropic distribution in Π is defined as

$$b^H := \arg \max_{b \in \Pi} H(b).$$

Suppose $B_2 = \{(-q \wedge \neg r) = 0.1\}$. Then the belief state $b \in \Pi^{B_2}$ satisfying the constraints posed by B_2 for which $H(b)$ is maximized is $b_{rep} = b^H = \langle 0.3, 0.3, 0.3, 0.1 \rangle$.

The above distribution can be found directly by applying the principle of maximum entropy: The true belief state is estimated to be the one consistent with known constraints, but is otherwise as unbiased as possible, or ‘‘Given no other knowledge, assume that everything is as random as possible. That is, the probabilities are distributed as uniformly as possible consistent with the available information,’’ [24]. Obviously world 00 must be assigned probability 0.1. And the remaining 0.9 probability mass should be uniformly spread across the other three worlds.

Applying Gl to b_{rep} on evidence $\neg q$ results in $b_{rep}^{\neg q} = \langle 0, 0, 0.6, 0.4 \rangle$.

Example 8 Suppose the vocabulary is $\{q, r\}$, $B_1 = \{(q) \geq 0.6\}$ and α is $(q \wedge \neg r) \vee (\neg q \wedge r)$. Then $b_{rep} = \arg \max_{b \in \Pi^{B_1}} H(b) = \langle 0.3, 0.3, 0.2, 0.2 \rangle$. Applying Gl to b_{rep} on α results in $b_{rep}^{\alpha} = \langle 0, 0.61, 0.39, 0 \rangle$. b_{rep}^{α} can be translated into B_{1rep}^{α} as $\{(q \wedge \neg r) = 0.61, (\neg q \wedge r) = 0.39\}$. \square

Still using $\alpha = (q \wedge \neg r) \vee (\neg q \wedge r)$, notice that $\Pi^{B_{1rep}^{\alpha}} \neq \Pi^{B_{1bnd}^{\alpha}}$. But how different are $B_{1rep}^{\alpha} = \{(q \wedge \neg r) = 0.61, (\neg q \wedge r) = 0.39\}$ and $B_{1bnd}^{\alpha} = \{(q \wedge r) \leq 0, (q \wedge \neg r) \geq 0.3, (\neg q \wedge r) \leq 0.7, (\neg q \wedge \neg r) \leq 0.0\}$? Perhaps one should ask, how different B_{1rep}^{α} is from the representative of B_{1bnd}^{α} : The least biased belief state satisfying B_{1bnd}^{α} is $\langle 0, 0.5, 0.5, 0 \rangle$. That is, How different are $\langle 0, 0.61, 0.39, 0 \rangle$ and $\langle 0, 0.5, 0.5, 0 \rangle$?

In the case of B_2 , we could compare $B_{2bnd}^{\neg q} = \{(\phi_{11}) \leq 0, (\phi_{10}) \leq 0, (\phi_{01}) \leq 0.9, (\phi_{00}) \geq 0.1\}$ with $b_{rep}^{\neg q} = \langle 0, 0, 0.6, 0.4 \rangle$. Or if we take the least biased belief state satisfying $B_{2bnd}^{\neg q}$, we can compare $\langle 0, 0, 0.5, 0.5 \rangle$ with $\langle 0, 0, 0.6, 0.4 \rangle$.

It has been extensively argued [11, 26, 23] that maximum entropy is a reasonable inference mechanism, if not the most reasonable one (w.r.t. probability constraints). On the other hand, the boundary belief states method also seems like a very reasonable inference mechanism for revising BBs as defined here, in the sense that it causes information loss. Resolving this misalignment in the results of the two methods is an obvious task for future research.

