

Revising by an Inconsistent Set of Formulas

James Delgrande

School of Computing Science

Simon Fraser University

Burnaby, B.C. V5A 1S6

Canada

jim@cs.sfu.ca

Abstract

This paper presents an approach to belief revision in which revision is a function from a belief state and a finite set of formulas to a new belief state. In the interesting case, the set for revision S may be inconsistent but individual members of S are consistent. We argue that S will still contain interesting information regarding revision; in particular, maximum consistent subsets of S will determine candidate formulas for the revision process, and the agent's associated faithful ranking will determine the plausibility of such candidate formulas. Postulates and semantic conditions characterizing this approach are given, and representation results are provided. As a consequence of this approach, we argue that revision by a sequence of formulas, usually considered as a problem of iterated revision, is more appropriately regarded as revision by the possibly-inconsistent set of these formulas. Hence we suggest that revision by a sequence of formulas is foremost a problem of (uniterated) set revision.

1 Introduction

The area of *belief revision* studies how an agent may incorporate new information about a domain into its set of beliefs. The original and best-known approach to belief revision is the so-called *AGM approach* [Alchourrón *et al.*, 1985; Gärdenfors, 1988]. In this approach, an agent's belief state is characterised by a deductively-closed set of formulas, or *belief set*, K . The revision of K by formula ϕ , denoted $K * \phi$, is a belief set K' such that $\phi \in K'$, and where K' is consistent whenever ϕ is. Revision functions may be characterised in one of two ways: on the one hand, *rationality criteria* or postulates may be given that any "reasonable" revision function should satisfy; on the other hand, formal constructions for revision functions may be developed. Ideally, classes of functions conforming to a set of postulates are linked to a corresponding formal construction via a *representation result*.

Several assumptions about revision functions appear to be incontrovertible. First, it is assumed that the domain of application is static¹ and the task of revision is to improve an

agent's knowledge about the domain. As well, a principle of *informational economy* requires that in changing its beliefs, an agent will retain as many of the old beliefs as possible. Furthermore, if one revises K by ϕ , ϕ is believed in the result, i.e. $\phi \in K * \phi$. Thus in belief revision, the prerequisite question of whether the agent *should* accept ϕ has been resolved in favour of acceptance, and the issue is *how* the formula is to be incorporated into the belief corpus.

One aspect of revision not addressed in the AGM approach is that, even in an unchanging domain, an agent will typically receive not just a single item of information, but rather a sequence of items. (For example, an agent may be exploring a room while leaving the room unaltered.) Subsequently, there has been a great deal of attention paid to *iterated belief revision*, which addresses logical relations among sequences of revisions. Hence, if an agent with belief set K were to receive the information that ϕ followed by μ , this would be represented by $(K * \phi) * \mu$. Crucially we would have $\mu \in (K * \phi) * \mu$.

However, we suggest that these assumptions don't sit well together. First, if the domain is static then there is no a priori reason to prefer more recent information. That is, there *may* sometimes be a reason to prefer more recently received information but it is unreasonable to require that this is necessarily *always* the case. In fact, [Nayak *et al.*, 2003] presents an approach where $(K * \phi) * \mu$ is the same as $K * (\phi \wedge \mu)$ when $\phi \wedge \mu$ is consistent. But what if $\phi \wedge \mu$ is not consistent? There are two alternatives in this case, both of which we argue are unpalatable. First, one may preferentially accept the more recently received information. But again, if an agent is receiving independent reports about an unchanging domain, such an assumption is unjustified. Second, and worse, the agent may accept the conjunction $\phi \wedge \mu$, in which case it falls into inconsistency.

The approach presented here takes the view that in a static domain, the order in which formulas are received is irrelevant, and so revision should be with respect to the set of such formulas. Thus if the agent received the information that ϕ and then μ , the revision would be expressed as $K * \{\phi, \mu\}$. While there has been previous work in revising by a set of formulas (see the next section), crucially, all such work assumes that if the set of formulas is inconsistent then the result of revision is an inconsistent belief set.

accurate is to say that the underlying language is static; see [Friedman and Halpern, 1999] or [Lang, 2006].