An extended version of maximum entropy is *minimum cross-entropy* (MCE) [17, 5]:

Definition 7 (Minimum Cross-Entropy) The ‘directed divergence’ of distribution c from distribution b is defined as

$$R(c, b) := \sum_{w \in W} c(w) \ln \frac{c(w)}{b(w)}.$$

$R(c, b)$ is undefined when $b(w) = 0$ while $c(w) > 0$; when $c(w) = 0$, $R(c, b) = 0$, because $\lim_{x \rightarrow 0} \ln(x) = 0$. Given new evidence

$\phi \in L^{prob}$, the distribution c satisfying ϕ diverging least from current belief state b is

$$\arg \min_{c \in \Pi, c \Vdash \phi} R(c, b).$$

Definition 8 (MCI) Then MCE inference (denoted (MCI)) is defined as

$$b \text{ MCI } \alpha := \arg \min_{b' \in \Pi, b' \Vdash (\alpha)=1} R(b', b).$$

In the following example, we interpret revision as MCE inference.

Example 9 Suppose the vocabulary is $\{q, r\}$ and $B_1 = \{(q) \geq 0.6\}$. Let α be $(q \wedge \neg r) \vee (\neg q \wedge r)$. Then

$$\begin{aligned} \Pi_{bnd}^{B_1} = & \{ \{(11, 1.0), (10, 0.0), (01, 0.0), (00, 0.0)\}, \\ & \{(11, 0.0), (10, 1.0), (01, 0.0), (00, 0.0)\}, \\ & \{(11, 0.6), (10, 0.0), (01, 0.4), (00, 0.0)\}, \\ & \{(11, 0.6), (10, 0.0), (01, 0.0), (00, 0.4)\}, \\ & \{(11, 0.0), (10, 0.6), (01, 0.4), (00, 0.0)\}, \\ & \{(11, 0.0), (10, 0.6), (01, 0.0), (00, 0.4)\} \}, \end{aligned}$$

$$\begin{aligned} (\Pi_{bnd}^{B_1})_{\alpha}^{\text{MCI}} = & \{ \{(11, 0), (10, 0), (01, 1), (00, 0)\}, \\ & \{(11, 0), (10, 1), (01, 0), (00, 0)\}, \\ & \{(11, 0), (10, 0.6), (01, 0.4), (00, 0)\} \} \text{ and} \end{aligned}$$

$$B_{1bnd}^{\alpha} = \{(\phi_{11}) \leq 0, (\phi_{00}) \leq 0\}.$$

Note that if we let $B' = \{(q \wedge \neg r) \vee (\neg q \wedge r) = 1\}$, then $\Pi^{B'} = \Pi_{bnd}^{B_1}$. \square

Recall from Example 6 that B' included $(q \wedge \neg r) \geq 0.3$. Hence, in this particular case, combining the boundary belief states approach with MCI results in a less informative revised belief base than when GI is used. The reason for the loss of information might be due to $R(\cdot, \{(11, 1.0), (10, 0.0), (01, 0.0), (00, 0.0)\})$ and $R(\cdot, \{(11, 0.6), (10, 0.0), (01, 0.0), (00, 0.4)\})$ being undefined: Recall that $R(c, b)$ is undefined when $b(w) = 0$ while $c(w) > 0$. But then there is no belief state c for which $c \Vdash \alpha$ and $R(\cdot)$ is defined (with these two belief states as arguments). Hence, there are no revised counterparts of these two belief states in $(\Pi_{bnd}^{B_1})_{\alpha}^{\text{MCI}}$. We would like to analyse MCI more within this framework. In particular, in the future, we would like to determine whether a statement like Theorem 1 holds for MCI too.

In MCE inference, b -consistency of evidence ϕ is defined as: There exists a belief state c such that $c \Vdash \phi$ and c is *totally continuous* with respect to b (i.e., $b(w) = 0$ implies $c(w) = 0$). MCE is undefined when the evidence is not b -consistent. This is analogous to Bayesian conditioning being undefined for $b(\alpha) = 0$. Obviously, this is a limitation of MCE because some belief states may not be considered as candidate revised belief states.