¹This is slightly misleading, but serves for our purposes. More

Our point of departure then is to address revision by a set of formulas S where S may be inconsistent. Consider for example where we wish to revise by a set of formulas $\{a \wedge b, a \wedge \neg b\}$ or $\{a, b, \neg b\}$ or $\{a, b \wedge \neg b\}$. Each set is inconsistent, but in each case an argument can be made for accepting a . We address this via an appeal to *informational economy* with respect to the set S for revision: If the agent is to incorporate S consistently into its belief corpus, but S cannot be consistently so incorporated, then maximal consistent subsets of S comprise suitable candidates for revision. Thus for the set $\{a, b, \neg b\}$, candidates for revision are $\{a, b\}$ and $\{a, \neg b\}$. Moreover, there is a means in place for adjudicating among these candidate sets, in that an agent will have an implicit plausibility ordering over formulas as part of its belief state. We give a set of postulates for revision by a finite, possibly-inconsistent set of formulas, and a representation result is given in terms of *faithful rankings*. We suggest that this approach is interesting because it generalises the AGM framework. As well, we argue that iterated belief revision should in fact be treated as (uniterated) revision by a set of formulas.

The next section introduces terminology, and reviews the area of belief revision. Section 3 further motivates the approach and presents the formal details, while the next section discusses implications of the approach with respect to iterated revision. We finish with a brief conclusion.

2 Background

2.1 Formal Preliminaries

Let $\mathcal{P} = \{a, b, c, \dots\}$ be a fixed, finite set of propositional variables. $\mathcal{L}$ is the language of classical propositional logic over $\mathcal{P}$, with the usual connectives $\wedge, \vee, \supset, \neg$. The classical consequence relation is denoted $\vdash$. $Ch(S)$ is the set of logical consequences of a set of formulas S , that is $Ch(S) = \{\phi \in \mathcal{L} \mid S \vdash \phi\}$. $\top$ stands for some arbitrary tautology, and $\perp$ is defined to be $\neg\top$.

A propositional *interpretation* (or *possible world*) is a mapping from $\mathcal{P}$ to $\{\text{true}, \text{false}\}$. The set of all interpretations is denoted by $\mathcal{M}$. A *model* of a formula ϕ is an interpretation w that makes ϕ true according to the usual definition of truth, and is denoted by $w \models \phi$. For $W \subseteq \mathcal{M}$, we also write $W \models \phi$ if $w \models \phi$ for every $w \in W$. For a set of sentences S , $Mod(S)$ is the set of all models of S . $Mod(\{\phi\})$ is also written as $Mod(\phi)$. For $W \subseteq \mathcal{M}$, we denote by $\mathcal{T}(W)$ the set of sentences that are true in all elements of W , that is $\mathcal{T}(W) = \{\phi \in \mathcal{L} \mid w \models \phi \text{ for all } w \in W\}$.

For sets of formulas S and S' , $S + S'$ denotes the *expansion* of S by S' , that is $S + S' = Ch(S \cup S')$. For formula ϕ , $S + \phi$ is defined to be $S + \{\phi\}$. Formulas ϕ and ψ are *logically equivalent*, written $\phi \equiv \psi$, iff $\phi \vdash \psi$ and $\psi \vdash \phi$. This extends to sets of formulas by: $S \equiv S'$ iff $S \vdash \phi$ for every $\phi \in S'$ and $S' \vdash \psi$ for every $\psi \in S$. Two sets of formulas S and S' are *strongly equivalent* just if for every $\phi \in S$ there is $\phi' \in S'$ such that $\phi \equiv \phi'$, and vice versa. A set of formulas S (or formula) is *inconsistent* if $S \vdash \perp$, and *strongly inconsistent* just if $S \neq \emptyset$ and for every $\phi \in S$, ϕ is inconsistent.

For a (finite) set of formulas S , $\wedge S$ is the conjunction of members of S , and $\bar{S} = \{\neg\phi \mid \phi \in S\}$. If S is a finite set

of finite sets of formulas, define $\bigvee S$ to be $\bigvee_{S \in \mathbf{S}} \wedge S$. As a degenerate case, define $\bigvee \{\emptyset\}$ to be $\perp$.

A total preorder $\preceq$ is a reflexive, transitive binary relation, such that either $\phi \preceq \psi$ or $\psi \preceq \phi$ for every ϕ, ψ . The strict part of $\preceq$ is denoted by $\prec$, that is, $\phi \prec \psi$ and $\psi \not\prec \phi$. As usual, $\phi = \psi$ abbreviates $\phi \preceq \psi$ and $\psi \preceq \phi$. Given a set S and total preorder $\preceq$ defined on members of S , we denote the set of minimal elements of S in $\preceq$ by $\min(S, \preceq)$.

2.2 Belief Revision

In the AGM approach, an agent's beliefs are modelled by a deductively closed set of formulas called a belief set. However, various researchers have subsequently argued that, in order to address iterated belief revision, it is more appropriate to consider *belief states* or *epistemic states* as objects of revision. An epistemic state $\mathcal{K}$ effectively encodes information regarding how the revision function itself changes under a revision. The belief set corresponding to belief state $\mathcal{K}$ is denoted $Bel(\mathcal{K})$. A revision operator $*$ then maps a belief state $\mathcal{K}$ and formula ϕ to a revised belief state $\mathcal{K} * \phi$. In the spirit of [Darwiche and Pearl, 1997], the AGM postulates for revision can be expressed as follows:

- ($\mathcal{K} * 1$) $Bel(\mathcal{K} * \phi) = Ch(Bel(\mathcal{K} * \phi))$
- ($\mathcal{K} * 2$) $\phi \in Bel(\mathcal{K} * \phi)$
- ($\mathcal{K} * 3$) $Bel(\mathcal{K} * \phi) \subseteq Bel(\mathcal{K}) + \phi$
- ($\mathcal{K} * 4$) If $\neg\phi \notin Bel(\mathcal{K})$ then $Bel(\mathcal{K}) + \phi \subseteq Bel(\mathcal{K} * \phi)$
- ($\mathcal{K} * 5$) $Bel(\mathcal{K} * \phi)$ is inconsistent, only if $\vdash \neg\phi$
- ($\mathcal{K} * 6$) If $\phi \equiv \psi$ then $Bel(\mathcal{K} * \phi) \equiv Bel(\mathcal{K} * \psi)$
- ($\mathcal{K} * 7$) $Bel(\mathcal{K} * (\phi \wedge \psi)) \subseteq Bel(\mathcal{K} * \phi) + \psi$
- ($\mathcal{K} * 8$) If $\neg\psi \notin Bel(\mathcal{K} * \phi)$ then
 $Bel(\mathcal{K} * \phi) + \psi \subseteq Bel(\mathcal{K} * (\phi \wedge \psi))$

We will call a revision operator an *AGM revision operator* if it satisfies the reformulated AGM postulates. Katsuno and Mendelzon [1991] have shown that a necessary and sufficient condition for constructing an AGM revision operator is that a belief state $\mathcal{K}$ can induce, as its preferential information, a total preorder on the set of possible worlds.

Definition 1 A faithful assignment is a function that maps each belief state $\mathcal{K}$ to a total preorder $\preceq_{\mathcal{K}}$ on $\mathcal{M}$ such that for any possible worlds w_1, w_2 :

1. If $w_1, w_2 \models Bel(\mathcal{K})$ then $w_1 =_{\mathcal{K}} w_2$
2. If $w_1 \models Bel(\mathcal{K})$ and $w_2 \not\models Bel(\mathcal{K})$, then $w_1 \prec_{\mathcal{K}} w_2$

The resulting total preorder is referred to as the *faithful ranking corresponding to, or induced by* $\mathcal{K}$. Intuitively, $w_1 \preceq_{\mathcal{K}} w_2$ if w_1 is at least as plausible as w_2 .

It follows from the results of [Katsuno and Mendelzon, 1991] that a revision operator $*$ satisfies ($\mathcal{K} * 1$)–($\mathcal{K} * 8$) iff there exists a faithful assignment that maps $\mathcal{K}$ to the faithful ranking $\preceq_{\mathcal{K}}$ such that for any sentence ϕ :²

$$Bel(\mathcal{K} * \phi) = \mathcal{T}(\min(Mod(\phi), \preceq_{\mathcal{K}})). \quad (1)$$

²Katsuno and Mendelzon deal with formulas instead of belief sets. Since the language is finite, this difference is immaterial. As well, their postulate set is more compact. We use the given definitions in order to adhere more closely to the original AGM approach.

2.3 Belief Change via Sets of Formulas

The idea of changing an agent's beliefs with respect to a set of formulas isn't new. [Fuhrmann and Hansson, 1994] surveys *multiple contraction*, and proposes *package contraction* for removing a set of formulas from a belief set. Similarly, the contraction introduced by [Zhang *et al.*, 1997] studies how to contract a belief set so that it is consistent with a set of formulas, while [Fermé *et al.*, 2003] examines a construction for multiple contraction in a belief base. *Set revision* (also called *multiple revision*) is developed in [Zhang and Foo, 2001; Peppas, 2004], although these papers primarily study infinite sets. A postulate set may be given as follows, analogous to the postulate set in [Peppas, 2004] but adapted for belief states; we refer to this set as the (AGM) *set revision postulates*. We use $\otimes$ to denote a revision operator, where $\otimes$ maps a belief state $\mathcal{K}$ and a finite nonempty set of formulas S to a revised belief state $\mathcal{K} \otimes S$.