6 RATIONALITY POSTULATES

In this section, we assess the operation of revising a belief base B by α via GI (denoted B^{α}) with respect to several rationality postulates. Katsuno and Mendelzon [12] modified the eight AGM belief revision postulates [1] to the following six ((R1)-(R6)). The postulates are intended to be an ideal standard of rationality with respect to the behavior of any revision operator. We shall simply translate (\implies) each postulate into one appropriate for belief bases (in the notation of this paper) in order to get a sense of which ones are satisfied

and which not. We then intend to use this information to guide us in drawing up appropriate generalised postulates in the future. ⁶ In their notation [12], given a knowledge base represented by a propositional sentence ψ and an observation represented by a propositional sentence μ , $\psi \circ \mu$ denotes the revision of ψ by μ .⁷

$$\begin{aligned} \text{(R1)} \quad & \psi \circ \mu \text{ implies } \mu \implies \\ \text{(RB1)} \quad & B^{\alpha} \models (\alpha) = 1. \end{aligned}$$

Proposition 1 Postulate (RB1) holds.

Proof:

By definition of GI, all probability mass is shifted to closest α -worlds. For every revised boundary belief state b_{bnd} , it is thus the case that $b_{bnd}(\alpha) = 1$. By the process of inducing B^{α} , it must be that B^{α} entails $(\alpha) = 1$. \square

$$\text{(R2)} \quad \text{If } \psi \wedge \mu \text{ is satisfiable, then } \psi \circ \mu \equiv \psi \wedge \mu \implies$$

$$\text{(RB2)} \quad \text{If } B \cup \{(\alpha) = 1\} \text{ is satisfiable, then } B^{\alpha} \equiv B \cup \{(\alpha) = 1\}.$$

Proposition 2 Postulate (RB2) does not hold.

Proof:

Recall that $B_1 = \{(q) \geq 0.6\}$. Example 6 shows that $B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)} \equiv B' = \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1, (q \wedge \neg r) \geq 0.3\}$. Therefore, $B' \models (q \wedge \neg r) \geq 0.3$. But $B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}$, although satisfiable, does not entail $(q \wedge \neg r) \geq 0.3$. Hence, $B' \not\equiv B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}$. \square

Let $B_4 = \{(q \wedge r) = 0.2, (q \wedge \neg r) = 0.8\}$. Then $B_4^{\alpha} \equiv B_4 \equiv B_4 \cup \{(q) = 1\}$. Notice that B_4 specifies a particular belief state. One can see that Postulate (RB2) will hold whenever the BB to be revised specifies a particular belief state (i.e., whenever $|\Pi^B| = 1$).

$$\text{(R3)} \quad \text{If } \mu \text{ is satisfiable, then } \psi \circ \mu \text{ is also satisfiable} \implies$$

$$\text{(RB3)} \quad \text{If } (\alpha) = 1 \text{ is satisfiable, then } B^{\alpha} \text{ is also satisfiable.}$$

Proposition 3 Postulate (RB3) holds.

Proof:

$(\alpha) = 1$ is satisfiable iff α is. And if α is satisfiable, then every boundary belief state of B revised by α is defined. Hence, B^{α} must be defined, i.e., satisfiable. \square

$$\text{(R4)} \quad \text{If } \psi \equiv \psi' \text{ and } \mu \equiv \mu', \text{ then } \psi \circ \mu \equiv \psi' \circ \mu' \implies$$

$$\text{(RB4)} \quad \text{If } B \equiv C \text{ and } \alpha \equiv \beta, \text{ then } B^{\alpha} \equiv C^{\beta}.$$

Proposition 4 Postulate (RB4) holds.

Proof:

By definition, if $B \equiv C$, then $\Pi^B = \Pi^C$. Hence, the boundary belief states for B and C are the same. And given $\alpha \equiv \beta$, it must be the case that $B^{\alpha} \equiv C^{\beta}$. \square

$$\text{(R5)} \quad (\psi \circ \mu) \wedge \phi \text{ implies } \psi \circ (\mu \wedge \phi) \implies$$

$$\text{(RB5)} \quad B^{\alpha} \cup \{(\beta) = 1\} \models B^{\alpha \wedge \beta}.$$

Proposition 5 Postulate (RB5) does not hold.

Proof:

Let α be \top . Then $B^{\alpha} \equiv B$ and $B^{\alpha \wedge \beta} \equiv B^{\beta}$. Now let β be $(q \wedge \neg r) \vee (\neg q \wedge r)$. We can thus ask whether $B \cup \{(\beta) = 1\} \models B^{\beta}$.

Consider Example 6. Recall that $B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)} \models (q \wedge \neg r) \geq 0.3$. That is, $\forall b \in \Pi_{bnd}^{B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)}}$, $b \models (q \wedge \neg r) \geq 0.3$. And recall

⁶ In these postulates, it is sometimes necessary to write an observation α as a BB, i.e., as $\{(\alpha) = 1\}$ – in the present framework, observations are regarded as certain.

⁷ \circ is some revision operator.

that $B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\} \not\models (q \wedge \neg r) \geq 0.3$. That is, $\exists b \in \Pi^{B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}}$ s.t. $b \not\models (q \wedge \neg r) \geq 0.3$. Hence, $\Pi^{B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}} \not\subseteq \Pi^{B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)}}$ and by the definition of \models , $B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\} \not\models B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)}$. \square

If $B^\alpha \cup \{(\beta) = 1\}$ is not satisfiable, then (RB5) holds trivially. If $B^\alpha \cup \{(\beta) = 1\}$ is satisfiable, then α logically entails β , implying that $\alpha \wedge \beta \equiv \alpha$. Then we can ask whether $B^\alpha \cup \{(\beta) = 1\} \models B^\alpha$? Assuming B specifies a particular belief state, clearly $\Pi^{B^\alpha \cup \{(\beta) = 1\}} \subseteq \Pi^{B^\alpha}$ and (RB5) holds. So the difficulties come in when $|\Pi^B| > 1$.

(R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable, then $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$
 \implies

(RB6) If $B^\alpha \cup \{(\beta) = 1\}$ is satisfiable, then $B^{\alpha \wedge \beta} \models B^\alpha \cup \{(\beta) = 1\}$.

Proposition 6 *Postulate (RB6) does not hold.*

Proof:

Let α be \top . Then (RB6) becomes: If $B \cup \{(\beta) = 1\}$ is satisfiable, then $B^\beta \models B \cup \{(\beta) = 1\}$.

Consider Example 6. Let β be $(q \wedge \neg r) \vee (\neg q \wedge r)$. Note that $B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}$ is satisfiable. Then by (RB6), $B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)} \models B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}$. But this is false:

Let $b = \langle 0, 0.5, 0.5, 0 \rangle$. Then

$$b \in \Pi^{B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)}} = \Pi^{\{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1, (q \wedge \neg r) \geq 0.3\}}, \text{ but}$$

$$b \notin \Pi^{B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}} = \Pi^{\{(q \geq 0.6, ((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}}.$$

Hence, $\Pi^{B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)}} \not\subseteq \Pi^{B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}}$. Therefore, $B_1^{(q \wedge \neg r) \vee (\neg q \wedge r)} \not\models B_1 \cup \{((q \wedge \neg r) \vee (\neg q \wedge r)) = 1\}$. \square

From the discussion of Postulate (RB5), we notice that given the antecedent of (RB6), $B^\alpha \cup \{(\beta) = 1\} \equiv B^\alpha$. Thus, the consequent of (RB6) becomes $B^\alpha \models B^\alpha$. This means that (RB6) holds when $|\Pi^B| = 1$. Again, the difficulties come in when $|\Pi^B| > 1$.

7 RELATED WORK

Voorbraak [27] proposed the partial probability theory (PTT), which allows probability assignments to be partially determined, and where there is a distinction between probabilistic information based on (i) hard background evidence and (ii) some assumptions. He does not explicitly define the “constraint language”, however, from his examples and discussions, one can infer that he has something like the language $L^{P^{TT}}$ in mind: it contains all formulae which can be formed with sentences in our L^{prob} in combination with connectives \neg, \wedge and \vee . A “belief state” in PTT is defined as the quadruple $\langle \Omega, \mathcal{B}, \mathcal{A}, \mathcal{C} \rangle$, where Ω is a sample space, $\mathcal{B} \subset L^{P^{TT}}$ is a sets of probability constraints, $\mathcal{A} \subset L^{P^{TT}}$ is a sets of assumptions and $\mathcal{C} \subseteq W$ “represents specific information concerning the case at hand” (an observation or evidence).⁸ Our epistemic state can be expressed as a restricted PTT “belief state” by letting $\Omega = W$, $\mathcal{B} = B$, $\mathcal{A} = \emptyset$ and $\mathcal{C} = \{w \in W \mid w \models \alpha\}$, where B is a belief base and α is an observation in our notation.