- ($\mathcal{K} \otimes 1$) $Cn(Bel(\mathcal{K} \otimes S)) = Bel(\mathcal{K} \otimes S)$
- ($\mathcal{K} \otimes 2$) $S \subseteq Bel(\mathcal{K} \otimes S)$
- ($\mathcal{K} \otimes 3$) $Bel(\mathcal{K} \otimes S) \subseteq Bel(\mathcal{K}) + S$
- ($\mathcal{K} \otimes 4$) If $Bel(\mathcal{K}) \cup S$ is consistent, then $Bel(\mathcal{K}) + S \subseteq Bel(\mathcal{K} \otimes S)$
- ($\mathcal{K} \otimes 5$) $Bel(\mathcal{K} \otimes S)$ is inconsistent only if S is inconsistent.
- ($\mathcal{K} \otimes 6$) If $S_1 \equiv S_2$, then $Bel(\mathcal{K} \otimes S_1) = Bel(\mathcal{K} \otimes S_2)$
- ($\mathcal{K} \otimes 7$) $Bel(\mathcal{K} \otimes (S_1 \cup S_2)) \subseteq Bel(\mathcal{K} \otimes S_1) + S_2$
- ($\mathcal{K} \otimes 8$) If $Bel(\mathcal{K} \otimes S_1) \cup S_2$ is consistent, then $Bel(\mathcal{K} \otimes S_1) + S_2 \subseteq Bel(\mathcal{K} \otimes (S_1 \cup S_2))$

None of these approaches are concerned with relations among subsets of formulas after revision. An approach that does deal with this issue is so-called *parallel revision* [Delgrande and Jin, 2008]. The idea is that in revising by a set of formulas S , if some members of S are subsequently found to be false, the remaining elements of S (where logically possible) will remain believed to be true. Thus, assuming that a , b , and c are logically independent, one would obtain that $a \in Bel((\mathcal{K} * \{a, b, c\}) * \{\neg c\})$. The set revision postulates are extended by the following postulate:³

- ($\mathcal{K} \otimes PP$) Let $S_1 \subseteq S$ where $S_1 \cup (\overline{S \setminus S_1}) \not\models \perp$. Then $Bel(\mathcal{K} \otimes S \otimes (\overline{S \setminus S_1})) = Bel(\mathcal{K} \otimes (S_1 \cup (\overline{S \setminus S_1})))$.

($\mathcal{K} \otimes PP$) expresses the condition that after revising by a set of formulas, and then by the negations of some of those formulas, the remaining formulas, where possible, will persist in being believed to be true. The following condition on a faithful ranking is provided, analogous to ($\mathcal{K} \otimes PP$).

- (PP) Let $S_1 \subseteq S$ where $S_1 \cup (\overline{S \setminus S_1}) \not\models \perp$. Then $\min(Mod(S_1 \cup (\overline{S \setminus S_1})), \preceq_{\mathcal{K}}) = \min(Mod(\overline{S \setminus S_1}), \preceq_{\mathcal{K} \otimes S})$.

A representation result is provided:

Theorem 1 A revision operator $\otimes$ satisfies ($\mathcal{K} \otimes 1$) – ($\mathcal{K} \otimes 8$), and ($\mathcal{K} \otimes PP$) iff there is a faithful assignment whose

corresponding faithful ranking $\preceq_{\mathcal{K}}$ satisfies (PP), and where for any finite set of sentences S :

$$Bel(\mathcal{K} \otimes S) = \mathcal{T}(\min(Mod(S), \preceq_{\mathcal{K}}))$$

In the next section, after we present our approach for revising by a possibly-inconsistent set of formulas, we suggest that this approach augmented by the conditions for parallel revision ($(\mathcal{K} \otimes PP)/(PP)$) constitute a suitable basic approach for revising by a set of formulas.

Finally, in set revision, revision by $\phi \wedge \psi$ is treated differently from revision by $\{\phi, \psi\}$. This distinction between a conjunction and its set of conjuncts has been explored in [Konieczny *et al.*, 2005], where *comma* is treated an additional propositional connective. However, their goals (to study this extension of propositional logic) are quite different from our's here.

3 Revising by an Inconsistent Set of Formulas

3.1 Intuitions

In any extant approach to revision by a set of formulas, as with the AGM approach, if the set for revision is inconsistent then the resulting belief set is also inconsistent. This is due to the fact that in a revision $\mathcal{K} * S$, the success postulate ($\mathcal{K} \otimes 2$) requires that *every* element of S be believed. Informally however it seems that, if a set S is inconsistent, then one still may be able to do more with respect to revision.

Consider again the assumptions underlying AGM belief revision that we discussed. Our interpretation of *success*, as given in ($\mathcal{K} * 2$), is that in a revision $\mathcal{K} * \phi$ the agent has elected to accept ϕ . The formula ϕ is some report, observation, or similar piece of information about the domain. In the set-based approach, the agent has received several such pieces of information, and the agent is presumably prepared to accept any of these individual items of information. However, if these items conflict, then a truly rational agent will not accept all items (and so fall into inconsistency) but rather will treat the items as *candidates* for inclusion in its belief set. If we apply an analogue of the principle of *informational economy* to the set of input formulas S , then an agent will in some fashion incorporate a maximal number of such candidate items of information into its belief set. Consider the following example:

Example 1

$$\begin{aligned} S_1 &= \{a \wedge b, a \wedge \neg b\} \\ S_2 &= \{c, b, \neg b\} \\ S_3 &= \{a \wedge b, a \wedge \neg b, \neg a \wedge b\} \end{aligned}$$

Each set is inconsistent, but intuitively the sets carry different information. In the first case, the agent has received two reports, that $a \wedge b$ and $a \wedge \neg b$. Clearly both formulas cannot be simultaneously and consistently accepted. Nor is it rational to accept both formulas. This might suggest that one *merge* the input formulas in some fashion (as for example in [Konieczny and Pino Pérez, 2002]). However we suggest that such a step is not necessary, nor even appropriate: If the agent had any information about how to select among input formulas, then this would be reflected in its epistemic state. Thus the agent

³[Delgrande and Jin, 2008] also had a postulate to deal with revision by the empty set. Here we simply exclude this case.

already has sufficient resources to determine what aspects of S_1 to incorporate: For formulas $a \wedge b$ and $a \wedge \neg b$, the faithful ranking $\preceq_K$ associated with the agent's epistemic state indicates which formula is more plausible, according to the agent. Thus if the minimal $a \wedge b$ worlds were ranked below the minimal $a \wedge \neg b$ worlds, then this would provide a solid rationale for the revision to be characterised by the minimal $a \wedge b$ worlds. And if the minimal $a \wedge b$ and $a \wedge \neg b$ worlds were at the same rank, then the agent's revised beliefs would be characterised by the union of these worlds.

Similarly with S_2 : Intuitively there is no reason to not incorporate c into the agent's belief corpus, and while it cannot incorporate both b and $\neg b$, it seems that it should accept the more plausible of b , $\neg b$. This suggests that if the least $\{b, c\}$ worlds are more plausible than the least $\{\neg b, c\}$ worlds, then these worlds should characterise the revision. If the least $\{b, c\}$ worlds are as plausible as the least $\{\neg b, c\}$ worlds, then revision would be characterised by their union. In either case, the result would be a consistent belief set in which c is believed. This line of argument has the consequence that for revision by a set S , one should consider the maximal consistent subsets of S as candidate sets for revision, and then select the most plausible worlds among these various subsets to characterise the revision. Thus, in the case of S_3 , which of the candidate formulas will be accepted depends on how the agent ranks their plausibility. It may be that $a \wedge b$ is believed, or a is believed and b is not. As well, if $\neg a$ is believed then so is b . In any case, in the resulting belief state, at least $a \vee b$ will be believed.

3.2 The Approach

We start by reviewing the definition of the set of maximum consistent subsets of a set of formulas.

Definition 2 Let S be a nonempty set of formulas. $Con(S)$, the set of maximum consistent subsets of S , is given as follows. $S_1 \in Con(S)$ iff

1. $S_1 \subseteq S$.
2. $S_1 \not\vdash \perp$.
3. For any S_2 where $S_1 \subset S_2 \subseteq S$, we have $S_2 \vdash \perp$.

The following results are straightforward.

Proposition 1

1. $Con(S) = \{S\}$ iff $S \not\vdash \perp$.
2. $Con(S) = \{\emptyset\}$ iff $\forall \phi \in S, \phi$ is inconsistent.
3. Let $S_1 \in Con(S)$. Then $S_1 \cup (\overline{S \setminus S_1}) \not\vdash \perp$.

We next characterise set-based revision. Consider the following set of postulates which extend the AGM set postulates $(K \otimes 1)$ – $(K \otimes 8)$ to allow for inconsistent sets. The numbering reflects that of the AGM set postulates; for example $(K \otimes 1)$ is the first AGM set revision postulate while $(K \otimes 2')$ is a variation on $(K \otimes 2)$. We will refer to the following as the *extended set revision postulates*.