Voorbraak [27] mentions that he will only consider conditioning where the evidence does not contradict the current beliefs. He defines the set of belief states corresponding to the conditionalized PPT

⁸ Voorbraak’s “belief state” would rather be called an *epistemic state* or *knowledge structure* in our language.

“belief state” as $\{b(\cdot \mid C) \in \Pi \mid b \in \Pi^{B \cup \mathcal{A}}, b(C) > 0\}$. In our notation, this corresponds to $\{(b \text{BC } \alpha) \in \Pi \mid b \in \Pi^B, b(\alpha) > 0\}$, where α corresponds to C . Voorbraak [27] proposes *constraining* as an alternative to conditioning: Let $\phi \in L^{prob}$ be a probability constraint. In our notation, constraining Π^B on ϕ produces $\Pi^{B \cup \{\phi\}}$.

Note that expanding a belief set reduces the number of models (worlds) and expanding a PPT “belief state” with extra constraints also reduces the number of models (belief states / probability distributions).

In the context of belief sets, it is possible to obtain any belief state from the ignorant belief state by a series of expansions.

In PPT, constraining, but not conditioning, has the analogous property. This is one of the main reasons we prefer to constraining and not conditioning to be the probabilistic version of expansion. [27, p. 4]

But Voorbraak does not address the issue that C and ϕ are different kinds of observations, so constraining, as defined here, cannot be an alternative to conditioning. C cannot be used directly for constraining and ϕ cannot be used directly for conditioning.

W.l.o.g., we can assume C is represented by α . If we take $b \text{GI } \alpha$ to be an expansion operation whenever $b(\alpha) > 0$, then one might ask, Is it possible to obtain any belief base B' from the ignorant belief base $B = \emptyset$ by a series of expansions, using our approach? The answer is, No. For instance, there is no observation or series of observations which can change $B = \{\}$ into $B' = \{(q) \geq 0.6\}$. But if we were to allow sentences (constraints) in L^{prob} to be observations, then we could obtain any B' from the ignorant B .

Grove and Halpern [9] investigate what “update” (incorporation of an observation with current beliefs, such that the observation does not contradict the beliefs) means in a framework where beliefs are represented by a set of belief states. They state that the main purpose of their paper is to illustrate how different the set-of-distributions framework can be, “technically”, from the standard single-distribution framework. They propose six postulates characterizing what properties an update function should have. They say that some of the postulates are obvious, some arguable and one probably too strong. Out of seven (families of) update functions only the one based on conditioning ($Upd_{cond}(\cdot)$) and the one based on constraining ($Upd_{constrain}(\cdot)$) satisfy all six postulates, where $Upd_{cond}(\Pi^B, \alpha) := \{(b \text{BC } \alpha) \in \Pi \mid b \in \Pi^B, b(\alpha) > 0\}$ and where they interpret Voorbraak’s constraining [27] as $Upd_{constrain}(\Pi^B, \alpha) := \{b \in \Pi^B \mid b(\alpha) = 1\}$. Grove and Halpern [9] do not investigate the case when an observation must be incorporated while it is (possibly) inconsistent with the old beliefs (i.e., revision). It would be interesting to analyse the present work against their six postulates.

Kern-Isberner [14] develops a new perspective of probabilistic belief change. Based on the ideas of Alchourrón et al. [1] and Katsuno and Mendelzon [12] (KM), the operations of revision and update, respectively, are investigated within a probabilistic framework. She employs as basic knowledge structure as belief base (b, \mathcal{R}) , where b is a probability distribution (belief state) of background knowledge and \mathcal{R} is a set of probabilistic conditionals of the form $A \rightsquigarrow B[x]$ meaning ‘The probability of B , given A , is x ’. A universal inference operation – based on the techniques of optimum entropy – is introduced as an “adequate and powerful method to realize probabilistic belief change”.