Postulates:

- $(K \otimes 1)$ $Cn(Bel(K \otimes S)) = Bel(K \otimes S)$
- $(K \otimes 2')$ $\bigvee Con(S) \in Bel(K \otimes S)$

$$(K \otimes 3') \quad Bel(K \otimes S) \subseteq \bigcap_{S' \in Con(S)} (Bel(K) + S')$$

$(K \otimes 4')$ If $Bel(K) + \bigvee Con(S)$ is consistent, then

$$\bigcap_{S' \in Con(S)} (Bel(K) + S') \subseteq Bel(K \otimes S)$$

$(K \otimes 5')$ $Bel(K \otimes S)$ is inconsistent only if S is strongly inconsistent.

$(K \otimes 6)$ If $S_1 \not\vdash \perp$ and $S_1 \equiv S_2$, then

$$Bel(K \otimes S_1) = Bel(K \otimes S_2)$$

$(K \otimes 6')$ If S_1 and S_2 are strongly equivalent, then

$$Bel(K \otimes S_1) = Bel(K \otimes S_2)$$

$(K \otimes 7')$ If $S_1 \not\vdash \perp$ then

$$Bel(K \otimes (S_1 \cup S_2)) \subseteq Bel(K \otimes S_1) + S_2$$

$(K \otimes 8')$ If $S_1 \not\vdash \perp$ and $Bel(K \otimes S_1) \cup S_2$ is consistent, then

$$Bel(K \otimes S_1) + S_2 \subseteq Bel(K \otimes (S_1 \cup S_2))$$

With regards $(K \otimes 2')$, consider the following example.

Example 2

$$\begin{aligned} c &\in Bel(K * \{a, b, \neg b, c\}). \\ c \vee d &\in Bel(K * \{a \wedge b \wedge c, a \wedge \neg b \wedge d\}). \end{aligned}$$

Formulas entailed by all maximum consistent subsets of S (e.g. c and $c \vee d$ respectively) are believed following revision. As well, a is believed following each revision in the example.

A consequence of $(K \otimes 3')$, $(K \otimes 4')$ is the following:

Proposition 2

Let $S = \{S' \in Con(S) \mid Bel(K) + S' \text{ is consistent}\}$.

If $S \neq \emptyset$ then $Bel(K \otimes S) = \bigcap_{S' \in S} (Bel(K) + S')$.

Example 3 If $Bel(K) = Cn(a)$ and $S = \{a \wedge b, \neg a \wedge \neg b\}$ then $Bel(K \otimes S) = Cn(a \wedge b)$.

$(K \otimes 5')$ deals with the situation in which every formula in S is inconsistent. The decision that $Bel(K * S)$ is also inconsistent in this case is arbitrary; one could equally well argue that a rational agent would reject incorporating any inconsistent information, on the grounds that it is impossible. Hence by this argument the belief set should remain unchanged. On the other hand, one might decide that $Bel(K * S)$ should be inconsistent, since this is closer to the original AGM approach. For this reason, and because it makes the representation result tidier, we elect to have $Bel(K * S)$ be inconsistent when S is strongly inconsistent.

It can be noted that $(K \otimes 6)$ and $(K \otimes 6')$ are independent.

Example 4

We have that $Bel(K \otimes \{a, b\}) \equiv Bel(K \otimes \{a \wedge b\})$ by $(K \otimes 6)$, while $(K \otimes 6')$ has nothing to say in this case.

$(K \otimes 6')$ asserts that $Bel(K \otimes \{a \wedge b, \neg a\}) \equiv Bel(K \otimes \{a \wedge (a \supset b), \neg a\})$ while $(K \otimes 6)$ is inapplicable.

The proviso $S_1 \not\vdash \perp$ for $(K \otimes 7')$ and $(K \otimes 8')$ basically makes these postulates the same as their counterparts $(K \otimes 7)$ and $(K \otimes 8)$.

We also obtain the following factoring result for maximum consistent sets:

Proposition 3

$(\mathcal{K} \otimes F)$ Let $Con(S) = \{S_1, S_2\} \neq \emptyset$. Then

$$Bel(\mathcal{K} * S) = \begin{cases} Bel(\mathcal{K} * S_1) & \text{or} \\ Bel(\mathcal{K} * S_2) & \text{or} \\ Bel(\mathcal{K} * S_1) \cap Bel(\mathcal{K} * S_2). \end{cases}$$

The case where $S \equiv S_1 \equiv S_2$ is trivial. Otherwise S is inconsistent, but there are two maximum consistent subsets. Factoring states that revision is a function of these two maximum consistent subsets. This result also generalises via a straightforward induction:

Proposition 4

$(\mathcal{K} \otimes F^n)$ Let $Con(S) = \{S_1, \dots, S_n\} \neq \emptyset$. Then for some $\mathbf{S}$ where $\mathbf{S} \subseteq Con(S)$,

$$Bel(\mathcal{K} * S) = \bigcap_{S' \in \mathbf{S}} Bel(\mathcal{K} * S').$$

This postulate set can be tied to faithful assignments via an extension of the Katsuno and Mendelzon [1991] result:

Theorem 2 A revision operator $\otimes$ satisfies the extended set postulates iff there exists a faithful assignment that maps $\mathcal{K}$ to the faithful ranking $\preceq_{\mathcal{K}}$ such that for any finite set of formulas S :

$$Bel(\mathcal{K} \otimes S) = \mathcal{T}(\min(\text{Mod}(\bigvee Con(S)), \preceq_{\mathcal{K}})).$$

Finally, we can combine this approach with that of [Delgrande and Jin, 2008] to obtain a suggested “preferred” basic approach to revising by a set of formulas:

Theorem 3 A revision operator $\otimes$ satisfies the extended set postulates together with $(\mathcal{K} \otimes PP)$ iff there is a faithful assignment whose corresponding faithful ranking $\preceq_{\mathcal{K}}$ satisfies (PP) , and where for any finite set of sentences S :

$$Bel(\mathcal{K} \otimes S) = \mathcal{T}(\min(\text{Mod}(\bigvee Con(S)), \preceq_{\mathcal{K}})).$$

This in turn yields a revision operator on sets of formulas in which inconsistency of a set of formulas is handled appropriately and, for future revisions, elements of the set are retained where possible.

We also obtain the result:

Proposition 5

If $S' \in Con(S)$ then $Bel(\mathcal{K} \otimes S \otimes S') = Bel(\mathcal{K} \otimes S')$.

So revising by a set of formulas and then by one of the maximum consistent subsets gives the same belief set as revising by that maximum consistent subset.

Consider the following examples, where $\otimes$ is a revision operator as given in Theorem 3.

Example 5

$$\begin{aligned} a, c &\in Bel(\mathcal{K} \otimes \{a, b, c\} \otimes \{\neg b\}) \\ a &\in Bel(\mathcal{K} \otimes \{a, b, c, \neg c\} \otimes \{\neg b\}) \\ a &\in Bel(\mathcal{K} \otimes \{a \wedge b, a \wedge \neg b\} \otimes \{\neg b\}) \end{aligned}$$

The first part of the example illustrates the basic approach in parallel revision. The second part shows that this extends to the case of inconsistent sets. The last part is interesting, in that there are revision operators in which $a \wedge b \in Bel(\mathcal{K} \otimes$

$\{a \wedge b, a \wedge \neg b\})$; however, even in this case, in subsequently revising by $\neg b$, a is retained. On the other hand, there are revision functions where $a \notin Bel(\mathcal{K} \otimes \{a \wedge b, \neg b\})$, and similarly in subsequently revising by $\neg b$.

As noted in [Delgrande and Jin, 2008], the basic approach to parallel revision doesn’t fully address iterated revision. In fact [Delgrande and Jin, 2008] augments the basic approach with that of [Jin and Thielscher, 2007] to obtain an approach to parallel revision that also handles iterated revision appropriately. While one could do the same here, with this or another approach to iterated revision, we suggest that in fact that such an iterated revision operator would play only a minor role with respect to maintaining an agent’s corpus of beliefs. We develop this thesis in the next section.

4 Belief Change and the Role of Iterated Revision

Consider again assumptions underlying classical AGM revision, and how they are interpreted in the present framework. Most pertinently, AGM revision assumes that:

1. The underlying domain (or: domain language) is static.
2. The agent receives a sequence of input formulas for revision about this (static) domain.
3. Each formula for revision is accepted.
4. Implicitly all input formulas have equal weight. This is a consequence of the success postulate; for example for contingent ϕ , one will never have $\phi \in \mathcal{K} * \phi * \neg\phi$.

In the present approach, we also assume the first two points. However, then things differ: if a domain is static, and there is a sequence of formulas that an agent has (tentatively) decided to accept, there is no reason that the most recent *must* be accepted. Thus, if the order in which formulas are accepted is irrelevant, then the input to revision *should* be the set of such formulas. This set then consists of formulas that individually the agent would accept. Since it is not rational to accept all formulas if the set is inconsistent, and since an agent would want to accept a maximal amount of information, the set of maximal consistent subsets of the input set constitute candidate sets for revision. Furthermore, the agent has sufficient information to adjudicate among these maximum consistent subsets, as given by its associated faithful ranking, which assigns a plausibility level to every formula.

We emphasize this difference in approaches since it leads to a very different view of so-called iterated belief revision: a revision sequence $\mathcal{K} * \phi_1 * \dots * \phi_n$ is more appropriately regarded as the uniterated revision $\mathcal{K} \otimes \{\phi_1, \dots, \phi_n\}$. In theory then there is no need for iterated revision as understood in the belief change community.

So does this mean that there is no role for iterated revision? Arguably, there is a role, but it is a pragmatic one. In the set revision approach, an agent will need to successively recompute a new belief state based on an increasing set of input formulas. That is, if the agent has computed $\mathcal{K} \otimes S$, and then receives a new item of information ϕ , it will need to compute $\mathcal{K} \otimes (S \cup \{\phi\})$. Clearly, the process of computing a revision over all input formulas each time a new formula is received

will be expensive. Thus from a practical point of view, after a sufficient number of input formulas have been received, one may elect to “compile” these formulas into a new belief state, and so adopt a state $\mathcal{K}'$ given by $\mathcal{K} \otimes S$. Then further inputs would again be collected into another set, and revision computed with respect to $\mathcal{K}'$. The advantage of the extended approach, as given in Theorem 3, is that for $(\mathcal{K} \otimes S) \otimes S'$, elements of S will be retained in the subsequent revision by S' unless specifically contradicted by S' .

5 Conclusion

This paper has presented an approach to belief revision in which consequences of the assumption of an underlying static domain are explored. In particular, we argue that the success postulate is too strong when there is more than one formula for revision. Instead we suggest that revision should be set-based, and that an agent will revise its corpus of beliefs by the set of input formulas that it has received. Given that the agent’s full knowledge about the domain is encoded in its epistemic state, then this along with a principle of informational economy has the consequence that for revision by a set S , one should consider the maximal consistent subsets of S as candidate sets for revision, and then select the most plausible worlds among these various subsets to characterise the revision. This suggests in turn that “one shot” AGM set-based revision is in fact of primary interest, and that iterated revision plays a relatively minor role.

In the approach, input formulas are effectively given the same weight, in that they are all candidates for acceptance. However, in general, formulas may come with different degrees of reliability; for example, an agent may have information that a report of ϕ is more likely to be correct than one of ψ . This notion of reliability is distinct from the agent’s plausibility ordering on formulas. Hence an interesting direction for future research concerns dealing with formulas with different degrees of reliability, and reconciling such an implicit ordering on observations with the plausibility ordering implicit in the agent’s ranking function on worlds.

References

- [Alchourrón *et al.*, 1985] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence* 89, 1–29, 1997.
- [Delgrande and Jin, 2008] J. Delgrande and Y. Jin. Parallel belief revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 317–322, 2008.
- [Fermé *et al.*, 2003] Eduardo L. Fermé, Karina Saez, and Pablo Sanz. Multiple kernel contraction. *Studia Logica*, 73(2):183–195, 2003.
- [Friedman and Halpern, 1999] N. Friedman and J. Halpern. Belief revision: A critique. *Journal of Logic, Language, and Information*, 8(4):401–420, 1999.
- [Fuhrmann and Hansson, 1994] A. Fuhrmann and S.O. Hansson. A survey of multiple contractions. *Journal of Logic, Language, and Information*, 3:39–76, 1994.
- [Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
- [Jin and Thielscher, 2007] Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [Katsuno and Mendelzon, 1991] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.
- [Konieczny and Pino Pérez, 2002] S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
- [Konieczny *et al.*, 2005] S. Konieczny, J. Lang, and P. Marquis. Reasoning under inconsistency: The forgotten connective. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 484–489, Edinburgh, 2005.
- [Lang, 2006] J. Lang. About time, revision, and update. In J. Dix and A. Hunter, editors, *Proceedings of the Eleventh International Workshop on Non-Monotonic Reasoning (NMR 2006)*, 2006.
- [Nayak *et al.*, 2003] Abhaya C. Nayak, Maurice Pagnucco, and Pavlos Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.
- [Peppas, 2004] Pavlos Peppas. The limit assumption and multiple revision. *Journal of Logic and Computation*, 14(3):355–371, 2004.
- [Zhang and Foo, 2001] D. Zhang and N. Foo. Infinitary belief revision. *Journal of Philosophical Logic*, 30(6):525–570, 2001.
- [Zhang *et al.*, 1997] Dongmo Zhang, Shifu Chen, Wujia Zhu, and H. Li. Nonmonotonic reasoning and multiple belief revision. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 95–100, Nagoya, Japan, August 1997. Morgan Kaufmann Publishers.