By having a belief state available in the belief base, minimum cross-entropy can be used. The intention is then that an agent with belief base (b, \mathcal{T}) should always reason w.r.t. belief state $b^T := \arg \min_{c \in \Pi, c \models \mathcal{T}} R(c, b)$. Kern-Isberner [14] then defines the prob-

abilistic belief revision of (b, \mathcal{R}) by evidence \mathcal{S} as $(b, \mathcal{R} \cup \mathcal{S})$. And the probabilistic belief update of (b, \mathcal{R}) by evidence \mathcal{S} is defined as $(b^{\mathcal{R}}, \mathcal{S})$.⁹ She distinguishes between revision as a knowledge adding process, and updating as a change-recording process. Kern-Isberner [14] sets up comparisons of maximum cross-entropy belief change with AGM revision and KM update. Cases where, for update, new information \mathcal{R} is inconsistent with the prior distribution b , or, for revision, is inconsistent with b or the context \mathcal{R} , are not dealt with [14, p. 399, 400].

Having a belief state available for modification when new evidence is to be adopted is quite convenient. As Voorbraak [27] argues, however, an agent's ignorance can hardly be represented in an epistemic state where a single belief state must always be chosen.

We would like to investigate the representation of conditional probabilistic information such as is done in the work of Kern-Isberner [14, 15] and Yue and Liu [29], for instance.

Yue and Liu [29] propose a probabilistic revision operation for imprecise probabilistic beliefs in the framework of Probabilistic Logic Programming (PLP). New evidence may be a probabilistic (conditional) formula and needs not be consistent with the original beliefs. Revision via imaging (e.g., GI) also overcomes this consistency issue. Essentially, their *probabilistic epistemic states* Ψ are induced from a PLP program which is a set of formulae, each formula having the form $(\psi \mid \phi)[l, u]$, meaning that the probability of the conditional $(\psi \mid \phi)$ lies in the interval $[l, u]$. The operator they propose has the characteristic that if an epistemic state Ψ represents a single probability distribution, revising collapses to Jeffrey's rule and Bayesian conditioning. They mention that it is required that the models (distributions) of Ψ is a convex set. There might thus be an opportunity to employ their revision operation on a representative set of boundary distributions as proposed in this paper.

Another PLP, proposed by Michels et al. [21], also allow for incomplete specification of probabilities. The language of their logic is however more expressive than ours, and they focus on inference (probabilistic query answering). They mention that they want to allow for learning in their system – one could possibly interpret some kinds of learning in this setting as belief revision.

8 CONCLUSION AND FUTURE DIRECTIONS

In this paper, we propose an approach how to generate a new probabilistic belief base from an old one, given a new piece of non-probabilistic information, where a belief base is a finite set of sentences, each sentence stating the likelihood of a proposition about the world. In this framework, an agent's belief base represents the set of belief states compatible with the sentences in it. In this sense, the agent is able to represent its knowledge *and* ignorance about the true state of the world.

We used a version of the so-called *imaging* approach to implement the revision operation.

Three methods were proposed: revising a finite set of 'boundary belief states' via generalized imaging, revising a finite set of 'boundary belief states' via minimum cross-entropy and revising a least biased belief state. We focussed on the first method and showed that the latter two give different results.

There were two main contribution of this paper. The first was to prove that the set of belief states satisfying B_{new} is exactly those belief states satisfying the original belief base, revised. The second was to uncover an interesting conflict in the results of the three belief

base revision methods. It is worth further understanding the reasons behind such a difference, as such an investigation could give more insight about the mechanisms behind the methods and indicate possible pros and cons of each. Importantly, further analysis with respect to rationality postulated is necessary, as mentioned in § 8. Such an analysis may also bring insights into the differing results.

The computational complexity of $Min(\cdot)$ is in $O(|W|^2)$ and the complexity of GI is thus in $O(|W|^2|W|^2) = O(|W|^4)$ in the worst case. However, this complexity is highly dependent on the observation and the distance function. Note that $|\Pi_{bnd}^B| \leq |W^{perm}| = |W|!$. GI is applied to every belief state in Π_{bnd}^B . Hence, the complexity of the method, in the worst case, is in $O(|W|!|W|^4)$.

As far as we know, there is no analytic solution to determine the distribution in Π^B with maximum entropy / minimum cross-entropy. To narrow in on the computational complexity of these methods, we would have to know what class of optimization problem they are (convex?), and thus what techniques are used to solve them. Our knowledge in this area is lacking and it would require some time for investigation.

The proposal or design, and justification of rationality postulates similar to (RB1)-(RB6) in the section above is called for. An analysis of the postulates with respect to the revision operation must then be carried out. An attempt may be made to design the revision operation so as to make more of the postulates hold, or we may attempt to justify why our revision process does not / need not satisfy postulates it fails at.

Given that we have found that the belief base resulting from revising via the boundary-belief-states approach differs from the belief base resulting from revising via the representative-belief-state approach, the question arises, When is it appropriate to use a representative belief state defined as the most entropic belief state of a given set Π^B ? This is an important question, especially due to the popularity of employing the Maximum Entropy principle in cases of underspecified probabilistic knowledge [11, 8, 10, 27, 14, 16] and the principle's well-behavedness [26, 22, 13].

As far as we know, imaging for belief change has never been applied to (conditional) probabilistic evidence. Due to issues with many revision methods required to be consistent with prior beliefs, and imaging not having this limitation, it might be worthwhile investigating.

The translation from the set of belief states back to a belief base is a mapping from every belief state to a probability formula. The size of the belief base is thus in the order of $|W^{perm}|$, where $|W|$ is already exponential in the size of \mathcal{P} , the set of atoms. As we saw in several examples in this paper, the new belief base often has a more concise equivalent counterpart. It would be useful to find a way to consistently determine more concise belief bases than our present approach does.

Does a similar result as Theorem 1 holds for Bayesian conditioning? This is an important question we would like to answer and which credal set theory [19, 28] may answer.

ACKNOWLEDGEMENTS

The work of Giovanni Casini has been supported by the Fonds National de la Recherche, Luxembourg, and cofunded by the Marie Curie Actions of the European Commission (FP7-COFUND) (AFR/9181001). Thomas Meyer has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 690974. We thank the reviewers for their valuable comments and suggestions.

⁹ This is a very simplified version of what she presents. Please refer to the paper for details.

REFERENCES

- [1] C. Alchourrón, P. Gärdenfors, and D. Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic*, **50**(2), 510–530, (1985).
- [2] C. Boutilier, 'A unified model of qualitative belief change: A dynamical systems perspective', *Artif. Intell.*, **98**(1–2), 281–316, (1998).
- [3] K. Chhogyal, A. Nayak, R. Schwitter, and A. Sattar, 'Probabilistic belief revision via imaging', in *PRICAI 2014: Trends in Artif. Intell.: Thirteenth Pacific Rim Intl. Conf. on Artif. Intell.*, eds., D.-N. Pham and S.-B. Park, volume 8862 of *LNAI*, 694–707, Springer Intl. Publishing, Switzerland, (2014).
- [4] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, New York, 1991.
- [5] I. Csiszár, 'I-divergence geometry of probability distributions and minimization problems', *Annals of Probability*, **3**, 146–158, (1975).
- [6] D. Dubois and H. Prade, 'Belief revision and updates in numerical formalisms: An overview, with new results for the possibilistic framework', in *Proceedings of the Thirteenth Intl. Joint Conf. on Artif. Intell. (IJCAI93)*, pp. 620–625, San Francisco, CA, USA, (1993). Morgan Kaufmann Publishers Inc.
- [7] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, Massachusetts/England, 1988.
- [8] M. Goldszmidt, P. Morris, and J. Pearl, 'A maximum entropy approach to nonmonotonic reasoning', in *Proceedings of the Eighth Natl. Conf. on Artif. Intell. (AAAI-90)*, pp. 646–652. AAAI Press, (1990).
- [9] A. Grove and J. Halpern, 'Updating sets of probabilities', in *Proceedings of the Fourteenth Conf. on Uncertainty in Artif. Intell.*, UAI'98, pp. 173–182, San Francisco, CA, USA, (1998). Morgan Kaufmann.
- [10] D. Hunter, 'Maximum entropy updating and conditionalization', in *Existence and Explanation*, eds., W. Spohn, B. Van Fraassen, and B. Skyrms, volume 49 of *The University of Western Ontario Series in Philosophy of Science*, 45–57, Springer Netherlands, (1991).
- [11] E. Jaynes, 'Where do we stand on maximum entropy?', in *The Maximum Entropy Formalism*, 15–118, MIT Press, (1978).
- [12] H. Katsuno and A. Mendelzon, 'On the difference between updating a knowledge base and revising it', in *Proceedings of the Second Intl. Conf. on Principles of Knowledge Representation and Reasoning*, pp. 387–394, (1991).
- [13] G. Kern-Isberner, 'Characterizing the principle of minimum cross-entropy within a conditional-logical framework', *Artif. Intell.*, **98**(12), 169 – 208, (1998).
- [14] G. Kern-Isberner, 'Revising and updating probabilistic beliefs', in *Frontiers in Belief Revision*, eds., M.-A. Williams and H. Rott, volume 22 of *Applied Logic Series*, 393–408, Kluwer Academic Publishers, Springer Netherlands, (2001).
- [15] G. Kern-Isberner, 'Linking iterated belief change operations to non-monotonic reasoning', in *Proceedings of the Eleventh Intl. Conf. on Principles of Knowledge Representation and Reasoning*, pp. 166–176, Menlo Park, CA, (2008). AAAI Press.
- [16] G. Kern-Isberner and W. Rödder, 'Belief revision and information fusion on optimum entropy', *Intl. Journal of Intelligent Systems*, **19**(9), 837–857, (2004).
- [17] S. Kullback, *Information theory and statistics*, volume 1, Dover, New York, 2nd edn., 1968.
- [18] D. Lehmann, M. Magidor, and K. Schlechta, 'Distance semantics for belief revision', *Journal of Symbolic Logic*, **66**(1), 295–317, (2001).
- [19] I. Levi, *The enterprise of knowledge*, The MIT press, 1983.
- [20] D. Lewis, 'Probabilities of conditionals and conditional probabilities', *Philosophical Review*, **85**(3), 297–315, (1976).
- [21] S. Michels, A. Hommersom, P. Lucas, M. Velikova, and P. Koopman, 'Inference for a new probabilistic constraint logic', in *Proceedings of the Twenty-Third Intl. Joint Conf. on Artif. Intell. (IJCAI-13)*, ed., F. Rossi, pp. 2540–2543, Menlo Park, CA, (August 2013). AAAI Press.
- [22] J. Paris, *The Uncertain Reasoner's Companion: A Mathematical Perspective*, Cambridge University Press, Cambridge, 1994.
- [23] J. Paris and A. Vencovsk, 'In defense of the maximum entropy inference process', *Intl. Journal of Approximate Reasoning*, **17**(1), 77–103, (1997).
- [24] D. Poole and A. Mackworth, *Artif. Intell.: Foundations of Computational Agents*, Cambridge University Press, New York, USA, 2010.
- [25] G. Rens and T. Meyer, 'A new approach to probabilistic belief change', in *Proceedings of the Intl. Florida AI Research Society Conf. (FLAIRS)*, eds., I. Russell and W. Eberle, pp. 582–587. AAAI Press, (2015).
- [26] J. Shore and R. Johnson, 'Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy', *Information Theory, IEEE Transactions on*, **26**(1), 26–37, (Jan 1980).
- [27] F. Voorbraak, 'Partial Probability: Theory and Applications', in *Proceedings of the First Intl. Symposium on Imprecise Probabilities and Their Applications*, pp. 360–368, (1999). url: decsai.ugr.es/ smc/isipta99/proc/073.html.
- [28] P. Walley, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, 1991.
- [29] A. Yue and W. Liu, 'Revising imprecise probabilistic beliefs in the framework of probabilistic logic programming', in *Proceedings of the Twenty-third AAAI Conf. on Artif. Intell. (AAAI-08)*, pp. 590–596, (2008).